Title
Data-Driven Decision Making for Last-Mile Delivery and Online Platform Operations

Author
Cao, Junyu

Publication Date
2020

Peer reviewed|Thesis/dissertation

Data-Driven Decision Making for Last-Mile Delivery and Online Platform Operations

by

Junyu Cao

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Zuo-Jun (Max) Shen, Chair
Professor Mariana Olvera-Cravioto, Co-chair
Professor Rhonda Righter
Professor Terry Taylor

Spring 2020

Data-Driven Decision Making for Last-Mile Delivery and Online Platform Operations

Abstract

Data-Driven Decision Making for Last-Mile Delivery and Online Platform Operations

by

Junyu Cao

Doctor of Philosophy in Engineering – Industrial Engineering and Operations Research

University of California, Berkeley

Professor Zuo-Jun (Max) Shen, Chair

Professor Mariana Olvera-Cravioto, Co-chair

This dissertation focuses on two main areas: 1) data-driven stochastic modeling and applied probability, with applications to the sharing economy and networks; 2) machine learning, with a focus on sequential decision making for recommender systems and revenue management. Using tools from probability, statistics, and learning theory, this dissertation emphasizes fundamental contributions to both theory and methodology.

Chapter 2 focuses on the design of stochastic models and mechanisms to increase efficiency of urban transportation and logistics systems. Trends in global urbanization require innovation, especially in the area of urban transportation networks. At the same time, the sharing economy provides opportunities to enhance efficient resource utilization. This Chapter addresses the new operational challenges that arise from emerging technologies within smart cities, focusing on last-mile delivery and smart mobility. Last-mile delivery may take up to 28% of the total transportation costs, and is arguably one of the biggest challenges in logistics management. We propose a new business model for optimizing the last-mile delivery of packages, using a strategy that combines the use of ride-sharing platforms (e.g. Uber or Lyft) with traditional in-house van delivery systems. To make the proposed solution tractable, we develop new theoretical results by approaching the problem from a probabilistic perspective. Our approach of determining the optimal reward to private drivers for delivering packages is computationally efficient. Using synthetic and real data, we show that our approach reduces cost by as much as 30% compared with the van-only strategy.

In Chapter 3, motivated by the observation that overexposure to unwanted marketing activities can lead to customer dissatisfaction, we consider a setting where a platform offers a sequence of messages to its users and is penalized when users abandon the platform due to marketing fatigue. We propose a novel sequential choice model to capture multiple interactions taking place between the platform and its users: upon receiving a message, a user

decides on whether to accept or reject the message. If she chooses to reject, she would then decide to either receive the next message in the sequence or abandon the platform. Based on user feedback, the platform dynamically learns users' abandonment distribution and the relevance of the recommended content. With a goal to maximize the cumulative payoff over a time horizon, the platform dynamically adjusts the sequence of messages and the order in which the messages are shown to a user. We refer to this online learning task as the sequential choice bandit (SC-Bandit) problem. For the offline combinatorial optimization problem, we show a polynomial-time algorithm. For the online problem, we consider two variants, depending on whether contexts are included, and propose algorithms that balance exploration and exploitation. Lastly, we evaluate the performance of our algorithms with both synthetic and real-world datasets.

Complex networks appear in essentially all branches of science and engineering, and people from various fields have used random graphs to model, explain and predict some of the properties commonly observed in real-world networks. In Chapter 4, we study a family of directed random graphs whose arcs are sampled independently of each other, and are present in the graph with a probability that depends on the attributes of the vertices involved. In particular, this family of models includes as special cases the directed versions of the Erdös-Rényi model, graphs with given expected degrees, the generalized random graph, and the Poissonian random graph. We establish the phase transition for the existence of a giant strongly connected component and provide some other basic properties, including the limiting joint distribution of the degrees and the mean number of arcs. In particular, we show that by choosing the joint distribution of the vertex attributes according to a multivariate regularly varying distribution, one can obtain scale-free graphs with arbitrary in-degree/outdegree dependence.

To Lizhong and Xihong

Contents

List of Figures

List of Tables

Acknowledgments

I am fortunate to be co-advised by Professor Zuo-Jun (Max) Shen and Professor Mariana Olvera-Cravioto. They introduced me to the field of Operations Research and guided me to do research. The advice and tremendous support I obtained from them would impact my entire research life. Their encouragement leads me to the academic road. I would like to gratefully thank my committee member Professor Rhonda Righter for her incredible mentorship both on my research and teaching. Through many conversations with her, I am able to improve my presentation skills and critical thinking abilities. I want to thank Professor Terry Taylor, Professor Candace Yano, and Professor Paul Grigas for valuable suggestions on my research. I am also very lucky to be advised by Professor Van M. Savage during my undergraduate years. He provided enormous help on writing my first research paper.

I appreciate the internship opportunity at Thomas J. Watson Research Center for two summers and the financial support from IBM Ph.D. Fellowship. I would like to thank my mentor Wei Sun and my team manager Markus Ettl for the great help on my internship projects. I also have many fun memories with my friends Xiufan Yu, Hao Wang, and Lam Nguyen at IBM, just to name a few.

I feel very fortunate meeting amazing people at Berkeley. I treasure the time spent with Mengyu Li and Danqing Zhang for fun conversations and appreciate their great support standing behind me. I also spent memorable times with Sheng Liu, Wei Qi, Ying Cao, Tianyi Lin, Auyon Siddiq, Hansheng Jiang, Xu Rao, Haoyang Cao, Mengxin Wang, Quico Spaen, Salar Fattahi, Ruijie Zhou, Renyuan Xu, Heyuan Liu, Anran Hu, Alfonso Lobos, Pedro Ivo Bastos Hespanhol, Georgios Patsakis, Nan Yang, Erik Bertelli, Han Feng, Marie Pelagie Elimbi Moudio, Yuhao Ding.

Lastly, I want to thank my parents and grandparents for supporting me all the time. Their endless love is my Infinity Stones Power Ring – I am not afraid anymore.

Chapter 1

Last-mile shared delivery: A discrete sequential packing approach

1.1 Introduction.

Last-mile delivery, the last step of the delivery process from a distribution center to its final destination, is arguably one of the biggest challenges in logistics management, and it may take up to 28% of the total transportation costs [97]. Therefore, minimizing the cost of this last step of the order fulfillment process is of great importance. Historically, the delivery of items/packages to their final recipients from a centrally located warehouse has been done by either investing in a fleet of vans or trucks to operate regionally, or by outsourcing it to third-party logistics companies specializing in such services. However, today's popular "sharing economy" has provided a new alternative that can take advantage of ride-sharing platforms to offset some of the last-mile delivery costs. In this paper, we propose a model for analyzing a delivery process that combines the use of private drivers (e.g., Uber or Lyft drivers) with a more traditional van delivery system, with the idea of reducing the need to invest in and maintain a large fleet of delivery vans or trucks.

We study here a problem where a warehouse or distribution center has to deliver a large number of packages on a given day. Traditionally, this would be done by deploying a fleet of vans along many different routes, with the routes being computed efficiently using some version of the capacitated vehicle routing problem (CVRP). One of the main features to focus on regarding this approach is the number of vans that are needed to deliver all the packages during the allotted time. Instead, we propose a framework to encourage private drivers to pick up bundles of packages from the warehouse and deliver them during time interval $[0, T]$, with any remaining packages at time T being delivered by the warehouse's van system. We can think of the time threshold T as the end of the day, after which any undelivered packages will be loaded into vans for delivery the next day. Alternatively, we can envision T as representing a time in the early afternoon that can allow vans to deliver packages during a late afternoon shift (e.g. UPS and Amazon vans deliver packages until around 8 pm). The purpose of setting up a time threshold T is to guarantee a promised delivery time, and its use is common in practice. For example, UPS Next and UPS Ground packages must be received prior to 4:00 pm[1], and any packages arriving after 4:00 pm will be considered to be part of next day's packages. In general, we believe the length of the period $[0, T]$ is dictated by external circumstances to the warehouse (e.g., a promised delivery time like the ones UPS and Amazon offer, the work schedules of van drivers, or the boundaries of "peak" and "off-peak" hours in a multiperiod implementation). In the latter case, the fleet of vans would be delivering priority packages, packages with tight schedules, or oversize items, prior to time T, which would be unsuitable for private drivers anyways. The key difference between our proposed approach compared to the traditional van-only one, is that the number of vans needed will be considerably smaller, which translates into important savings[2] for the warehouse.

The main challenges of using a mixed strategy are: 1) the modeling of the behavior of private drivers, and 2) the design of a payment scheme that will incentivize private drivers to pick up packages from the warehouse efficiently. The problem of using shared-mobility to deliver packages has been studied in [54] and [67], however, the approaches proposed by the authors do not scale well with the size of the problem (i.e., the number of packages that need to be delivered). Our framework provides an easily computable approach to determine whether using shared mobility for delivering packages in addition to traditional van deliveries is desirable, and if so, to choose the optimal payment scheme to offer private drivers for delivering each package. The main modeling contribution of this paper lies precisely on the computation of the latter. The key idea is that the number of packages that can be picked up by private drivers can be modeled in a way that is independent of their destinations, provided that all packages are equally desirable to the drivers. To achieve this, we design a payment scheme that makes the profit for delivering packages the same across packages, and we control this profit through a common "incentive rate". Once the expected number of packages that can be picked up during $[0, T]$ is computed (as a function of the incentive rate) we find the optimal incentive rate by solving a single-variable continuous optimization problem. The daily computations are done in two phases: one at the beginning of the day that ends with the assignment of the payment rewards for each individual package using the optimal incentive rate, and another one at time T that computes an optimal solution to the CVRP for the leftover packages. Our proposed approach keeps the computational cost low, and the framework can be seen as a first step towards an online implementation that can react in real time to the randomness in the supply of private drivers.

The mathematical contributions of the paper are centered around the computation of the expected number of packages that can be picked up by private drivers during the time period $[0, T]$, which is done both exactly (via a recursive computation) and asymptotically (as the number of packages grows to infinity). The exact computation is used in the formulation of the optimization problem that we will solve, while the asymptotic result is needed to ensure that the computational complexity of the optimization problem remains small for large numbers of packages. Our methodology is closely related to the analysis of Rényi's classical "parking/packing problem" [91], which studies a model where unit-length "cars" arrive sequentially to random points of a "sidewalk" of fixed length and park if they can "fit", with no parked cars ever leaving. The main question studied in [91] is the distribution of the unoccupied space once no more cars can fit, i.e., when there are no remaining unoccupied subintervals of length one. In the context of our package delivery process, we consider a discrete version of the problem where the cars are replaced by bundles of packages and parking attempts are replaced by bundle requests from the private drivers. Since we need to model the randomness in the times at which private drivers' requests arrive, we incorporate a time element to our formulation in the spirit of the packing problem considered in [30], where the cars/packages (requests for bundles of packages in our case) arrive according to a Poisson process. Our theoretical results are therefore of independent interest within the literature on Rényi's parking/packing problem, and include a novel convergence rate theorem for the expected fraction of packages that can be picked-up by time T that was previously unknown even for special cases of our model.

To better tie our technical results with the more applied optimization problem we propose to solve, it is worth mentioning that the scale of the parcel delivery industry is very large. For example, Holland et al. [52] described the UPS delivery system, which operates about 1,400 package distribution centers in the United States. On a typical day, more than 16 million packages need to be delivered, which indicates an average of 11,429 packages per distribution center per day. In 2015, UPS delivered 34 million packages on its peak day. It follows that in our problem formulation, we could be working with over 10,000 packages per day, which is a large enough number to make the exact computation of the expected fraction of packages that can be picked up during $[0, T]$ computationally intensive (since it is done recursively). It follows, that having an asymptotic expression that can replace the exact computation without sacrificing accuracy (a consequence of its fast convergence rate) is extremely valuable to the everyday computation of the optimal payment rate.

The paper is organized as follows. In Section 1.2 we include an overview of the existing literature on last-mile delivery problems in a shared economy. In Section 1.3.1 we introduce a *discrete sequential packing* (DSP) problem that will be used to compute the expected number of packages that can be picked-up from the warehouse by private drivers during the time period $[0, T]$. In Section 1.3.2 we describe a way to estimate the cost of delivering n packages with known destinations $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\} \subseteq \mathbb{R}^2$, using a strategy that combines the use of private drivers and in-house vans. We also include in that section a comparison between our proposed strategy and a more traditional van-only strategy. Section 1.4 contains all the proofs of our theoretical results. Finally, to illustrate our methodology we include, as an online companion to the paper, a numerical experiments section (Section 1.5) that provides some insights into the computational effort needed to implement our strategy and its potential cost benefit.

1.2 Literature Review.

The idea of using shared mobility to deliver packages is still in its infancy, but there are already a few relevant studies worth mentioning. Li et al. [67] proposed a framework to use taxis to transport people and deliver parcels at the same time, where the goal is to determine an optimal pick-up/drop-off sequence. They first consider a static version of the problem where all the people and parcel locations are known and formulate a mixed integer linear program (MILP) to find the optimal pick-up/drop-off sequence; then, they repeatedly solve their static formulation over small periods of time to obtain a more realistic dynamic scenario over a longer time window. In their dynamic approach, parcel locations are revealed at the beginning of the day while passenger locations are updated throughout the day, which means that a full schedule for picking up and delivering parcels cannot be computed upfront. Moreover, the dynamic version requires that the sequence be recomputed every time a passenger request appears.

Qi et al. [88] consider the problem of delivering packages from a warehouse to locations in a given region, and propose to first subdivide the region into a number of sub-regions, use a van system to distribute the packages among the centers of the sub-regions (referred to as "terminals") , and then use shared mobility to do the last-mile delivery within each sub-region. The goal is to determine the optimal size (number) of the sub-regions, which is done using a continuous-approximation for the total average cost, and once the sub-regions are determined, the routes for delivering packages within each sub-region are computed by solving an open vehicle routing problem (OVRP). Note that the optimization problem to identify the regions is meant to be solved only once and the OVRP is solved once per day, unlike the work in [67] where the optimization problem is solved many times every day.

Perhaps the closest to our approach is the work done in Kafle et al. [54], where the authors consider using cyclists and pedestrians as crowdsources to help an in-house van system to do the last-leg delivery of parcels originating from a central warehouse. Their model assumes that parcel destinations are revealed at the beginning of the day, at which point cyclists and pedestrians are expected to submit bids to deliver them; then, once all bids have been received, the warehouse solves a mixed integer non-linear program to compute the optimal assignment, with any unassigned parcels being delivered by the in-house vans. Since solving exactly the mixed integer non-linear program is computationally very expensive, the authors propose a tabu search approximate algorithm instead. Compared to our model, we can think of the approach in [54] as looking at a multi-player game where the cyclists and pedestrians compete with each other with their bids and the warehouse optimizes over the entire set of bids, while in our setup there is only one player, the warehouse, that computes the optimal rewards for each package based on the expected response of the private drivers. In [54] both parcel destinations and bids need to be known at the beginning of the day while in our approach only package destinations need to be revealed while private driver requests remain stochastic. In both [54] and our proposed framework the optimization problem is meant to be solved once every day.

The theoretical contributions of our work are related to Rényi's parking/packing problem [91], which examines the problem of filling up an interval $I = [0, n]$ with subintervals of length one whose left endpoints are uniformly chosen at random, and computes the asymptotic proportion of filled space (as $n \to \infty$). A discretized version of this problem was introduced by Page [84], by considering n points instead of the interval I and replacing the subintervals of unit length with pairs of adjacent points. The main results in [84] include the mean and other moments of the distribution of the remaining isolated points. Page's work was later generalized, both theoretically and numerically, to the case of m adjacent points in [29, 77, 86]. In our context, the number of adjacent points corresponds to the bundle sizes, which are allowed to be random from any finite support discrete distribution; hence, our setting includes as special cases those of the problems studied in [29, 77, 84, 86].

We point out that none of the references mentioned above require the use of a "time" component for the random selection of subintervals/points, since only the distribution of unfilled space is of interest. In this regard, our work is more closely related to that of [30], which assumes that the times at which the subintervals are chosen in Rényi's parking/packing problem occur according to a Poisson process, in which case one can focus on the proportion of filled space by time t. The results in [30] include the asymptotic proportion of filled space by time t, as $n \to \infty$, and a corresponding central limit theorem. We emphasize that the convergence rate result for the proportion of packages that can be delivered in $[0, T]$ (alternatively, proportion of occupied space in the parking problem) as $n \to \infty$ is completely new, and was unknown even for the case of deterministic m in [29, 77, 84, 86].

It is worth reiterating that the methodology we propose to compute the payment (reward) offered to private drivers for each individual package is simple and flexible, since it only requires that we optimize a single-variable cost function at the beginning of the day and it allows private drivers to pick up bundles of packages, i.e., not only one package at a time. As our main results show, this is enough to guarantee a lower total expected cost for the delivery of all packages under very natural cost assumptions. Our modeling approach could also become a building block for constructing improved online pricing models, where the rewards offered to private drivers change throughout the day as more packages are delivered and the allotted time decreases, which would very likely provide even better cost reductions at the expense of more computational effort.

1.3 Model Description.

In the first part of this section we model our last-mile delivery problem using private drivers as a *discrete sequential packing* problem (DSP), and provide both exact and asymptotic results for the expected number of packages that can be delivered during the time window $[0, T]$. In the second part of this section we describe a payment scheme for the private drivers based on this expectation and introduce a joint optimization problem to calculate the optimal incentive rate for minimizing the total expected cost.

1.3.1 The expected number of packages that can be picked.

At the beginning of the day, the destinations of n packages that need to be delivered are revealed. The requirement that package destinations be known at the beginning of the day is standard in today's UPS's operations, since delivery routes need to be computed the previous day and vans need to be packed before drivers can start their routes (see [52]). As mentioned in the introduction, we consider a time window of length T during which private drivers can pick up any of the n packages. The packages are assumed to have destinations spread over a region of a two-dimensional plane. This region is expected to be a population-dense area where package destinations are close to each other.

During the period $[0, T]$ the distribution center receives requests from private drivers to deliver available packages. Any undelivered packages at time T will be delivered by the distribution center's vans/trucks. Moreover, the private drivers are allowed to take a bundle of packages with them, which we assume would have destinations in close proximity to each other. To model this proximity of the destinations, we first arrange all n packages on a circle using the solution (exact or approximate) to the traveling salesman problem (TSP). The idea behind arranging the packages in this way is to reduce the complexity of dealing with the two-dimensional package destinations by reorganizing them into one dimension, and will let us to identify bundles of nearby destinations with segments of the optimal TSP route. Moreover, from a modeling point of view, this simplification allows us to disregard the package destinations in the computation of the expected number of packages that will be picked up by private drivers, and we will argue that we can do this since our pricing mechanism will make the profit rate for delivering each package equal for all packages.

From this point onwards, we can think of the packages as arranged on a circle with n points, where point i will be referred to as the "location" of package i (see Figure 1.1). A bundle of size k at location i is the set consisting of the packages at locations $\{i, i + 1, \ldots, i + k - 1\}$. Each location i is associated to a marked Poisson process with rate λ_i, which represents the arriving requests from private drivers to deliver a bundle whose first package is at location i. The marks of the Poisson process determine the size of the bundle that the driver would want to deliver, and is assumed to be distributed according to some distribution F, independent of the Poisson process and of any other marked Poisson processes at different locations. The arrival behavior of the drivers can be understood as "launching an app", which means the driver need not be physically at the warehouse. These marked Poisson processes model drivers' action of opening an app and using it to select the bundle they wish to pick up, which may occur prior to their actual arrival time to the distribution center. We assume that each driver has a preconceived idea of which package(s) they want to pick up, and that if at the time of the request one or more of the desired packages is no longer available, the driver abandons the idea of picking up packages altogether. Our assumption that drivers arrive at the distribution center with a "fixed set of packages to deliver in mind" is consistent with the behavior of Uber and Lyft drivers in the real world, since they tend to "reject" passengers when they discover where they need to be picked up. In order for all packages to be equally likely to be requested, we will assume that $\lambda_i = \lambda$ for all $1 \leq i \leq n$. Moreover, we assume that λ remains constant during $[0, T]$, which would be the case if the pick-up window is chosen to coincide with private drivers' off-peak hours. Once a bundle request is accepted, the driver is also given the segment of the optimal TSP route for delivering all the packages in the bundle, in case they want to use it.

A numerical analysis comparing our framework using a TSP optimal path to determine

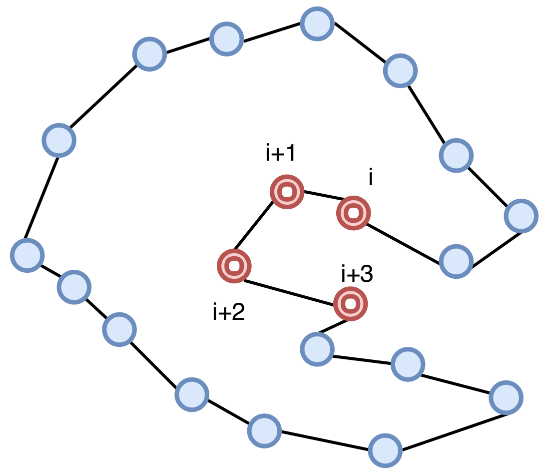

Figure 1.1: TSP tour. The picture illustrates a case with $n = 21$ packages, where the first bundle to be requested is of size four (double circular packages). After the first bundle is picked up, the remaining packages are arranged on a line.

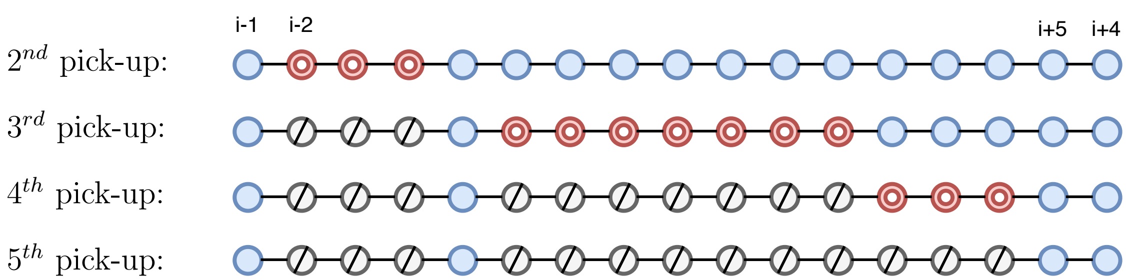

Figure 1.2: Depiction of the package pick-up process. Double circular packages correspond to accepted bundle requests; circular packages are yet to be picked up; circle-with-a-line packages are already taken.

the possible bundles with a more relaxed setting where bundles consist of nearby packages on the 2-dimensional plane (i.e., within a predetermined radius of the location where the request arrives) is given in Section 1.5.4. Our results there show that the difference in cost due to the restriction imposed by our definition of bundle is very small, and worth the analytical simplification it provides.

Throughout the time interval $[0, t]$, the packages at the n locations are picked up by private drivers, and our goal is to analyze the expected number of packages that can be delivered in this way over the interval $[0, t]$, which we denote by $\mathcal{C}(t, n, \lambda)$. Our analysis of $\mathcal{C}(t, n, \lambda)$ is based on the observation that once the first bundle, say of size B, is picked up, the remaining $n - B$ packages can be arranged on a line (see Figure 1.2). In fact, our main results for $\mathcal{C}(t,n,\lambda)$ are obtained by analyzing the expected number of packages that can be picked up during the interval $[0,t]$ when we start with n packages arranged on a line, which we denote by $\mathcal{K}(t,n,\lambda)$. Figures 1.1 and 1.2 show an example of the pick-up process. We refer to this model as a *discrete sequential packing* (DSP) problem since it has strong connections to Rényi's parking/packing problem [91].

The next subsection includes our main results for the computation of $\mathcal{C}(t,n,\lambda)$. We point out that the DSP formulation is based only on the following two assumptions:

i) The marks of the Poisson process at location i, corresponding to bundle size requests and denoted $\{B_j^{(i)}\}_{j\geq 1}$, are i.i.d. with common distribution F for all $1 \leq i \leq n$, and are independent of the Poisson processes at every location.

ii) The arrival rates of the different Poisson processes, which correspond to the arrival rates for requests from drivers to pick up bundles at location i, are all equal, i.e., $\lambda_i = \lambda$ for all $1 \leq i \leq n$.

As mentioned above, the first of these assumptions is justified in our package delivery context by the TSP route, where the distance between adjacent packages along the route are typically small, making all bundles of the same size (approximately) equally desirable. That the distance between adjacent points along an optimal TSP route is small follows from the work in [94], where it is shown that the shortest path through n points independently uniformly distributed over $[0,1]^2$ has the property that the number of edges of length at least $un^{-1/2}$ is at most $Kn\exp(-u^2/K)$ with high probability for some universal constant K. In other words, the neighboring distance in an optimal TSP route is of order $O(n^{-1/2})$ with high probability. The second assumption, which ensures that the arrival rate for requests at all n locations are the same, will be justified by the pricing mechanism described in Section 1.3.2, along with the assumption that our business targets a population-dense area where package destinations are close to each other. In particular, we propose a pricing scheme that incorporates the distance from the distribution center to each destination and the neighboring distance between packages along the TSP route, among other parameters, in such a way that the profit rate for delivering any of the n packages is essentially the same. Since intuitively the total number of packages that need to be delivered and the supply of drivers in a given region are both proportional to its size, we can reasonably assume that λ does not depend on n.

For implementation purposes, we also refer the reader to [7, 32, 51, 66] for further details on the TSP solution and existing algorithmic approaches for its computation [1, 81, 87].

1.3.1.1 Results for the discrete sequential packing problem.

Our main results for $\mathcal{C}(t,n,\lambda)$, the expected number of packages that can be delivered during a time period $[0,t]$, include an exact calculation obtained by solving a differential equation, and its asymptotic behavior as the number of packages goes to infinity. Throughout this section, the arrival rate λ_i at each location is assumed to be the same for all locations, say λ, and the distribution of the bundle sizes, denoted F, is assumed to have support on $\{1, 2, \ldots, m\}$ for some fixed $m \in \mathbb{N}_+$. Define $f(i) = F(i) - F(i-1)$ to be the probability mass function of F.

The first result gives a differential equation satisfied by $\mathcal{K}(t, n, \lambda)$, which we recall is the expected number of packages that can be picked up by time t when n packages are arranged on a line.

Theorem 1 *$\mathcal{K}(t, n, \lambda)$ satisfies the following recursive differential equation*

$$\frac{1}{\lambda}\frac{\partial \mathcal{K}(t,n,\lambda)}{\partial t} = -\sum_{j=1}^{n} F(j)\mathcal{K}(t,n,\lambda) + 2\sum_{i=1}^{n-1} F(n-i)\mathcal{K}(t,i,\lambda) + \sum_{i=1}^{n} f(i)i(n+1-i)$$

with boundary condition $\mathcal{K}(0, n, \lambda) = 0$. Moreover, $\mathcal{R}(t, n, \lambda) := n - \mathcal{K}(t, n, \lambda)$ satisfies the recursive differential equation:

$$\frac{1}{\lambda}\frac{\partial \mathcal{R}(t,n,\lambda)}{\partial t} = -\sum_{i=1}^{n} F(i)\mathcal{R}(t,n,\lambda) + 2\sum_{i=1}^{n-1} F(n-i)\mathcal{R}(t,i,\lambda) \tag{1.3.1}$$

with boundary condition $\mathcal{R}(0, n, \lambda) = n$.

The corresponding solution is given by the following theorem.

Proposition 1.3.1 *The solution to equation* (1.3.1) *is given by*

$$\mathcal{R}(t,n,\lambda) = \begin{cases} \sum_{i=1}^{n} \gamma_{n,i}\, e^{-\lambda \sum_{j=1}^{i} F(j)t} & \textit{if } F(j) > 0 \\ n & \textit{otherwise,} \end{cases} \tag{1.3.2}$$

where the constants $\gamma_{n,i}$ can be computed recursively according to

$$\gamma_{n,n} = n - \sum_{i=1}^{n-1} \gamma_{n,i} \quad \textit{and} \quad \gamma_{n,i} = 2 \cdot \frac{\sum_{j=1}^{n-i} F(j)\gamma_{n-j,i}}{\sum_{k=i+1}^{n} F(k)}, \quad 1 \le i < n,$$

with boundary value $\gamma_{i,j} = 1$ for all i, j such that $F(i) = 0$ and $j \le i$. Furthermore, $\mathcal{C}(t, n, \lambda)$ is given by

$$\mathcal{C}(t,n,\lambda) = n - \sum_{i=1}^{n-1} \tilde{\gamma}_{n,i} e^{-\lambda \sum_{j=1}^{i} F(j)t} - \tilde{\gamma}_{n,n} e^{-\lambda n t}, \ \textit{for } n \ge m$$

where

$$\tilde{\gamma}_{n,i} = \sum_{k=1}^{n-i} f(k) \frac{\gamma_{n-k,i}}{1 - \frac{1}{n}\sum_{j=1}^{i} F(j)} \quad \textit{and} \quad \tilde{\gamma}_{n,n} = n - \sum_{i=1}^{n-1} \tilde{\gamma}_{n,i}.$$

We point out that the exact computation of $\mathcal{C}(t,n,\lambda)$ involves solving recursively for the coefficients $\gamma_{n,i}$, which can be time consuming for very large n. Since n will indeed be large in real-world applications, it is desirable to have a good approximation. In particular, we propose replacing $\mathcal{C}(t,n,\lambda)/n$ by its asymptotic as $n \to \infty$, since we will show, both theoretically and numerically, that the convergence is very fast. Theorem 2 below provides an expression for the limit.

Theorem 2 *Define*

$$\alpha(t,\lambda) = \lim_{n\to\infty} \frac{\mathcal{K}(t,n,\lambda)}{n}$$

to be the asymptotic proportion of delivered packages during the time interval $[0,t]$*, with arrival rate* λ*. Then,*

$$\lim_{n\to\infty} \frac{\mathcal{C}(t,n,\lambda)}{n} = \alpha(t,\lambda)$$

and

$$\alpha(t,\lambda) = 1 - \int_{e^{-\lambda t}}^{1} e^{2(\varphi(u)-\varphi(1))}\hat{q}(t+(\ln u)/\lambda, u)\, du - (m-(m-1)e^{-\lambda t})e^{2(\varphi(e^{-\lambda t})-\varphi(1))-\lambda t(m-\varphi'(1))}, \tag{1.3.3}$$

where

$$\varphi(y) = \sum_{i=1}^{m-1} \frac{\overline{F}(i)}{i} y^i,$$

$\overline{F}(i) = 1 - F(i)$, $\mathcal{R}(t,n,\lambda) = n - \mathcal{K}(t,n,\lambda)$ *and*

$$\hat{q}(s,v) := 2v^{m-\varphi'(1)-1}(1-v)\sum_{i=1}^{m-1}\mathcal{R}(s,i) - 2(1-v)^2\sum_{i=1}^{m-1}\mathcal{R}(s,i)\sum_{j=m-i}^{m-1}\overline{F}(j)v^{i+j-\varphi'(1)-1}.$$

To compare the computational complexities of Proposition 1.3.1 and Theorem 2, note that to calculate $\alpha(t,\lambda)$, we only need to compute $\{\gamma_{i,j} : 1 \le j \le i\}$ for $i < m$, since only the terms $\mathcal{R}(t+(\ln u)/\lambda, i, \lambda)$ for $1 \le i \le m-1$ appear in the expression for $\hat{q}(t+(\ln u)/\lambda, u)$. To calculate $\mathcal{R}(t,n,\lambda)$ using Proposition 1.3.1, we need to calculate each $\gamma_{i,j}$ for $i \le j$ and $j \le n$. Therefore, the complexity using Proposition 1.3.1 is $\Theta(n^2)$ while that of using Theorem 2 is $O(1)$. This is a huge difference in the computational cost.

Moreover, whenever $P(B=1) > 0$ we have that $\lim_{t\to\infty}\mathcal{R}(t,i) = 0$ for all $i \ge 1$, which yields

$$\lim_{t\to\infty}\hat{q}(t+(\ln u)/\lambda, u) = 0$$

for all $u \in [0,1)$, and therefore,

$$\lim_{t\to\infty}\alpha(t,\lambda) = 1,$$

as expected. Note that Theorem 2 can also be used to compute the limiting proportion of remaining packages when $P(B=1) = 0$, in which case the quantity of interest is the number of remaining packages at time $t = \infty$. For the special case $P(B = m) = 1$ $(m > 1)$ Pinsky [86] obtained the formula

$$\alpha(\infty, \lambda) = m \exp\left(-2\sum_{j=1}^{m-1} \frac{1}{j}\right) \int_0^1 \exp\left(2\sum_{j=1}^{m-1} \frac{u^j}{j}\right) du = m \int_0^1 e^{2\varphi(u)-\varphi(1)} du$$

which can be derived from Theorem 2 by noting that $\varphi'(1) = m - 1$ and

$$\begin{aligned}\lim_{t\to\infty} \hat{q}(t + (\ln u)/\lambda, u) &= 2(1-u)\sum_{i=0}^{m-1} i\left(1 - (1-u)\sum_{j=m-i}^{m-1} u^{i+j-m}\right) \\ &= 2(1-u)\sum_{i=0}^{m-1} iu^i = 2(u\varphi'(u) - (m-1)u^m) \\ &= 2\varphi'(u)(mu - (m-1)u^2) - 2(m-1)u\end{aligned}$$

for all $\lambda > 0$, and therefore,

$$\begin{aligned}\lim_{t\to\infty} \alpha(t, \lambda) &= 1 - \int_0^1 e^{2(\varphi(u)-\varphi(1))} \lim_{t\to\infty} \hat{q}(t + (\ln u)/\lambda, u)\, du \\ &= 1 - \int_0^1 e^{2(\varphi(u)-\varphi(1))} 2\varphi'(u)(mu - (m-1)u^2) du + 2(m-1)\int_0^1 e^{2(\varphi(u)-\varphi(1))} u\, du \\ &= \int_0^1 e^{2(\varphi(u)-\varphi(1))}(m - 2(m-1)u)\, du + 2(m-1)\int_0^1 e^{2(\varphi(u)-\varphi(1))} u\, du \\ &= m\int_0^1 e^{2(\varphi(u)-\varphi(1))} du.\end{aligned}$$

In addition to the expression for $\alpha(t, \lambda)$ in Theorem 2, we also provide the corresponding rate of convergence of $\mathcal{C}(t, n, \lambda)/n$ to $\alpha(t, \lambda)$. To the best of our knowledge, this is the first result regarding the rate of convergence for the limiting proportion of packages picked up by time t. Although we do not include the details of the proof in this paper, a similar set of arguments as those used in the proof of Theorem 3 also yield the result for $t = \infty$. Throughout the paper we use $f(x) = O(g(x))$ as $x \to \infty$ if $\limsup_{x\to\infty} |f(x)/g(x)| < \infty$.

As pointed out earlier, we claim that one can safely substitute $\mathcal{C}(t, n, \lambda)/n$ with $\alpha(t, \lambda)$ when n is large. Theorem 3 below provides a theoretical justification by stating that the rate of convergence is of order $O(n^{-1})$, which will be considerably of smaller order than all other terms in the objective function we will optimize (see Section 1.3.2). In Section 1.5 we provide additional numerical evidence that the asymptotic can be used even for moderately large values of n.

Theorem 3 *For any fixed* t,

$$\left|\frac{\mathcal{C}(t, n, \lambda)}{n} - \alpha(t, \lambda)\right| = O(n^{-1}), \qquad \text{as } n \to \infty.$$

We now move on to the pricing part of our model.

1.3.2 Computing the reward for each package.

The overall goal of the proposed framework is to provide a pricing strategy for delivering n different packages using a combination of private drivers and in-house delivery vans. Section 1.3.1 provided analytical results for the expected number of packages delivered during time $[0, T]$ as a function of the arrival rate λ, under the assumption that all packages are equally desirable (λ is the same for all n locations). As mentioned earlier, it is through the pricing mechanism that we will justify the modeling assumption on λ, since we would naturally expect that packages with remote destinations would receive fewer requests. Our payment scheme is based on the idea that the amount of money that a driver can make per unit of time should be the same for all packages, and we accomplish this by separating the costs associated to the destination of each package from those of a common "incentive rate". The package specific costs will take into account factors such as the distance between the destination and the warehouse (long-haul distance) and the distance between neighboring packages with respect to the TSP route (local distance). Once we have provided an expression for the cost of delivering packages through the use of private drivers, we will need to estimate the cost of delivering the remaining packages using in-house vans. The detailed description of our pricing for the delivery using private drivers is given in Section 1.3.2.1, and the corresponding vans' cost is given in Section 1.3.2.2.

1.3.2.1 Rewards for private drivers.

For the delivery process, assume that both private drivers and in-house vans must pick up the packages they will be delivering from the distribution center. At the beginning of the day, the destinations of the n packages to be delivered that day are revealed and an optimal TSP tour is computed. Let $\{\mathbf{x}_1, ..., \mathbf{x}_n\} \subseteq \mathbb{R}^2$ denote the destinations of the n packages, where their indexes correspond to their "locations" within the TSP route. The distribution center is assumed to be at the origin. We denote by $\hat{B}^{(i)}$ the size of the bundle of the first request at location i to be accepted. Note that the distribution of $\hat{B}^{(i)}$ is not F since the acceptance depends on the configuration of available packages at the time of the request.

The overall cost for private drivers to deliver packages consists of transportation costs and opportunity costs, since drivers can also choose to work for Uber-like companies or take other jobs instead. We use the following quantities in our cost estimation; $d(\mathbf{x}, \mathbf{y})$ denotes the distance on $\mathbb{R}^2$:

- $r_i = d(\mathbf{x}_i, \mathbf{0})$ denotes the distance from the depot to the destination of package i (the long-haul distance).
- $d_i = (d(\mathbf{x}_{i-1}, \mathbf{x}_i) + d(\mathbf{x}_i, \mathbf{x}_{i+1}))/2$ is the average distance between the destinations of packages $i-1$ and i, and i and $i+1$, for $2 \leq i \leq n-1$, $d_1 = (d(\mathbf{x}_n, \mathbf{x}_1) + d(\mathbf{x}_1, \mathbf{x}_2))/2$, $d_n = (d(\mathbf{x}_{n-1}, \mathbf{x}_n) + d(\mathbf{x}_n, \mathbf{x}_1))/2$ (the local distance).
- ζ_P is the per-mile transportation cost.

- h_P is the opportunity cost per unit of time, i.e., the profit earned while doing other jobs such as transporting people for Uber or Lyft.
- τ_P is the end-point delivery time.
- v_P is the average speed of private cars.

The decision variable in our pricing model will be an incentive rate z that each driver will receive in addition to the opportunity cost, i.e., the total payment rate that a driver receives per unit of time is $h_P + z$. Note that the only quantities in the cost that depend on the geographic location of the package destinations are the $\{r_i\}$ and the $\{d_i\}$. The $\{r_i\}$ can be computed as soon as the destinations $\{\mathbf{x}_i\}$ are revealed, while the $\{d_i\}$ are determined by the TSP route.

The traveling distance associated with a bundle of packages at locations $\{i, i+1, \ldots, i+k-1\}$ is $r_i + \sum_{j=i}^{i+k-2} d(\mathbf{x}_j, \mathbf{x}_{j+1})$. Thus, the price set for the bundle should be

$$\text{price}_1 = \zeta_P \left(r_i + \sum_{j=i}^{i+k-2} d(\mathbf{x}_j, \mathbf{x}_{j+1}) \right) + (h_P + z) \left(r_i/v_P + \sum_{j=i}^{i+k-2} d(\mathbf{x}_j, \mathbf{x}_{j+1})/v_P + k\tau_P \right). \tag{1.3.4}$$

However, since the number of possible bundles increases geometrically with the number of total packages, it is computationally expensive to set a price for every possible bundle. Therefore, we consider instead a pricing scheme for each individual package, regardless of which bundle it will be included in. To derive this price, suppose that package j is delivered as part of a bundle of size k, and start by prorating the long-haul cost among all the k packages, and separate the contribution of package j to the local distance. To incorporate into the pricing longer neighboring distances between adjacent packages in the TSP tour, we determine the contribution of package j to the local distance to be the average of the distances to both the neighbor to the left and the neighbor to the right, i.e., $d_j = (d(\mathbf{x}_{j-1}, x_j) + d(\mathbf{x}_j, \mathbf{x}_{j+1}))/2$. We also argue that the long-haul cost of package j is approximately the same as that of other packages in the same bundle, and therefore if j is part of a bundle accepted at location i, then $r_j \approx r_i$. We then propose the payment reward for package j to be:

$$\zeta_P \left(\frac{r_j}{k} + d_j \right) + (h_P + z) \left(\frac{r_j}{kv_P} + \frac{d_j}{v_P} + \tau_P \right). \tag{1.3.5}$$

Using (1.3.5) we obtain a price for a bundle of size k accepted at location i of the form:

$$\begin{aligned}\text{price}_2 &= \zeta_P\left(\frac{1}{k}\sum_{j=i}^{i+k-1} r_j + \sum_{j=i}^{i+k-1} d_j\right) + (h_P+z)\left(\frac{1}{kv_P}\sum_{j=i}^{i+k-1} r_j + \sum_{j=i}^{i+k-1} d_j/v_P + k\tau_P\right) \quad (1.3.6)\\ &= \text{price}_1 + \left(\zeta_P + \frac{h_P+z}{v_P}\right)\left(\frac{1}{k}\sum_{j=i}^{i+k-1} r_j - r_i\right)\\ &\quad + \frac{1}{2}(d(\mathbf{x}_{i-1},\mathbf{x}_i) + d(\mathbf{x}_{i+k-2},\mathbf{x}_{i+k-1}))\left(\zeta + \frac{h_P+z}{v_P}\right).\end{aligned}$$

Note that the difference between (1.3.4) and (1.3.6) is small whenever adjacent packages in the TSP tour are small, which we expect to be true for large n, as suggested by [94].

To obtain our proposed expression for the cost to deliver each of the n packages we also need to take into account that package j could be delivered as part of a number of different bundles, e.g., it could be delivered in bundle $\{j, \dots, j+\hat{B}^{(j)}-1\}$ or it could be contained in a bundle of the form $\{i, \dots, j, \dots, i+\hat{B}^{(i)}-1\}$ for some $i<j$. Since the exact computation of the distribution of the size of the bundle containing j is too complex, we approximate it with $E[B]$ to obtain the following price for package j:

$$p_j := \zeta_P\left(\frac{r_j}{E[B]} + d_j\right) + (h_P+z)\left(\frac{r_j}{E[B]v_P} + \frac{d_j}{v_P} + \tau_P\right). \quad (1.3.7)$$

Using the same type of arguments, we estimate the time required to deliver package j to be:

$$t_j := \frac{r_j}{E[B]v_P} + \frac{d_j}{v_P} + \tau_P.$$

Note that by deriving our pricing mechanism the way we did we have made the profit rate for the drivers the same regardless of which package(s) they choose to deliver. This profit is determined in our pricing scheme by the incentive rate z, which is the same for all n packages and is linear in the total traveling distance. Moreover, the incentive rate will be used to control the arrival rate for requests in our calculations from Section 1.3.1 by setting $\lambda = \lambda(z)$ to be a non-decreasing function. It remains to set up an optimization problem for determining the best incentive rate to use.

To this end, start by noting that the sum of the prices for all n packages after their destinations are revealed satisfies:

$$\sum_{i=1}^n p_i = \left(\zeta_P + \frac{h_P+z}{v_P}\right)\left(\frac{1}{E[B]}\sum_{i=1}^n r_i + L(\text{TSP}(\mathbf{x}^{(n)}))\right) + n(h_P+z)\tau_P,$$

where $L(\text{TSP}(\mathbf{x}^{(n)}))$ is the length of an optimal TSP route for points with destinations $\mathbf{x}^{(n)} := \{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n\}$. Since the probability that package i will be picked up before time T given incentive rate z is $\mathcal{C}(T, n, \lambda(z))/n$ (recall that packages are arranged on a circle, and are therefore undistinguishable), the aggregate expected payment for private drivers (conditional on $\mathbf{x}^{(n)}$) is

$$\frac{\mathcal{C}(T,n,\lambda(z))}{n}\left(\zeta_P+\frac{h_P+z}{v_P}\right)\left(\frac{1}{E[B]}\sum_{i=1}^{n}r_i+L(\mathrm{TSP}(\mathbf{x}^{(n)}))\right)+\mathcal{C}(T,n,\lambda(z))(h_P+z)\tau_P. \tag{1.3.8}$$

It remains to compute the cost associated to delivering the remaining packages using the in-house van service, which we do in the following section.

1.3.2.2 Van's cost to deliver leftover packages.

After time T, all leftover packages will be delivered by vans owned by the distribution center. The delivery route for a van is designed by an optimal CVRP, which is also NP-hard. For more detailed information about the algorithms that can be used to solve the CVRP problem, we refer readers to [43, 45, 111].

The van's operating cost includes the per-mile cost for vans and time-based wages for the drivers. The per-mile cost includes the cost of fuel, maintenance, repairs, depreciation, etc., and is denoted by ζ_V. To compute the time-based wages for the drivers note that the time they spend delivering packages includes the time driving along an optimal CVRP route and the time on the end-point delivery. Let h_V, v_V, τ_V denote the drivers' payment rate, the vans' average speed, and the end-point delivery time, respectively. Then, the total cost for delivering k packages using vans is

$$\left(\zeta_V+\frac{h_V}{v_V}\right)L(\mathrm{CVRP}(\mathbf{y}^{(k)}))+kh_V\tau_V,$$

where $L(\mathrm{CVRP}(\mathbf{y}^{(k)}))$ is the length of capacitated vehicle routing through the points $\mathbf{y}^{(k)}=\{\mathbf{y}_1,\mathbf{y}_2,..,\mathbf{y}_k\}$.

In the context of our problem, the number of packages that will need to be delivered after time T is random, and so are their destinations, which we will denote by $\mathbf{Y}^{(n)}(\lambda)=\{\mathbf{Y}_1,\dots,\mathbf{Y}_{n-\tilde{N}(T,n,\lambda)}\}$, where $\tilde{N}(T,n,\lambda)$ is the number of packages that can be delivered during $[0,T]$ when we start with n packages arranged on a circle. It follows that the expected cost to deliver the remaining packages, conditionally on the destinations $\mathbf{x}^{(n)}$, is given by

$$\begin{aligned}&\left(\zeta_V+\frac{h_V}{v_V}\right)\mathbb{E}_n\left[L(\mathrm{CVRP}(\mathbf{Y}^{(n)}(\lambda)))\right]+\mathbb{E}_n[n-\tilde{N}(T,n,\lambda(z))]h_V\tau_V\\&=\left(\zeta_V+\frac{h_V}{v_V}\right)\mathbb{E}_n\left[L(\mathrm{CVRP}(\mathbf{Y}^{(n)}(\lambda)))\right]+\left(1-\frac{\mathcal{C}(T,n,\lambda(z))}{n}\right)nh_V\tau_V,\end{aligned} \tag{1.3.9}$$

where $\mathbb{E}_n[\,\cdot\,]=E[\,\cdot\,|\mathbf{x}^{(n)}]$.

Putting together our cost estimations for both the private drivers (1.3.8) and that of delivering the remaining packages using the in-house vans (1.3.9), we obtain the following expected cost function for delivering n packages with destinations $\mathbf{x}^{(n)} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ using the incentive rate z:

$$\begin{aligned}\frac{\mathrm{Cost}_P(z;\mathbf{x}^{(n)})}{n} := {} & \frac{\mathcal{C}(T,n,\lambda(z))}{n}\left(\zeta_P + \frac{h_P+z}{v_P}\right)\left(\frac{\overline{r}(n)}{E[B]} + \frac{L(\mathrm{TSP}(\mathbf{x}^{(n)}))}{n}\right) \\ & + \frac{\mathcal{C}(T,n,\lambda(z))}{n}(h_P+z)\tau_P + \left(1 - \frac{\mathcal{C}(T,n,\lambda(z))}{n}\right)h_V\tau_V \\ & + \left(\zeta_V + \frac{h_V}{v_V}\right)\frac{1}{n}\mathbb{E}_n\left[L(\mathrm{CVRP}(\mathbf{Y}^{(n)}(\lambda)))\right],\end{aligned}$$

where $\overline{r}(n) = n^{-1}\sum_{i=1}^n r_i$. Note that the computation of the length of the TSP route, $L(\mathrm{TSP}(\mathbf{x}^{(n)}))$, is done as part of our proposed approach, however, we still need to compute the expected length of the CVRP route for the leftover packages, i.e., $\mathbb{E}_n\left[L(\mathrm{CVRP}(\mathbf{Y}^{(n)}(\lambda)))\right]$.

Since the exact computation of $\mathbb{E}_n\left[L(\mathrm{CVRP}(\mathbf{Y}^{(n)}(\lambda)))\right]$ is intractable, both from the point of view of the distribution of the leftover packages and also from that of the length of an optimal CVRP route, we instead estimate it using the continuous approximation given in [35]. Specifically, for a set of k points $\{\mathbf{y}_1, \ldots, \mathbf{y}_k\}$ distributed uniformly over a region of area A, and under the assumption that the vehicle capacity V satisfies $V \ll k$, the continuous approximation given in [35] states that

$$L(\mathrm{CVRP}(\mathbf{y}^{(k)})) \approx \frac{2k}{V}\overline{r}(k) + \beta_{\mathrm{VRP}}\sqrt{kA}, \tag{1.3.10}$$

where β_{VRP} is estimated to be 0.82 when using the L_1 distance (see Appendix A in [34]). We point out that the assumption of $V \ll k$ is realistic, since we expect the number of remaining packages at the end of the time window $[0,T]$ to be large. To justify this belief, we mention that in the study of UPS delivery routes given in [52], each distribution center needs to deliver around 10,000 packages per day, and each van serves 140-160 customers. Our simulations suggest that by time T we have about 2156 and 4026 leftover packages when we start the day with $n = 8000$ and $n = 15,000$, respectively, which implies that $V/k \approx 0.070$ for $n = 8000$ and $V/k \approx 0.037$ for $n = 15,000$. Moreover, approximation (1.3.10) has been empirically shown in [105] to be accurate even for moderate values of k, e.g., $k \geq 100$. For our packing problem, we know that the expected number of leftover packages is $n - \mathcal{C}(T,n,\lambda(z))$ and the expected sum of their long-haul distances is $\mathbb{E}_n\left[\sum_{i=1}^n r_i 1(\mathbf{x}_i \in \mathbf{Y}^{(n)}(\lambda))\right] = (n - \mathcal{C}(T,n,\lambda(z)))\overline{r}(n)$, so we use the approximation

$$\frac{1}{n}\mathbb{E}_n\left[L(\mathrm{CVRP}(\mathbf{Y}^{(n)}(\lambda)))\right] \approx \frac{2(1-\mathcal{C}(T,n,\lambda(z))/n)\overline{r}(n)}{V} + \frac{\beta_{\mathrm{VRP}}}{\sqrt{n}}\sqrt{(1-\mathcal{C}(T,n,\lambda(z))/n)A}.$$

Based on our numerical experiments, this approximation seems to work well, even though the destinations of the leftover packages are not necessarily uniformly distributed. Adding the lower order term $(\beta_{\mathrm{VRP}}/\sqrt{n})\sqrt{(1-\mathcal{C}(T,n,\lambda(z))/n)A}$ seems to improve the overall cost optimization, especially for moderate values of n. For large values of n, we can go further in our simplification of the cost function and replace $\mathcal{C}(T, n, \lambda(z))/n$ with $\alpha(T, \lambda(z))$, which in view of Theorem 3 incurs an error of order $O(n^{-1})$, which is negligible with respect to any of the terms in $\text{Cost}_P(z, \mathbf{x}^{(n)})/n$, and significantly reduces its computing time (see Table 1.3 in Section 1.5.3).

We finally arrive at the following approximation for $\text{Cost}_P(z, \mathbf{x}^{(n)})/n$:

$$\begin{aligned}\frac{\text{Cost}'_P(z; \mathbf{x}^{(n)})}{n} &:= \alpha(T, \lambda(z)) \left(\zeta_P + \frac{h_P + z}{v_P}\right) \left(\frac{\overline{r}(n)}{E[B]} + \frac{L(\text{TSP}(\mathbf{x}^{(n)}))}{n}\right) \\ &\quad + \alpha(T, \lambda(z))(h_P + z)\tau_P + (1 - \alpha(T, \lambda(z)))\, h_V \tau_V \\ &\quad + \left(\zeta_V + \frac{h_V}{v_V}\right) \left(\frac{2(1 - \alpha(T, \lambda(z)))\overline{r}(n)}{V} + \frac{\beta_{\text{VRP}}}{\sqrt{n}} \sqrt{(1 - \alpha(T, \lambda(z)))A}\right).\end{aligned}$$

Our proposed solution to the problem of selecting an optimal incentive rate for the private drivers is to compute

$$\hat{z} = \underset{z}{\text{argmin}}\ \text{Cost}'_P(z; \mathbf{x}^{(n)}), \tag{1.3.11}$$

as a proxy for the intractable

$$z^* = \underset{z}{\text{argmin}}\ \text{Cost}_P(z; \mathbf{x}^{(n)}). \tag{1.3.12}$$

The following lemma provides theoretical justification for our proposed approach when n is large. We also include in Section 1.5.3 (see Table 1.4) numerical evidence supporting the use of our approximation even for moderate values of n.

Lemma 1.3.2 *Define z^* and $\hat{z}$ according to* (1.3.12) *and* (1.3.11), *respectively. Then,*

$$0 \le \frac{Cost_P(\hat{z}; \mathbf{x}^{(n)})}{n} - \frac{Cost_P(z^*; \mathbf{x}^{(n)})}{n} = O(n^{-1/2}), \qquad n \to \infty.$$

1.3.2.3 Joint optimization problem.

We will now argue that it suffices to minimize $\text{Cost}'_P(z; \mathbf{x}^{(n)})$ over a bounded interval. The arrival rate $\lambda(z)$ can be taken to be any non-negative monotone non-decreasing and differentiable almost everywhere function on the real line, e.g., linear or piecewise linear.

Note that the distribution center will be paying private drivers $\zeta_P + (h_P + z)/v_P$ per mile travelled, plus $(h_P + z)\tau_P$ per package delivered, while it will pay its van drivers $\zeta_V + h_V/v_V$ per mile travelled and $h_V \tau_V$ per package delivered. Hence, in order for a strategy using private drivers to even make sense, we would need at least one of the following conditions to hold:

$$\zeta_V + \frac{h_V}{v_V} > \zeta_P + \frac{h_P + z}{v_P} \qquad \text{or} \qquad h_V \tau_V > (h_P + z)\tau_P.$$

In other words, if the optimal $\hat{z}$ where to violate both conditions, then the optimal strategy would be to use vans only. This implies that it suffices to consider values of z such that

$$z \le \max\left\{ \left(\zeta_V + \frac{h_V}{v_V} - \zeta_P\right) v_P,\ \frac{h_V \tau_V}{\tau_P} \right\} - h_P.$$

On the other hand, since the payment rate $(h_p + z)$ to private drivers (not including the per-mile transportation cost ζ_P) must be nonnegative, we have that $h_P + z \geq 0$, which implies that

$$z \geq -h_P.$$

In view of the above, we propose to compute the approximately optimal incentive rate $\hat{z}$ by solving:

$$\min_{-h_P \leq z \leq \max\left\{\left(\zeta_V + \frac{h_V}{v_V} - \zeta_P\right)v_P,\, \frac{h_V \tau_V}{\tau_P}\right\} - h_P} \text{Cost}'_P(z; \mathbf{x}^{(n)}). \tag{1.3.13}$$

The optimization problem is meant to be solved at the beginning of the day, once the destinations $\mathbf{x}^{(n)}$ are revealed. Once $\hat{z}$ has been found the reward offered to private drivers for delivering package i is p_i, as given by (1.3.7).

We point out that since $\mathcal{C}(T, n, \lambda(z))$ is infinitely differentiable in $z \in \mathbb{R}$ whenever $\lambda(z)$ is (see Proposition 1.3.1), solving the minimization problem in (1.3.13) can be done very efficiently. The problem is to be solved on a daily basis as soon as the destinations $\mathbf{x}^{(n)} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ are revealed.

We conclude the main results by providing sufficient conditions under which the expected cost to deliver n packages with destinations $\mathbf{x}^{(n)} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ using a combination of private drivers and in-house vans is smaller than that of using only in-house vans. To this end, note that the expected cost of the van-only strategy is given by:

$$\text{Cost}_V(\mathbf{x}^{(n)}) := \left(\zeta_V + \frac{h_V}{v_V}\right) L(\text{CVRP}(\mathbf{x}^{(n)})) + n h_V \tau_V.$$

Hence, a mixed strategy that uses both vans and private drivers can only be better provided there exists some z for which $\text{Cost}_P(z; \mathbf{x}^{(n)}) < \text{Cost}_V(\mathbf{x}^{(n)})$. The following lemma provides the desired condition.

Lemma 1.3.3 *Suppose the destinations $\mathbf{x}^{(n)} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ are contained in a compact region $R \subseteq \mathbb{R}^2$ and are such that the limit $r^* = \lim_{n\to\infty} \overline{r}(n)$ exists. Then, for any $z \in \mathbb{R}$ and $\lambda = \lambda(z)$ we have*

$$\limsup_{n\to\infty} \frac{1}{n} \left(Cost_P(z; \mathbf{x}^{(n)}) - Cost_V(\mathbf{x}^{(n)})\right)$$
$$\leq -\alpha(T, \lambda) \left(\left(\zeta_V + \frac{h_V}{v_V}\right)\frac{2r^*}{V} - \left(\zeta_P + \frac{h_P}{v_P}\right)\frac{r^*}{E[B]} - (h_P\tau_P - h_V\tau_V) - z\left(\frac{r^*}{v_P E[B]} + \tau_P\right)\right) \quad a.s.$$

Moreover, whenever

$$\left(\zeta_V + \frac{h_V}{v_V}\right)\frac{2r^*}{V} - \left(\zeta_P + \frac{h_P}{v_P}\right)\frac{r^*}{E[B]} - (h_P\tau_P - h_V\tau_V) > 0, \tag{1.3.14}$$

there exists a $z \in \mathbb{R}$ for which the upper bound for the limit superior is strictly negative, i.e., for which the private drivers strategy is better than the van-only strategy.

We point out that the condition given in (1.3.14) does not depend on z, since the non-negativity of $\alpha(T,\lambda)$ for all $\lambda \geq 0$ implies that whenever (1.3.14) is satisfied there will exist a z for which $\text{Cost}_P(z;\mathbf{x}^{(n)}) < \text{Cost}_V(\mathbf{x}^{(n)})$ with high probability. The magnitude of the cost reduction will of course depend on the specific value of z we choose, and it is maximized when we choose $z = z^*$.

Until now, we provided a framework to deal with the scenario where packages only arrive once during the day and the optimal incentive rate is set as soon as the package destinations are revealed. However, our framework can be easily extended to multiple periods, i.e., to situations where packages arrive twice or more times during the day and/or the supply of private drivers fluctuates throughout different time periods. The flexibility to change the incentive rate may be important, for example, if the distribution center wants to compete with the Uber or Lyft passenger population during peak hours. If this were the case, the distribution center would need to update the opportunity cost, the end-point delivery time and/or the average speed of private drivers in the objective function. Recomputing the optimal TSP route may be unnecessary if the number of packages is large enough at the beginning of the day and continue being large even after some packages have been picked up. Hence, our proposed framework can scale very well to a multi-period setting, without incurring expensive computational costs.

The remainder of the paper is devoted to the proofs of all the results in the paper.

1.4 Proofs.

In this section, we give all the proofs of the theorems in Section 1.3.1.1 and Section 1.3.2. The analysis of $\mathcal{C}(t,n,\lambda)$ is based on the observation that once the first package is picked up (which is guaranteed to be accepted whenever $n \geq m$), the remaining packages can be arranged on a line. Therefore, the proofs of all our theoretical results are based on the analysis of $\mathcal{K}(t,n,\lambda)$, the expected number of packages that can be picked up by private drivers during the interval $[0,t]$, when there are n packages arranged on a line. Moreover, we point out that if B has distribution F and T_n^* denotes the time of the first request when we have n packages arranged on a circle, then

$$\mathcal{C}(t,n,\lambda) = E\left[(B + \mathcal{K}(t - T_n^*, n - B))\, 1(T_n^* \leq t)\right],$$

with T_n^* exponentially distributed with rate λn.

Throughout this section we simplify the notation by omitting λ from $\mathcal{C}(t,n,\lambda)$, $\mathcal{K}(t,n,\lambda)$, $\mathcal{R}(t,n,\lambda)$, etc., and simply write $\mathcal{C}(t,n)$, $\mathcal{K}(t,n)$, $\mathcal{R}(t,n)$; all the proofs in this section are valid for any fixed $\lambda > 0$. We also use $f(k) = P(B = k)$ to denote the bundle size probability mass function. The first proof corresponds to Theorem 1, which gives a differential equation for $\mathcal{K}(t,n)$ and $\mathcal{R}(t,n) = n - \mathcal{K}(t,n)$.

Proof. Proof of Theorem 1. Define $N(t,n)$ to be the number of packages that can be delivered over the period $[0,t]$ when packages are arranged on a line, i.e., $\mathcal{K}(t,n) = E[N(t,n)]$.

Looking at the first Δ units of time after time t, we obtain

$$
\begin{aligned}
N(t+\Delta, n) = {} & N(t+\Delta, n)1(\text{no requests arrive in } [t, t+\Delta]) \\
& + \sum_{w=1}^{m} N(t+\Delta, n)1(\text{bundle of size } w \text{ arrives in } [t, t+\Delta]) \\
& + N(t+\Delta, n)1(\text{two or more requests arrive in } [t, t+\Delta]).
\end{aligned} \tag{1.4.1}
$$

By conditioning on the location of the arrival we further get

$$
\begin{aligned}
& N(t+\Delta, n)1(\text{bundle of size } w \text{ arrives in } [t, t+\Delta]) \\
&= \sum_{i=1}^{n} N(t+\Delta, n)1(\text{bundle of size } w \text{ arrives at position } i \text{ in } [t, t+\Delta]) \\
&= 1(n \geq w) \sum_{i=1}^{n-w+1} N(t+\Delta, n)1(\text{bundle of size } w \text{ arrives at position } i \text{ in } [t, t+\Delta]) \\
&\quad + 1(n \geq w) \sum_{i=n-w+2}^{n} N(t, n)1(\text{bundle of size } w \text{ arrives at position } i \text{ in } [t, t+\Delta]), \quad (1.4.2) \\
&\quad + 1(n < w) \sum_{i=1}^{n} N(t, n)1(\text{bundle of size } w \text{ arrives at position } i \text{ in } [t, t+\Delta]), \quad (1.4.3)
\end{aligned}
$$

where (1.4.2) and (1.4.3) correspond to the cases where we reject the arrival since either some packages in the requested bundle are not available or the bundle size is bigger than the number of remaining packages. Taking expectation on both sides of (1.4.1), we have

$$
\begin{aligned}
&\mathcal{K}(t+\Delta, n) \\
&= E\left[N(t,n)1(\text{no requests arrive in } [t,t+\Delta])\right] \\
&\quad + E\left[N(t+\Delta,n)1(\text{two or more requests arrive in } [t,t+\Delta])\right] \\
&\quad + \sum_{w=1}^{m} 1(n \geq w) \sum_{i=1}^{n-w+1} E\left[N(t+\Delta,n)1(\text{bundle of size } w \text{ arrives at position } i \text{ during } [t,t+\Delta])\right] \\
&\quad + \sum_{w=1}^{m} 1(n \geq w) \sum_{i=n-w+2}^{n} E[N(t,n)1(\text{bundle of size } w \text{ arrives at position } i \text{ during } [t,t+\Delta])] \\
&\quad + \sum_{w=1}^{m} 1(n < w) \sum_{i=1}^{n} E[N(t,n)1(\text{bundle of size } w \text{ arrives at position } i \text{ during } [t,t+\Delta])] \\
&= \mathcal{K}(t,n)P(\text{no requests arrive in } [t,t+\Delta]) + O(\Delta^2 n^3) \\
&\quad + \sum_{w=1}^{m} 1(n \geq w) \sum_{i=1}^{n-w+1} \{\mathcal{K}(t,i-1) + \mathcal{K}(t,n-i-w+1) + w\} \\
&\qquad\qquad \cdot P(\text{bundle of size } w \text{ arrives at position } i \text{ during } [t,t+\Delta]) \\
&\quad + \sum_{w=1}^{m} 1(n \geq w) \sum_{i=n-w+2}^{n} \mathcal{K}(t,n)P(\text{bundle of size } w \text{ arrives at position } i \text{ during } [t,t+\Delta]) \\
&\quad + \sum_{w=1}^{m} 1(n < w) \sum_{i=1}^{n} \mathcal{K}(t,n)P(\text{bundle of size } w \text{ arrives at position } i \text{ during } [t,t+\Delta])
\end{aligned}
$$

where we have used the following three observations: first, that

$$
\begin{aligned}
&E\left[N(t+\Delta,n)1(\text{two or more requests arrive in } [t,t+\Delta])\right] \\
&\leq nP(\text{two or more requests arrive in } [t,t+\Delta]) = O(\Delta^2 n^3)
\end{aligned}
$$

as $\Delta \to 0$; second, that

$$
\begin{aligned}
&E[N(t+\Delta,n)|\text{bundle size } w \text{ arrives at position } i \text{ during } [t,t+\Delta]] \\
&= \mathcal{K}(t,i-1) + \mathcal{K}(t,n-i-w+1) + w;
\end{aligned}
$$

and third, that $N(t+\Delta,n) = N(t,n)$ whenever the bundle request is rejected (i.e., when

$n - w + 2 \leq i \leq n$ for $n \geq w$ or when $1 \leq i \leq n$ for $n < w$). We thus have

$$\begin{aligned}\mathcal{K}(t+\Delta, n) = \mathcal{K}(t,n)(1 - \lambda n \Delta) + O(\Delta^2 n^3) \\ + \sum_{w=1}^{m} 1(n \geq w) \sum_{i=1}^{n-w+1} \{\mathcal{K}(t, i-1) + \mathcal{K}(t, n-i-w+1) + w\} f(w) \cdot \frac{1}{n} \cdot (\lambda n \Delta) \\ + \sum_{w=1}^{m} 1(n \geq w) \sum_{i=n-w+2}^{n} \mathcal{K}(t,n) f(w) \cdot \frac{1}{n} \cdot (\lambda n \Delta) \\ + \sum_{w=1}^{m} 1(n < w) \sum_{i=1}^{n} \mathcal{K}(t,n) f(w) \cdot \frac{1}{n} \cdot (\lambda n \Delta).\end{aligned}$$

We have thus shown that

$$\begin{aligned}\mathcal{K}(t+\Delta, n) - \mathcal{K}(t,n) &= -\lambda n \Delta \mathcal{K}(t,n) + O(\Delta^2 n^3) \\ &\quad + \sum_{w=1}^{m} 1(n \geq w) \sum_{i=1}^{n-w+1} \{\mathcal{K}(t, i-1) + \mathcal{K}(t, n-i-w+1) + w\} f(w) \lambda \Delta \\ &\quad + \sum_{w=1}^{m} \sum_{i=(n-w+2) \vee 1}^{n} \mathcal{K}(t,n) f(w) \lambda \Delta \\ &= -\lambda n \Delta \mathcal{K}(t,n) + \sum_{w=1}^{m} 1(n \geq w) \left\{ 2 \sum_{i=0}^{n-w} \mathcal{K}(t,i) + w(n-w+1) \right\} f(w) \lambda \Delta \\ &\quad + \sum_{w=1}^{m} (n \wedge (w-1)) \mathcal{K}(t,n) f(w) \lambda \Delta,\end{aligned}$$

which yields the differential equation:

$$\begin{aligned}\frac{\partial}{\partial t}\mathcal{K}(t,n) &= -\lambda n\mathcal{K}(t,n) + \sum_{w=1}^{m} 1(n \geq w)\left\{2\sum_{i=0}^{n-w}\mathcal{K}(t,i) + w(n-w+1)\right\} f(w)\lambda \\ &\quad + \sum_{w=1}^{m}(n \wedge (w-1))\mathcal{K}(t,n)f(w)\lambda \\ &= -\lambda n\mathcal{K}(t,n) + 2\lambda\sum_{w=1}^{m\wedge n}\sum_{i=0}^{n-w}\mathcal{K}(t,i)f(w) + \lambda\sum_{w=1}^{m\wedge n} w(n-w+1)f(w) \\ &\quad + \sum_{w=1}^{m}(n \wedge (w-1))\mathcal{K}(t,n)f(w)\lambda \\ &= -\lambda n\mathcal{K}(t,n) + 2\lambda\sum_{w=1}^{n}\sum_{i=0}^{n-w}\mathcal{K}(t,i)f(w) \\ &\quad + \lambda\sum_{w=1}^{n} w(n-w+1)f(w) + \lambda\mathcal{K}(t,n)\sum_{w=1}^{m}(n \wedge (w-1))f(w),\end{aligned}$$

and in the third equality we used the observation that since $f(w) = 0$ for $w > 0$, we can replace the upper limit in the sum by n.

To further simplify the expression, exchange the order of the sums to obtain that

$$2\lambda\sum_{w=1}^{n}\sum_{i=0}^{n-w}\mathcal{K}(t,i)f(w) = 2\lambda\sum_{i=0}^{n-1}\mathcal{K}(t,i)\sum_{w=1}^{n-i}f(w) = 2\lambda\sum_{i=0}^{n-1}\mathcal{K}(t,i)F(n-i).$$

Also,

$$
\begin{aligned}
& -n+\sum_{w=1}^{m}(n \wedge (w-1)) f(w) \\
& =-n+\sum_{w=1}^{n \wedge m}(w-1) f(w)+1(m>n) \sum_{w=n+1}^{m} n f(w) \\
& =-n+\sum_{w=1}^{n \wedge m}(w-1)(F(w)-F(w-1))+1(m>n) n(1-F(n)) \\
& =-n+\sum_{w=1}^{n \wedge m} w F(w)-\sum_{w=0}^{(n \wedge m)-1} w F(w)-\sum_{w=1}^{n \wedge m} F(w)+1(m>n) n(1-F(n)) \\
& =-n+(n \wedge m) F(n \wedge m)-\sum_{w=1}^{n \wedge m} F(w)+1(m>n) n(1-F(n)) \\
& =\begin{cases}-\sum_{w=1}^{n} F(w), & n<m, \\ -n+m-\sum_{w=1}^{m} F(w), & n \geq m,\end{cases} \\
& =-\sum_{w=1}^{n} F(w).
\end{aligned}
$$

We conclude that

$$
\frac{1}{\lambda} \frac{\partial \mathcal{K}(t, n)}{\partial t}=-\sum_{i=1}^{n} F(i) \mathcal{K}(t, n)+2 \sum_{i=1}^{n-1} F(n-i) \mathcal{K}(t, i)+\sum_{i=1}^{n} i(n-i+1) f(i) \tag{1.4.4}
$$

with boundary condition $\mathcal{K}(0, n)=0$.

Finally, let $\mathcal{R}(t, n)=n-\mathcal{K}(t, n)$ denote the expected number of undelivered packages at time t. Then, writing Equation (1.4.4) in terms of $\mathcal{R}(t,n)$, we obtain

$$\begin{aligned}
\frac{1}{\lambda}\frac{\partial \mathcal{R}(t,n)}{\partial t} &= -\sum_{i=1}^{n} F(i)\mathcal{R}(t,n) + 2\sum_{i=1}^{n-1} F(n-i)\mathcal{R}(t,i) + n\sum_{i=1}^{n} F(i) \\
&\quad - 2\sum_{i=1}^{n-1} F(n-i)i - \sum_{i=1}^{n} i(n-i+1)f(i) \\
&= -\sum_{i=1}^{n} F(i)\mathcal{R}(t,n) + 2\sum_{i=1}^{n-1} F(n-i)\mathcal{R}(t,i) + n\sum_{j=1}^{n} f(j)(n-j+1) \\
&\quad - \sum_{j=1}^{n-1} f(j)(n-j)(n-j+1) - \sum_{i=1}^{n} i(n-i+1)f(i) \\
&= -\sum_{i=1}^{n} F(i)\mathcal{R}(t,n) + 2\sum_{i=1}^{n-1} F(n-i)\mathcal{R}(t,i).
\end{aligned}$$

The corresponding boundary condition is $\mathcal{R}(0,n) = n$. This completes the proof. ■

After obtaining the differential equation, we use induction to prove Proposition 1.3.1, which provides the explicit solution to the ODE from Theorem 1.

Proof. Proof of Proposition 1.3.1. Note that for any n such that $F(n) = 0$, $\mathcal{R}(t,n) = n$. Thus the boundary condition for $\gamma_{n,j} = 1$ satisfies

$$\mathcal{R}(t,n) = \sum_{j=1}^{n} \gamma_{n,j} e^{-\lambda \sum_{k=1}^{j} F(k)t} = \sum_{j=1}^{n} \gamma_{n,j} = n.$$

It remains to prove the result for n such that $F(n) > 0$. To start, let

$$\phi_n(t) = 2\lambda \sum_{i=1}^{n-1} F(n-i)\mathcal{R}(t,i)$$
$$\text{and} \quad \theta_n = \lambda \sum_{j=1}^{n} \overline{F}(j)$$

then the above differential equation becomes

$$\frac{d\mathcal{R}(t,n)}{dt} + (\lambda n - \theta_n)\mathcal{R}(t,n) = \phi_n(t)$$

with boundary condition $\mathcal{R}(0,n) = n$. We will prove by induction in n that

$$\mathcal{R}(t,n) = \sum_{i=1}^{n} \gamma_{n,i} e^{-(\lambda i - \theta_i)t}, \tag{1.4.5}$$

where

$$\gamma_{n,n} = n - \sum_{i=1}^{n-1} \gamma_{n,i} \quad \text{and} \quad \gamma_{n,i} = 2 \cdot \frac{\sum_{j=1}^{n-i} F(j)\gamma_{n-j,i}}{\sum_{k=i+1}^{n} F(k)}, \quad 1 \le i < n, \qquad \gamma_{1,1} = 1.$$

Suppose now that (1.4.5) holds and consider $\mathcal{K}(t, n+1)$. Solving the differential equation satisfied by $\mathcal{R}(t, n+1)$ directly (see, e.g., [102]), we obtain that

$$\begin{aligned}
&\mathcal{R}(t, n+1) \\
&= e^{-(\lambda(n+1)-\theta_{n+1})t}\left(n+1+\int_0^t \phi_{n+1}(v)e^{v(\lambda(n+1)-\theta_{n+1})}dv\right) \\
&= e^{-(\lambda(n+1)-\theta_{n+1})t}(n+1) + \lambda e^{-(\lambda(n+1)-\theta_{n+1})t}\int_0^t 2\sum_{i=1}^{n} F(n+1-i)\mathcal{R}(v,i)e^{(\lambda(n+1)-\theta_{n+1})v}dv \\
&= e^{-(\lambda(n+1)-\theta_{n+1})t}(n+1) + \lambda e^{-(\lambda(n+1)-\theta_{n+1})t}\int_0^t 2\sum_{i=1}^{n} F(n+1-i)\sum_{j=1}^{i}\gamma_{i,j}e^{-(\lambda j-\theta_j)v}e^{(\lambda(n+1)-\theta_{n+1})v}dv \\
&= e^{-(\lambda(n+1)-\theta_{n+1})t}(n+1) + e^{-(\lambda(n+1)-\theta_{n+1})t}2\lambda\sum_{i=1}^{n} F(n+1-i)\sum_{j=1}^{i}\gamma_{i,j}\cdot\frac{e^{(\lambda(n+1-j)-\theta_{n+1}+\theta_j)t}-1}{\lambda(n+1-j)-\theta_{n+1}+\theta_j} \\
&= e^{-(\lambda(n+1)-\theta_{n+1})t}(n+1) + 2\sum_{i=1}^{n} F(n+1-i)\sum_{j=1}^{i}\gamma_{i,j}\cdot\frac{e^{-(\lambda j-\theta_j)t}-e^{-(\lambda(n+1)-\theta_{n+1})t}}{\sum_{k=j+1}^{n+1}F(k)}.
\end{aligned}$$

where in the third equality we used the induction hypothesis (1.4.5).

It remains to analyze the last expression, for which we exchange the summation order to obtain that

$$\begin{aligned}
\mathcal{R}(t, n+1) &= e^{-(\lambda(n+1)-\theta_{n+1})t}(n+1) + 2\sum_{j=1}^{n}\sum_{i=j}^{n} F(n+1-i)\gamma_{i,j}\cdot\frac{e^{-(\lambda j-\theta_j)t}-e^{-(\lambda(n+1)-\theta_{n+1})t}}{\sum_{k=j+1}^{n+1}F(k)} \\
&= e^{-(\lambda(n+1)-\theta_{n+1})t}(n+1) + 2\sum_{j=1}^{n}\frac{e^{-(\lambda j-\theta_j)t}-e^{-(\lambda(n+1)-\theta_{n+1})t}}{\sum_{k=j+1}^{n+1}F(k)}\sum_{r=1}^{n+1-j}F(r)\gamma_{n+1-r,j} \\
&= e^{-(\lambda(n+1)-\theta_{n+1})t}(n+1) + \sum_{j=1}^{n}\left(e^{-(\lambda j-\theta_j)t}-e^{-(\lambda(n+1)-\theta_{n+1})t}\right)\gamma_{n+1,j} \\
&= e^{-(\lambda(n+1)-\theta_{n+1})t}(n+1) + \sum_{i=1}^{n}\gamma_{n+1,i}e^{-(\lambda i-\theta_i)t} - e^{-(\lambda(n+1)-\theta_{n+1})t}\sum_{i=1}^{n}\gamma_{n+1,i}.
\end{aligned}$$

To complete the proof for $\mathcal{R}(t, n)$ note that

$$\sum_{i=1}^{n}\gamma_{n+1,i} = n+1-\gamma_{n+1,n+1}$$

gives

$$\begin{aligned}\mathcal{R}(t,n+1) &= e^{-(\lambda(n+1)-\theta_{n+1})t}(n+1) + \sum_{i=1}^{n} \gamma_{n+1,i} e^{-(\lambda i-\theta_i)t} - e^{-(\lambda(n+1)-\theta_{n+1})t}(n+1) \\ &\quad + e^{-(\lambda(n+1)-\theta_{n+1})t}\gamma_{n+1,n+1} \\ &= \sum_{i=1}^{n+1} \gamma_{n+1,i} e^{-(\lambda i-\theta_i)t}.\end{aligned}$$

We now use the explicit expression for $\mathcal{R}(t,n)$ to compute $\mathcal{C}(t,n)$. Recall that T_n^* denotes the time of the first request when we start with n packages arranged on a circle, and B is the size of the corresponding bundle. Moreover, since T_n^* is exponentially distributed with rate λn and B has distribution F, we have

$$\begin{aligned}\mathcal{C}(t,n) &= E[(B+\mathcal{K}(t-T_n^*, n-B))1(T_n^*<t)] \\ &= E\left[(n-\mathcal{R}(t-T_n^*, n-B))\, 1(T_n^*<t)\right] \\ &= nP(T_n^*<t) - \sum_{k=1}^{m} f(k) \sum_{i=1}^{n-k} \gamma_{n-k,i} e^{-\lambda \sum_{j=1}^{i} F(j) t} E\left[e^{\lambda \sum_{j=1}^{i} F(j) T_n^*} 1(T_n^*<t)\right] \\ &= n(1-e^{-\lambda n t}) - \sum_{k=1}^{m}\sum_{i=1}^{n-k} f(k)\gamma_{n-k,i} \frac{n}{n-\sum_{j=1}^{i} F(j)} \left(1-e^{\lambda(\sum_{j=1}^{i} F(j)-n)t}\right) e^{-\lambda \sum_{j=1}^{i} F(j) t} \\ &= n(1-e^{-\lambda n t}) - \sum_{i=1}^{n-1}\sum_{k=1}^{n-i} f(k)\gamma_{n-k,i} \frac{n}{n-\sum_{j=1}^{i} F(j)} \left(e^{-\lambda \sum_{j=1}^{i} F(j) t} - e^{-\lambda n t}\right) \\ &= n - \sum_{i=1}^{n-1} \tilde{\gamma}_{n,i} e^{-\lambda \sum_{j=1}^{i} F(j) t} - \tilde{\gamma}_{n,n} e^{-\lambda n t},\end{aligned}$$

where

$$\tilde{\gamma}_{n,i} = \sum_{k=1}^{n-i} f(k) \frac{\gamma_{n-k,i}}{1-\frac{1}{n}\sum_{j=1}^{i} F(j)} \quad \text{and} \quad \tilde{\gamma}_{n,n} = n - \sum_{i=1}^{n-1} \tilde{\gamma}_{n,i}.$$

■

The following technical result provides a monotonicity property for a function of $\mathcal{R}(t,n)$ that will be needed for the application of a Tauberian theorem in the proof of Theorem 2. Interestingly, $\mathcal{R}(t,n)$ is not generally monotone in n.

Definition 1.4.1 *We say that a sequence $\{a_n : n \geq 1\}$ is eventually increasing (decreasing) if there exists an $n_0 \in \mathbb{N}_+$ such that a_n is increasing (decreasing) for all $n \geq n_0$.*

Lemma 1.4.2 *For any $t \geq 0$,*

$$\mathcal{K}(t,n) + n \cdot \frac{m+1}{m-\mu+1}$$

is monotonically increasing in n when $n \geq m$, where $\mu = E[B]$ and m is the maximum bundle size.

Proof. Proof. Recall that $\mathcal{K}(t, n)$ is the expected number of packages that can be picked up by time t when we start with n packages arranged on a line. For $n \geq m$, note that arriving requests at location i occur according to a Poisson process with rate λ when $i \leq n - m + 1$ and with rate $\lambda F(n+1-i)$ when $i > n - m + 1$, the latter since close to the right end-point bundle sizes need to be smaller than $n - i$ to be accepted. Furthermore, the time T_n at which the first request is accepted is independent of the location where it occurs, and is exponentially distributed with rate $\lambda(n - m + 1) + \lambda \sum_{i=1}^{m-1} F(i) = \lambda(n - \mu + 1)$. Thus, the probability that the first accepted request occurs at location i, denoted as $p_i^{(n)}$, is

$$p_i^{(n)} = \frac{1}{n - \mu + 1}, \qquad \text{for } 1 \leq i \leq n - m + 1,$$

and

$$p_i^{(n)} = \frac{F(n+1-i)}{n - \mu + 1}, \qquad \text{for } n - m + 1 < i \leq n.$$

Let $\mathcal{Q}(t, k) = E[B + N(t, k - B)|B \leq k]$, where B is the bundle size (distributed according to F) and $N(t, k)$ is the number of packages picked up during $[0, t]$ when k packages are arranged on a line, and note that $\mathcal{Q}(t, k)$ denotes the expected number of packages that will be picked-up during the interval $[0, t]$ given that a request has been accepted at time zero at location 1 of a total of k packages. As in previous proofs, we have dropped the λ from the notation. Let L_n denote the location of the first request to be accepted when we start with n packages. To analyze $\mathcal{K}(t, n)$, we condition on T_n and L_n to obtain

$$\begin{aligned}
\mathcal{K}(t, n) &= E[E[N(t, n)|T_n]] = E\left[\sum_{i=1}^{n} P(L_n = i|T_n) E[N(t, n)|T_n, L_n = i]\right] \\
&= E\left[\sum_{i=1}^{n} p_i^{(n)} \big(E[N((t - T_n)^+, i - 1)|T_n]\right. \\
&\quad \left. + E\left[B + N((t - T_n)^+, n - i + 1 - B)|B \leq n - i + 1, T_n\right]\big)\right] \\
&= E\left[\sum_{i=1}^{n} p_i^{(n)} \left(\mathcal{K}((t - T_n)^+, i - 1) + \mathcal{Q}((t - T_n)^+, n - i + 1)\right)\right] \\
&= \frac{1}{n - \mu + 1} E\left[\sum_{i=1}^{n-m+1} \left(\mathcal{K}((t - T_n)^+, i - 1) + \mathcal{Q}((t - T_n)^+, n - i + 1)\right)\right] \\
&\quad + \sum_{j=1}^{m-1} \frac{F(m - j)}{n - \mu + 1} E\left[\mathcal{K}((t - T_n)^+, n - m + j) + \mathcal{Q}((t - T_n)^+, m - j)\right].
\end{aligned}$$

Therefore, we get the difference

$$
\begin{aligned}
&(n-\mu+2)\mathcal{K}(t,n+1)-(n-\mu+1)\mathcal{K}(t,n)\\
&=E\left[\sum_{i=1}^{n-m+1}\mathcal{K}((t-T_{n+1})^+,i-1)+\sum_{i=2}^{n-m+2}\mathcal{Q}((t-T_{n+1})^+,n-i+2)\right]\\
&\quad+E[\mathcal{K}((t-T_{n+1})^+,n-m+1)+\mathcal{Q}((t-T_{n+1})^+,n+1)]\\
&\quad+\sum_{j=1}^{m-1}F(m-j)E\left[\mathcal{K}((t-T_{n+1})^+,n-m+j+1)+\mathcal{Q}((t-T_{n+1})^+,m-j)\right]\\
&\quad-E\left[\sum_{i=1}^{n-m+1}\big(\mathcal{K}((t-T_n)^+,i-1)+\mathcal{Q}((t-T_n)^+,n-i+1)\big)\right]\\
&\quad-\sum_{j=1}^{m-1}F(m-j)E\left[\mathcal{K}((t-T_n)^+,n-m+j)+\mathcal{Q}((t-T_n)^+,m-j)\right]\\
&=\sum_{i=1}^{n-m+1}\Big(D(t,n,i-1)+\hat{D}(t,n,n-i+1)\Big)+\sum_{k=1}^{m-1}F(k)\hat{D}(t,n,k)\\
&\quad+E[\mathcal{K}((t-T_{n+1})^+,n-m+1)+\mathcal{Q}((t-T_{n+1})^+,n+1)]\\
&\quad+\sum_{k=1}^{m-1}F(k)E\left[\mathcal{K}((t-T_{n+1})^+,n-k+1)\right]-\sum_{k=1}^{m-1}F(k)E\left[\mathcal{K}((t-T_n)^+,n-k)\right]\\
&=\sum_{i=1}^{n-m+1}\Big(D(t,n,i-1)+\hat{D}(t,n,n-i+1)\Big)+\sum_{k=1}^{m-1}F(k)\Big(\hat{D}(t,n,k)+D(t,n,n-k)\Big)\\
&\quad+E[\mathcal{K}((t-T_{n+1})^+,n-m+1)+\mathcal{Q}((t-T_{n+1})^+,n+1)]\\
&\quad+\sum_{k=0}^{m-1}F(k)E\left[\mathcal{K}((t-T_{n+1})^+,n-k+1)-\mathcal{K}((t-T_{n+1})^+,n-k)\right],
\end{aligned}
$$

where

$$
\begin{aligned}
D(t,n,i)&:=E\left[\mathcal{K}((t-T_{n+1})^+,i)-\mathcal{K}((t-T_n)^+,i)\right]\quad\text{and}\\
\hat{D}(t,n,i)&:=E\left[\mathcal{Q}((t-T_{n+1})^+,i)-\mathcal{Q}((t-T_n)^+,i)\right].
\end{aligned}
$$

Furthermore, since $F(0) = 0$ and $F(m) = 1$ we have that

$$
\begin{aligned}
&E[\mathcal{K}((t - T_{n+1})^+, n - m + 1)] + \sum_{k=0}^{m-1} F(k) E\left[\mathcal{K}((t - T_{n+1})^+, n - k + 1) - \mathcal{K}((t - T_{n+1})^+, n - k)\right] \\
&= \sum_{k=0}^{m} F(k) E\left[\mathcal{K}((t - T_{n+1})^+, n - k + 1) - \mathcal{K}((t - T_{n+1})^+, n - k)\right] + E[\mathcal{K}((t - T_{n+1})^+, n - m)] \\
&= \sum_{k=1}^{m} F(k) E\left[\mathcal{K}((t - T_{n+1})^+, n - k + 1)\right] - \sum_{k=1}^{m+1} F(k-1) E\left[\mathcal{K}(t - T_{n+1})^+, n - k + 1)\right] \\
&\quad + E[\mathcal{K}((t - T_{n+1})^+, n - m)] \\
&= \sum_{k=1}^{m} f(k) E\left[\mathcal{K}((t - T_{n+1})^+, n - k + 1)\right].
\end{aligned}
$$

Now use the observation that for any $j \geq m$ and $t \geq 0$ we have

$$
\begin{aligned}
\mathcal{Q}(t, j) &= E\left[B + N(t, j - B) | B \leq j\right] \\
&= E[B] + E[\mathcal{K}(t, j - B)] = \mu + \sum_{k=1}^{m} f(k) \mathcal{K}(t, j - k),
\end{aligned}
$$

to obtain that

$$
E[\mathcal{Q}((t - T_{n+1})^+, n + 1)] = \mu + \sum_{k=1}^{m} f(k) \mathcal{K}((t - T_{n+1})^+, n - k + 1)]
$$

and

$$
\hat{D}(t, n, j) = \sum_{k=1}^{m} f(k) D(t, n, j - k) \qquad \text{for } j \geq m.
$$

We have thus derived that for $n \geq m$,

$$
\begin{aligned}
&(n - \mu + 2)\mathcal{K}(t, n + 1) - (n - \mu + 1)\mathcal{K}(t, n) \\
&= \sum_{i=1}^{n-m+1} \left(D(t, n, i - 1) + \sum_{k=1}^{m} f(k) D(t, n, n - i + 1 - k) \right) + \sum_{i=1}^{m-1} F(i) \left(D(t, n, n - i) + \hat{D}(t, n, i) \right) \\
&\quad + 2 \sum_{i=1}^{m} f(i) E\left[\mathcal{K}((t - T_{n+1})^+, n - i + 1)\right] + \mu. \qquad (1.4.6)
\end{aligned}
$$

Finally, note that $T_{n+1} \leq_{s.t.} T_n$, where (s.t.) denotes the standard stochastic order, which since both $\mathcal{K}$ and $\mathcal{Q}$ are non-decreasing in t, implies that $D(t, n, i) \geq 0$ and $\hat{D}(t, n, j) \geq 0$ for any $i, j \geq 1$. Hence, we immediately obtain that

$$
(n - \mu + 2)\mathcal{K}(t, n + 1) - (n - \mu + 1)\mathcal{K}(t, n) \geq 0.
$$

Dividing by $n - \mu + 1$ now gives

$$\mathcal{K}(t, n+1) - \mathcal{K}(t, n) \geq -\frac{1}{n - \mu + 1} \cdot \mathcal{K}(t, n+1),$$

and using the observation that $\mathcal{K}(t, n+1) \leq n + 1$ further gives

$$\mathcal{K}(t, n+1) - \mathcal{K}(t, n) \geq -\frac{n+1}{n - \mu + 1} \geq -\frac{m+1}{m - \mu + 1}$$

for all $n \geq m$. This in turn implies that

$$\mathcal{K}(t, n) + n \cdot \frac{m+1}{m - \mu + 1}$$

is monotonically increasing with $n \geq m$. ■

Next we calculate the formula for $\beta(t) = 1 - \alpha(t)$ through the generating function $G(t, x) := \sum_{n=m}^{\infty} \mathcal{R}(t, n) x^n$. The key tool in the analysis is the use of a Tauberian theorem that allows us to infer the behavior of $\mathcal{R}(t, n)$ from that of $G(t, x)$. We write $f(x) \sim g(x)$ as $x \to a$ to denote $\lim_{x \to a} f(x)/g(x) = 1$.

Proof. Proof of Theorem 2. Define $\varphi(y) = \sum_{i=1}^{m-1} \overline{F}(i) y^i / i$ and recall that $\theta_n = \lambda \sum_{j=1}^{n} \overline{F}(j)$. Note that when $n \geq m$ we have $\theta_n = \lambda \varphi'(1)$. From Theorem 1 we have

$$\begin{aligned} \frac{\partial \mathcal{R}(t, n)}{\partial t} &= -\lambda \sum_{i=1}^{n} F(i) \mathcal{R}(t, n) + 2\lambda \sum_{i=1}^{n-1} F(n-i) \mathcal{R}(t, i) \\ &= -(\lambda n - \theta_n) \mathcal{R}(t, n) + 2\lambda \sum_{i=1}^{n-1} F(n-i) \mathcal{R}(t, i). \end{aligned}$$

Next, multiply both sides by x^n and sum over n from m to infinity to obtain

$$\begin{aligned}
\sum_{n=m}^{\infty} \frac{\partial \mathcal{R}(t,n)}{\partial t} x^n &= -\sum_{n=m}^{\infty} (\lambda n - \lambda\varphi'(1))\mathcal{R}(t,n)x^n + 2\lambda \sum_{n=m}^{\infty}\sum_{i=1}^{n-1} F(n-i)\mathcal{R}(t,i)x^n \\
&= -\lambda x \sum_{n=m}^{\infty} \mathcal{R}(t,n)\frac{d}{dx}x^n + \lambda\varphi'(1)G(t,x) + 2\lambda \sum_{i=1}^{m-1}\sum_{n=m}^{\infty} F(n-i)\mathcal{R}(t,i)x^n \\
&\quad + 2\lambda \sum_{i=m}^{\infty}\sum_{n=i+1}^{\infty} F(n-i)\mathcal{R}(t,i)x^n \\
&= -\lambda x \frac{\partial}{\partial x}G(t,x) + \lambda\varphi'(1)G(t,x) + 2\lambda \sum_{i=1}^{m-1} \mathcal{R}(t,i)x^i \sum_{j=m-i}^{\infty} F(j)x^j \\
&\quad + 2\lambda \sum_{i=m}^{\infty} \mathcal{R}(t,i)x^i \sum_{j=1}^{\infty} F(j)x^j \\
&= -\lambda x \frac{\partial}{\partial x}G(t,x) + \lambda\varphi'(1)G(t,x) + 2\lambda \sum_{i=1}^{m-1} \mathcal{R}(t,i)x^i \left(\sum_{j=m-i}^{\infty} x^j - \sum_{j=m-i}^{\infty} \overline{F}(j)x^j \right) \\
&\quad + 2\lambda G(t,x) \left(\sum_{j=1}^{\infty} x^j - x\varphi'(x) \right) \\
&= -\lambda x \frac{\partial}{\partial x}G(t,x) + \lambda \left(\varphi'(1) + 2x(1-x)^{-1} - 2x\varphi'(x) \right) G(t,x) \\
&\quad + 2\lambda x^m (1-x)^{-1} \sum_{i=1}^{m-1} \mathcal{R}(t,i) - 2\lambda \sum_{i=1}^{m-1} \mathcal{R}(t,i)x^i \sum_{j=m-i}^{\infty} \overline{F}(j)x^j,
\end{aligned}$$

where the exchange of derivative and series in the third equality is justified by Theorem A.5.1 in [37] and we use the convention that $\sum_{i=a}^{b} x_i \equiv 0$ if $b < a$. Furthermore, Theorem A.5.1 in [37] also gives that

$$\sum_{n=m}^{\infty} \frac{\partial \mathcal{R}(t,n)}{\partial t} x^n = \frac{\partial G(t,x)}{\partial t},$$

and we obtain that $G(t,x)$ satisfies the differential equation:

$$\begin{aligned}
\frac{\partial G(t,x)}{\partial t} &= -\lambda x \frac{\partial}{\partial x}G(t,x) + \lambda \left(\varphi'(1) + 2x(1-x)^{-1} - 2x\varphi'(x) \right) G(t,x) \\
&\quad + 2\lambda x^m (1-x)^{-1} \sum_{i=1}^{m-1} \mathcal{R}(t,i) - 2\lambda \sum_{i=1}^{m-1} \mathcal{R}(t,i)x^i \sum_{j=m-i}^{\infty} \overline{F}(j)x^j.
\end{aligned}$$

To solve it, make the change of variables $r = \ln x - \lambda t$ and $s = t$, and define $\tilde{G}(s,r) =$

$G(s, e^{\lambda s+r})$ to obtain

$$\frac{\partial}{\partial t}G(t,x) = \frac{\partial}{\partial r}\tilde{G}(s,r)(-\lambda) + \frac{\partial}{\partial s}\tilde{G}(s,r), \quad \text{and}$$
$$\frac{\partial}{\partial x}G(t,x) = \frac{\partial}{\partial r}\tilde{G}(s,r)\frac{1}{x}.$$

Substituting in our expression for $\frac{\partial}{\partial t}G(t,x)$ we obtain

$$\begin{aligned}
&\frac{\partial}{\partial r}\tilde{G}(s,r)(-\lambda) + \frac{\partial}{\partial s}\tilde{G}(s,r) \\
&= -\lambda\frac{\partial}{\partial r}\tilde{G}(s,r) + \tilde{G}(s,r)\lambda\left(\varphi'(1) + \frac{2e^{\lambda s+r}}{1-e^{\lambda s+r}} - 2e^{\lambda s+r}\varphi'(e^{\lambda s+r})\right) \\
&\quad + \frac{2\lambda e^{m(\lambda s+r)}}{1-e^{\lambda s+r}}\sum_{i=1}^{m-1}\mathcal{R}(s,i) - 2\lambda\sum_{i=1}^{m-1}\mathcal{R}(s,i)\sum_{j=m-i}^{\infty}\overline{F}(j)e^{(i+j)(\lambda s+r)},
\end{aligned}$$

which by cancelling the terms $\lambda\frac{\partial}{\partial r}\tilde{G}(s,r)$ on both sides can be written as

$$\frac{d}{ds}H(s) + H(s)p(s) = q(s),$$

with $H(s) = \tilde{G}(s,r)$,

$$p(s) = -\lambda\left(\varphi'(1) + \frac{2e^{\lambda s+r}}{1-e^{\lambda s+r}} - 2e^{\lambda s+r}\varphi'(e^{\lambda s+r})\right),$$

and

$$q(s) = \lambda\left(\frac{2e^{m(\lambda s+r)}}{1-e^{\lambda s+r}}\sum_{i=1}^{m-1}\mathcal{R}(s,i) - 2\sum_{i=1}^{m-1}\mathcal{R}(s,i)\sum_{j=m-i}^{m-1}\overline{F}(j)e^{(i+j)(\lambda s+r)}\right).$$

The solution of this ODE (see [102] P.8 equation (1.18)) is

$$H(s) = e^{-\int_0^s p(h)dh}\left(C + \int_0^s q(v)e^{\int_0^v p(h)dh}dv\right),$$

where C is a constant determined by the boundary conditions. Since $H(0) = \tilde{G}(0,r) = G(0,e^r) = \sum_{n=m}^{\infty} ne^{rn}$, we have

$$H(s) = e^{-\int_0^s p(h)dh}\left(H(0) + \int_0^s q(v)e^{\int_0^v p(h)dh}dv\right).$$

Substituting the expression for $p(s)$ gives

$$\begin{aligned}
-\int_0^s p(h)dh &= \lambda \int_0^s \left(\varphi'(1) + \frac{2e^{\lambda h + r}}{1 - e^{\lambda h + r}} - 2e^{\lambda h + r}\varphi'(e^{\lambda h + r})\right) dh \\
&= \lambda\varphi'(1)s + 2\int_{e^r}^{e^{\lambda s + r}} \left(\frac{1}{1-y} - \varphi'(y)\right) dy \\
&= \lambda\varphi'(1)s + 2\left(\ln(1 - e^r) - \ln(1 - e^{r+\lambda s})\right) - 2\left(\varphi(e^{r+\lambda s}) - \varphi(e^r)\right).
\end{aligned}$$

This in turn implies that if we define $h(s, x, y) = (1 - y)^2 e^{-2(\varphi(x) - \varphi(y)) + \lambda\varphi'(1)s}$, then

$$e^{-\int_0^s p(h)dh} = \frac{(1 - e^r)^2}{(1 - e^{r+\lambda s})^2} e^{-2\left(\varphi(e^{r+\lambda s}) - \varphi(e^r)\right) + \lambda\varphi'(1)s} = \frac{1}{(1 - e^{r+\lambda s})^2} \cdot h(s, e^{r+\lambda s}, e^r).$$

Define also

$$\tilde{q}(x, y) := 2y^m(1 - y)\sum_{i=1}^{m-1} \mathcal{R}(\ln(y/x)/\lambda, i) - 2(1 - y)^2 \sum_{i=1}^{m-1} \mathcal{R}(\ln(y/x)/\lambda, i) \sum_{j=m-i}^{m-1} \overline{F}(j)y^{i+j}.$$

and note that $q(v) = \tilde{q}(e^r, e^{r+\lambda v})/(1 - e^{r+\lambda v})^2$. It follows that

$$\begin{aligned}
\tilde{G}(s, r)(1 - e^{r+\lambda s})^2 &= H(s)(1 - e^{r+\lambda s})^2 \\
&= h(s, e^{r+\lambda s}, e^r)\left(H(0) + \int_0^s \lambda\tilde{q}(e^r, e^{r+\lambda v}) \cdot \frac{1}{h(v, e^{r+\lambda v}, e^r)}\, dv\right) \\
&= h(s, e^{r+\lambda s}, e^r)\left(\sum_{n=m}^{\infty} ne^{rn} + \int_{e^{-\lambda s}}^1 \frac{\tilde{q}(e^r, e^{r+\lambda s}u)}{uh(s + (\ln u)/\lambda, e^{r+\lambda s}u, e^r)}\, du\right).
\end{aligned}$$

Substituting $t = s$ and $x = e^{r+\lambda s}$ we obtain:

$$\begin{aligned}
G(t, x)(1 - x)^2 &= h(t, x, xe^{-\lambda t})\left(\sum_{n=m}^{\infty} n(xe^{-\lambda t})^n + \int_{e^{-\lambda t}}^1 \frac{\tilde{q}(xe^{-\lambda t}, xu)}{uh(t + (\ln u)/\lambda, xu, xe^{-\lambda t})} du\right) \\
&= e^{-2(\varphi(x) - \varphi(xe^{-\lambda t})) + \lambda\varphi'(1)t}(xe^{-\lambda t})^m \left(m - (m - 1)xe^{-\lambda t}\right) \\
&\quad + e^{-2\varphi(x)}x^{\varphi'(1)+1}\int_{e^{-\lambda t}}^1 e^{2\varphi(xu)}\hat{q}(t + (\ln u)/\lambda, xu)\, du, \qquad (1.4.7)
\end{aligned}$$

where

$$\hat{q}(s, v) := 2v^{m-\varphi'(1)-1}(1 - v)\sum_{i=1}^{m-1} \mathcal{R}(s, i) - 2(1 - v)^2 \sum_{i=1}^{m-1} \mathcal{R}(s, i) \sum_{j=m-i}^{m-1} \overline{F}(j)v^{i+j-\varphi'(1)-1}.$$

Now take the limit as $x \nearrow 1$ to obtain:

$$\begin{aligned}
\beta(t) &:= \lim_{x\to 1} G(t, x)(1 - x)^2 \qquad (1.4.8) \\
&= e^{-2(\varphi(1) - \varphi(e^{-\lambda t})) + \lambda\varphi'(1)t}(e^{-\lambda t})^m \left(m - (m - 1)e^{-\lambda t}\right) + e^{-2\varphi(1)}\int_{e^{-\lambda t}}^1 e^{2\varphi(u)}\hat{q}(t + (\ln u)/\lambda, u)\, du.
\end{aligned}$$

Finally, to obtain the first statement of the theorem we use a Tauberian theorem to infer the asymptotic behavior of $\mathcal{R}(t,n)$ from that of $G(t,x)$. To this end, define

$$S(t,x) := \sum_{n=m}^{\infty} (n - \mathcal{R}(t,n) + n\kappa)\, x^n = -G(t,x) + c \cdot \frac{(1-x)mx^m + x^{m+1}}{(1-x)^2},$$

where $\kappa = (m+1)/(m-\mu+1)$, $c = \kappa + 1$ and $\mu = \varphi'(1) - 1$. From (1.4.8) we obtain

$$S(t,x) \sim (c - \beta(t))\, \frac{1}{(1-x)^2} \qquad \text{as } x \nearrow 1.$$

Moreover, by Lemma 1.4.2 we have that $cn - \mathcal{R}(t,n) = \mathcal{K}(t,n) + n\kappa$ is monotonically increasing with n, so Theorem 8.3 in [46] yields

$$\mathcal{K}(t,n) + n(c-1) = cn - \mathcal{R}(t,n) \sim n\,(c - \beta(t)) \qquad \text{as } n \to \infty,$$

which implies that

$$\lim_{n\to\infty} \frac{\mathcal{K}(t,n)}{n} = 1 - \beta(t) = \alpha(t).$$

To obtain the asymptotic behavior of $\mathcal{C}(t,n)$ recall that T_n^* denotes the time of the first request when we start with n packages arranged on a circle and B is the corresponding bundle size. Then, use Lemma 1.4.2 to get

$$\begin{aligned}
\mathcal{C}(t,n) &= E[B + \mathcal{K}(t - T_n^*, n - B)1(T_n^* < t)] \\
&\le \mu + E[\mathcal{K}(t, n-B)] \\
&= \mu + E[\mathcal{K}(t, n-B) + \kappa(n-B)] - \kappa(n-\mu) \\
&\le \mu + \mathcal{K}(t,n) + \kappa n - \kappa(n-\mu) \qquad \text{(by Lemma 1.4.2)} \\
&\le \mu + \mathcal{K}(t,n) + m\kappa.
\end{aligned} \tag{1.4.9}$$

Note that the expected number of packages that can be picked up during time Δ, when the total number of packages is n, is $\lambda n \Delta m$, so

$$\mathcal{K}(t+\Delta, n) \le \mathcal{K}(t,n) + m\lambda n\Delta. \tag{1.4.10}$$

Therefore, using Lemma 1.4.2 again and the observation that $E[T_n^*] = 1/(\lambda n)$ (since T_n^* is exponentially distributed with rate λn), we have

$$\begin{aligned}
\mathcal{C}(t,n) &\ge E[\mathcal{K}((t - T_n^*)^+, n - B)] \\
&= E[\mathcal{K}((t - T_n^*)^+, n - B) + (n-B)\kappa] - (n-\mu)\kappa \\
&\ge E[\mathcal{K}((t - T_n^*)^+, n - m) + (n-m)\kappa] - (n-\mu)\kappa \qquad \text{(by Lemma 1.4.2)} \\
&\ge E[\mathcal{K}(t, n-m) - m\lambda(n-m)T_n^*] - (m-\mu)\kappa \qquad \text{(by (1.4.10))} \\
&\ge \mathcal{K}(t, n-m) - m - m\kappa.
\end{aligned} \tag{1.4.11}$$

Combining (1.4.9) and (1.4.11) we obtain

$$\lim_{n\to\infty} \frac{\mathcal{C}(t,n)}{n} = \lim_{n\to\infty} \frac{\mathcal{K}(t,n)}{n} = \alpha(t).$$

This completes the proof. ∎

The last result in Section 1.3.1.1 is Theorem 3, which gives the convergence rate of $\mathcal{K}(t,n)/n$ to $\alpha(t)$. The proof will be based again on the use of a Tauberian theorem, however, the monotonicity required is more difficult to verify. To ease the reading of the proof we first state a couple of preliminary technical results.

Proposition 1.4.3 *Define $c_n(t) = \mathcal{K}(t,n) - n\alpha(t)$, where $\alpha(t)$ is defined as in Theorem 2. Then, for any $t \geq 0$, we have that $c_n(t)$ is either: i) bounded, ii) positive and eventually increasing, or iii) negative and eventually decreasing in n.*

Proof. Proof. Define T_n, $D(t,n,i)$ and $\hat{D}(t,n,j)$ as in the proof of Lemma 1.4.2, and let $\kappa = (m+1)/(m-\mu+1)$. Recall that $D(t,n,i) \geq 0$ and $\hat{D}(t,n,j) \geq 0$. Next, note that

$$E[T_{n+1}] = \frac{1}{(n-\mu+2)\lambda},$$

and use equation (1.4.6) to obtain

$$\begin{aligned}
&(n-\mu+2)\mathcal{K}(t,n+1) - (n-\mu+1)\mathcal{K}(t,n) \\
&\geq 2\sum_{i=1}^{m} f(i) E\left[\mathcal{K}((t-T_{n+1})^+, n+1-i)\right] \\
&\geq 2\sum_{i=1}^{m} f(i)\left(E\left[\mathcal{K}((t-T_{n+1})^+, n+1-m)\right] - (n+1-m)\kappa\right) + 2\sum_{i=1}^{m} f(i)(n+1-i)\kappa \\
&= 2E\left[\mathcal{K}((t-T_{n+1})^+, n+1-m)\right] + 2m\kappa - 2\kappa\mu \\
&\geq 2E\left[\mathcal{K}(t,n+1-m) - m\lambda(n+1-m)T_{n+1}\right] \\
&\geq 2\mathcal{K}(t,n+1-m) - 2m,
\end{aligned} \tag{1.4.12}$$

where in the second inequality we used Lemma 1.4.2, and in the third inequality we used (1.4.10).

To obtain an upper bound note that (1.4.10) also gives

$$\begin{aligned}
D(t,n,i) &= E\left[\mathcal{K}((t-T_n+T_n-T_{n+1})^+, i) - \mathcal{K}((t-T_n)^+, i)\right] \\
&\leq E\left[m\lambda i(T_n - T_{n+1})\right] = m\lambda i\left(\frac{1}{\lambda(n-\mu+1)} - \frac{1}{\lambda(n-\mu+2)}\right) \\
&= \frac{mi}{(n-\mu+1)(n-\mu+2)},
\end{aligned}$$

where we have used the observation that T_n is exponentially distributed with rate $\lambda(n-\mu+1)$. The same arguments also yield

$$\hat{D}(t,n,j) \leq \frac{mj}{(n-\mu+1)(n-\mu+2)}.$$

Substituting these estimates into (1.4.6) now gives

$$\begin{aligned}
&(n-\mu+2)\mathcal{K}(t,n+1) - (n-\mu+1)\mathcal{K}(t,n) \\
&\leq \sum_{i=1}^{n-m+1} \left(\frac{m(i-1)}{(n-\mu+1)(n-\mu+2)} + \sum_{k=1}^{m} f(k) \frac{m(n-i+1-k)}{(n-\mu+1)(n-\mu+2)} \right) \\
&\quad + \sum_{i=1}^{m-1} F(i) \left(\frac{m(n-i)}{(n-\mu+1)(n-\mu+2)} + \frac{mi}{(n-\mu+1)(n-\mu+2)} \right) \\
&\quad + 2\sum_{i=1}^{m} f(i) E\left[\mathcal{K}((t-T_{n+1})^+, n-i+1)\right] + \mu \\
&= \frac{m}{(n-\mu+1)(n-\mu+2)} \left((n-m+1)(n-\mu) + n(m-\mu)\right) \\
&\quad + 2\sum_{i=1}^{m} f(i) E\left[\mathcal{K}((t-T_{n+1})^+, n-i+1) + \kappa(n-i+1)\right] - 2\sum_{i=1}^{m} f(i)\kappa(n-i+1) + \mu \\
&\leq m + \frac{nm(m-\mu)}{(n-\mu+1)(n-\mu+2)} + 2E\left[\mathcal{K}((t-T_{n+1})^+, n) + \kappa n\right] - 2\kappa(n-\mu+1) + \mu \\
&\leq 2\mathcal{K}(t,n) + 3m(\kappa+1), \qquad (1.4.13)
\end{aligned}$$

where in the second inequality we used Lemma 1.4.2. Hence, combining (1.4.12) and (1.4.13) we obtain that

$$2\mathcal{K}(t,n+1-m) - \mathcal{K}(t,n) - 2m \leq (n-\mu+2)\left(\mathcal{K}(t,n+1) - \mathcal{K}(t,n)\right) \leq \mathcal{K}(t,n) + 3m(\kappa+1). \qquad (1.4.14)$$

Dividing by $n-\mu+2$, taking the limit as $n \to \infty$, and using Theorem 2, we obtain

$$\alpha(t) = \lim_{n\to\infty} \frac{2\mathcal{K}(t,n+1-m) - \mathcal{K}(t,n)}{n-\mu+2} \leq \lim_{n\to\infty} \left(\mathcal{K}(t,n+1) - \mathcal{K}(t,n)\right) \leq \lim_{n\to\infty} \frac{\mathcal{K}(t,n)}{n-\mu+2} = \alpha(t).$$

This in turn implies that

$$\lim_{n\to\infty} c_{n+1}(t) - c_n(t) = 0 \qquad (1.4.15)$$

for all $t \geq 0$. Furthermore, (1.4.14) can be written as:

$$\frac{2c_{n+1-m}(t) - c_n(t) - \alpha(t)(2m-\mu) - 2m}{n-\mu+2} \leq c_{n+1}(t) - c_n(t) \leq \frac{c_n(t) + \alpha(t)(\mu-2) + 3m(\kappa+1)}{n-\mu+2},$$

which combined with (1.4.15) gives that there exists $n_0 \in \mathbb{N}_+$ such that

$$c_{n-m+1}(t) \geq c_n(t) - m/2 \text{ for all } n \geq n_0.$$

This in turn implies that

$$\frac{c_n(t) - \alpha(t)(2m-\mu) - 3m}{n-\mu+2} \leq c_{n+1}(t) - c_n(t) \leq \frac{c_n(t) + \alpha(t)(\mu-2) + 3m(\kappa+1)}{n-\mu+2}$$

for all $n \geq n_0$.

To complete the proof, note that if there exists an $n' \geq n_0$ such that $c_{n'}(t) + \alpha(t)(\mu - 2) + 3m(\kappa+1) < 0$, then $c_{n'}(t) < 0$ and $c_{n'+1}(t) < c_{n'}(t)$, which in turn implies that,

$$c_{n'+1}(t) + \alpha(t)(\mu-2) + 3m(\kappa+1) < c_{n'}(t) + \alpha(t)(\mu-2) + 3m(\kappa+1) < 0.$$

Iterating in this way we obtain that

$$c_{n'+k}(t) < c_{n'+k-1}(t) < \cdots < c_{n'}(t) < 0$$

for all $k \geq 1$. This gives condition (iii) of the proposition.

Suppose now that there exists $n' \geq n_0$ such that $c_{n'}(t) - \alpha(t)(2m-\mu) - 3m > 0$. Similarly as above, we obtain that $c_{n'}(t) > 0$ and $c_{n'+1}(t) > c_{n'}(t)$, which in turn implies that,

$$c_{n'+1}(t) - \alpha(t)(2m-\mu) - 3m > c_{n'}(t) - \alpha(t)(2m-\mu) - 3m > 0.$$

Iterating as before, we obtain that

$$c_{n'+k}(t) > c_{n'+k-1}(t) > \cdots > c_{n'}(t) > 0$$

for all $k \geq 1$. This gives condition (ii) of the proposition.

Finally, neither of the two previous cases occurs we have that

$$-\alpha(t)(\mu-2) - 3m(\kappa+1) \leq c_n(t) \leq \alpha(t)(2m-\mu) + 3m$$

for all $n \geq n_0$, and condition (i) follows. ■

The second preliminary result provides the differentiability of the functions that determine $\alpha(t)$.

Lemma 1.4.4 *Define*

$$\omega_t(x) := e^{-2(\varphi(x)-\varphi(xe^{-\lambda t}))-\lambda t(m-\varphi'(1))} x^m \left(m - (m-1)xe^{-\lambda t}\right)$$
$$+ e^{-2\varphi(x)} x^{\varphi'(1)+1} \int_{e^{-\lambda t}}^{1} e^{2\varphi(xu)} \hat{q}(t + (\ln u)/\lambda, xu)\, du,$$

where

$$\hat{q}(s,v) := 2v^{m-\varphi'(1)-1}(1-v)\sum_{i=1}^{m-1}\mathcal{R}(s,i) - 2(1-v)^2\sum_{i=1}^{m-1}\mathcal{R}(s,i)\sum_{j=m-i}^{m-1}\overline{F}(j)v^{i+j-\varphi'(1)-1}.$$

Then, for any fixed $t \geq 0$, $\omega_t(x)$ is infinitely differentiable on $[0,\infty)$ and

$$\sup_{0\leq x\leq 1}|\omega_t''(x)| \leq H_F$$

for some constant $H_F < \infty$ that depends only on the distribution F.

Proof. Proof. Start by defining

$$\begin{aligned} p(s,v) := \hat{q}(s,v)v^{\varphi'(1)+1-m} &= 2(1-v)\sum_{i=1}^{m-1}\mathcal{R}(s,i) - 2(1-v)^2\sum_{i=1}^{m-1}\mathcal{R}(s,i)\sum_{k=0}^{i-1}\overline{F}(k-i+m)v^k \\ &= 2(1-v)\sum_{i=1}^{m-1}\mathcal{R}(s,i) - 2(1-v)^2\sum_{k=0}^{m-2}v^k\sum_{i=k+1}^{m-1}\mathcal{R}(s,i)\overline{F}(k-i+m) =: \sum_{k=0}^{m}c_k(s)v^k. \end{aligned} \tag{1.4.16}$$

Note that $p(s,v)$ is a polynomial of order m in v, and $x^{\varphi'(1)+1}\hat{q}(s,xu) = x^m u^{\varphi'(1)+1-m}p(s,xu)$, from where it follows that $x^{\varphi'(1)+1}\hat{q}(s,xu)$ is infinitely differentiable in x on $[0,\infty)$. Since $\varphi(y) = \sum_{i=1}^{m-1}\overline{F}(i)y^i/i$ is also infinitely differentiable on the real line, then $\omega_t(x)$ is infinitely differentiable in x on $(0,\infty)$. Moreover, by writing

$$\begin{aligned} \omega_t(x) &= a(t)e^{2(\varphi(xe^{-\lambda t})-\varphi(x))}x^m - b(t)e^{2(\varphi(xe^{-\lambda t})-\varphi(x))}x^{m+1} \\ &\quad + \int_{e^{-\lambda t}}^1 e^{2(\varphi(xu)-\varphi(x))}\sum_{k=0}^m c_k(t+(\ln u)/\lambda)u^{m-\varphi'(1)-1+k}x^{m+k}du, \end{aligned}$$

with $a(t) := me^{-\lambda t(m-\varphi'(1))}$ and $b(t) := (m-1)e^{-\lambda t(m-\varphi'(1)+1)}$, we obtain that

$$\begin{aligned} \omega_t''(x) &= a(t)\frac{\partial^2}{\partial x^2}e^{2(\varphi(xe^{-\lambda t})-\varphi(x))}x^m - b(t)\frac{\partial^2}{\partial x^2}e^{2(\varphi(xe^{-\lambda t})-\varphi(x))}x^{m+1} \\ &\quad + \sum_{k=0}^m\int_{e^{-\lambda t}}^1 c_k(t+(\ln u)/\lambda)u^{m-\varphi'(1)-1+k}\frac{\partial^2}{\partial x^2}e^{2(\varphi(xu)-\varphi(x))}x^{m+k}du \\ &= a(t)e^{2(\varphi(xe^{-\lambda t})-\varphi(x))}\nu_m(x,e^{-\lambda t}) - b(t)e^{2(\varphi(xe^{-\lambda t})-\varphi(x))}\nu_{m+1}(x,e^{-\lambda t}) \\ &\quad + \sum_{k=0}^m\int_{e^{-\lambda t}}^1 c_k(t+(\ln u)/\lambda)u^{m-\varphi'(1)-1+k}e^{2(\varphi(xu)-\varphi(x))}\nu_{m+k}(x,u)\,du, \end{aligned}$$

where

$$\begin{aligned} \nu_k(x,u) &= 4(\varphi'(xu)u-\varphi'(x))^2x^k + 4(\varphi'(xu)u-\varphi'(x))kx^{k-1}1(k\geq 1) \\ &\quad + 2(\varphi''(xu)u^2-\varphi''(x))x^k + k(k-1)x^{k-2}1(k\geq 2). \end{aligned}$$

Noting that $\varphi(y), \varphi'(y)$, and $\varphi''(y)$ are all increasing in y, gives that $e^{2(\varphi(xu)-\varphi(x))} \leq 1$ and

$$\begin{aligned} |\nu_k(x,u)| &\leq 4\varphi'(x)^2 x^k + 4\varphi'(x) k x^{k-1} 1(k \geq 1) + 2\varphi''(x) x^k + k(k-1) x^{k-2} 1(k \geq 2) \\ &\leq 4\varphi'(1)^2 + 4\varphi'(1) k + 2\varphi''(1) + k(k-1) =: H_k, \end{aligned}$$

for all $x, u \in [0,1]$, which in turn implies that for any $t \geq 0$,

$$\begin{aligned} |\omega_t''(x)| &\leq a(t)|\nu_m(x, e^{-\lambda t})| + b(t)|\nu_{m+1}(x, e^{-\lambda t})| + \sum_{k=0}^{m} \int_{e^{-\lambda t}}^{1} |c_k(t + (\ln u)/\lambda)\nu_{m+i}(x,u)| \, du \\ &\leq a(t) H_m + b(t) H_{m+1} + \sum_{k=0}^{m} H_{m+k} \int_{e^{-\lambda t}}^{1} |c_k(t + (\ln u)/\lambda)| \, du. \end{aligned}$$

To complete the proof note that $a(t) \leq m$ and $b(t) \leq m$ for all $t \geq 0$, and from (1.4.16) it can be verified that:

$$|c_k(s)| \leq 2 \sum_{i=(k-1)\vee 1}^{m-1} \mathcal{R}(s,i) \leq 2 \sum_{i=(k-1)\vee 1}^{m-1} i, \qquad 0 \leq k \leq m,$$

from where we obtain that

$$|\omega_t''(x)| \leq m H_m + m H_{m+1} + 2 \sum_{k=0}^{m} H_{m+k} \sum_{i=(k-1)\vee 1}^{m-1} i =: H_F$$

for all $t \geq 0$ and all $x \in [0,1]$. ■

We can now give the proof of Theorem 3.

Proof. Proof of Theorem 3. Fix t and let $\alpha(t) = \alpha(t, \lambda)$. Define $J(t,x) := \sum_{n=m}^{\infty} (\mathcal{K}(t,n) - n\alpha(t)) x^n$. Recall from the proof of Theorem 2 that $G(t,x) = \sum_{n=m}^{\infty} \mathcal{R}(t,n) x^n$, and note that from (1.4.7) we have that

$$G(t,x)(1-x)^2 = \omega_t(x),$$

where $\omega_t(x)$ is defined in Lemma 1.4.4. Moreover, by Lemma 1.4.4 we have that $\omega_t(x)$ is infinitely differentiable on $(0,\infty)$ and satisfies $\sup_{0 \leq x \leq 1} |\omega_t''(x)| \leq H_F$, for some constant $H_F < \infty$, independent of t. Also, by (1.4.8), we have $\omega_t(1) = \beta(t)$. Next, write $J(t,x)$ as:

$$\begin{aligned} J(t,x) &= -G(t,x) + \beta(t) \sum_{n=m}^{\infty} n x^n = \frac{\beta(t) - (1-x)^2 G(t,x)}{(1-x)^2} + \beta(t) \left(\sum_{n=m}^{\infty} n x^n - (1-x)^{-2} \right) \\ &= \frac{\beta(t) - \omega_t(x)}{(1-x)^2} + \beta(t) \frac{(1-m)x^{m+1} + m x^m - 1}{(1-x)^2}. \end{aligned}$$

Since

$$\lim_{x\to 1} \frac{(1-m)x^{m+1} + mx^m - 1}{1-x} = -1,$$

and

$$\omega_t(x) = \beta(t) + \omega_t'(1)(x-1) + O\left((x-1)^2\right) \text{ as } x \nearrow 1,$$

we have

$$\begin{aligned}\lim_{x\nearrow 1}(1-x)J(t,x) &= \lim_{x\nearrow 1} \frac{\beta(t) - \omega_t(x)}{1-x} + \beta(t)\frac{(1-m)x^{m+1} + mx^m - 1}{1-x} \\ &= \omega_t'(1) - \beta(t).\end{aligned}$$

Therefore, we have

$$G\sum_{n=0}^{\infty} \left(\mathcal{K}(t,n) - n\alpha(t)\right) x^n \sim \frac{\omega_t'(1) - \beta(t)}{1-x} \text{ as } x \uparrow 1.$$

From Proposition 1.4.3, $\mathcal{K}(t,n) - n\alpha(t)$ is either bounded, or eventually negative and decreasing, or eventually positive and increasing. If $\mathcal{K}(t,n) - n\alpha(t)$ is bounded, then

$$\left|\frac{\mathcal{K}(t,n)}{n} - \alpha(t)\right| = O(n^{-1}) \text{ as } n \to \infty$$

and the proof is complete. If $\mathcal{K}(t,n) - n\alpha(t)$ is eventually negative and decreasing, or eventually positive and increasing, then from the Tauberian Theorem (Theorem 8.3 in [46]), we have

$$\mathcal{K}(t,n) - n\alpha(t) \sim \frac{\omega_t'(1) - \beta(t)}{\Gamma(1)},$$

where $\Gamma(1) = 0!$ is the gamma function. It follows that

$$\left|\frac{\mathcal{K}(t,n)}{n} - \alpha(t)\right| = O(n^{-1}) \text{ as } n \to \infty.$$

To complete the proof, use inequalities (1.4.9) and (1.4.11) to obtain that

$$\mathcal{K}(t,n-m) - m - m\kappa \le \mathcal{C}(t,n) \le \mu + \mathcal{K}(t,n) + m\kappa,$$

where $\kappa = (m+1)/(m-\mu+1)$. and conclude that

$$\left|\frac{\mathcal{C}(t,n)}{n} - \alpha(t)\right| = O(n^{-1}) \text{ as } n \to \infty.$$

■

This completes the proofs of the main theoretical results from Section 1.3.1.1. We now proceed to prove the two lemmas is Section 1.3.2. We start with the proof of Lemma 1.3.2, which establishes the approximate optimality of $\hat{z}$.

Proof. Proof of Lemma 1.3.2. We start by defining the auxiliary cost function

$$\begin{aligned}\frac{\mathrm{Cost}''_P(z;\mathbf{x}^{(n)})}{n} := &\frac{\mathcal{C}(T,n,\lambda(z))}{n}\left(\zeta_P+\frac{h_P+z}{v_P}\right)\left(\frac{\overline{r}(n)}{E[B]}+\frac{L(\mathrm{TSP}(\mathbf{x}^{(n)}))}{n}\right)\\ &+\frac{\mathcal{C}(T,n,\lambda(z))}{n}(h_P+z)\tau_P+\left(1-\frac{\mathcal{C}(T,n,\lambda(z))}{n}\right)h_V\tau_V\\ &+\left(\zeta_V+\frac{h_V}{v_V}\right)\left(\frac{2(1-\mathcal{C}(T,n,\lambda(z))/n)\overline{r}(n)}{V}+\frac{\beta_{\mathrm{VRP}}}{\sqrt{n}}\sqrt{(1-\mathcal{C}(T,n,\lambda(z))/n)A}\right).\end{aligned}$$

Note that the only difference between $\mathrm{Cost}''_P(z;\mathbf{x}^{(n)})$ and $\mathrm{Cost}'_P(z;\mathbf{x}^{(n)})$ is that we have replaced $\mathcal{C}(T,n,\lambda(z))/n$ with $\alpha(T,\lambda(z))$. Recall that $z^*=\mathrm{argmin}_z\,\mathrm{Cost}_P(z;\mathbf{x}^{(n)})$ and $\hat{z}=\mathrm{argmin}_z\,\mathrm{Cost}'_P(z;\mathbf{x}^{(n)})$. It follows that

$$\begin{aligned}0 &\le \frac{\mathrm{Cost}_P(\hat{z};\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}_P(z^*;\mathbf{x}^{(n)})}{n}\\ &\le \left|\frac{\mathrm{Cost}_P(\hat{z};\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}''_P(\hat{z};\mathbf{x}^{(n)})}{n}\right|+\left|\frac{\mathrm{Cost}_P(z^*;\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}''_P(z^*;\mathbf{x}^{(n)})}{n}\right|\\ &\quad+\frac{\mathrm{Cost}''_P(\hat{z};\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}''_P(z^*;\mathbf{x}^{(n)})}{n}.\end{aligned}$$

Now note that for any z we have

$$\begin{aligned}\left|\frac{\mathrm{Cost}_P(z;\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}''_P(z;\mathbf{x}^{(n)})}{n}\right| &= \left(\zeta_V+\frac{h_V}{v_V}\right)\\ \times\left|\frac{1}{n}\mathbb{E}_n\left[L(\mathrm{CVRP}(\mathbf{Y}^{(n)}(\lambda)))\right]-\frac{2(1-\mathcal{C}(T,n,\lambda(z))/n)\overline{r}(n)}{V}\right.&\left.-\frac{\beta_{\mathrm{VRP}}}{\sqrt{n}}\sqrt{(1-\mathcal{C}(T,n,\lambda(z))/n)A}\right|.\end{aligned}$$

Using the deterministic bound $\left|L(\mathrm{CVRP}(\mathbf{y}^{(k)}))-2k\overline{r}(k)/V\right|\le L(\mathrm{TSP}(\mathbf{y}^{(k)}))+2\overline{r}(k)$ for any k points $\{\mathbf{y}_1,\ldots,\mathbf{y}_k\}$, regardless of how they are distributed (see Theorem 8 in the Appendix), and the asymptotic length of an optimal TSP route (see Theorem 9 in the Appendix), we obtain that

$$\left|\frac{\mathrm{Cost}_P(z;\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}''_P(z;\mathbf{x}^{(n)})}{n}\right|\le\frac{1}{n}\mathbb{E}_n\left[L(\mathrm{TSP}(\mathbf{Y}^{(n)}(\lambda)))\right]+O(n^{-1})=O(n^{-1/2})$$

as $n\to\infty$, uniformly in z. Now use the optimality of $\hat{z}$ and Theorem 3 to obtain that

$$\begin{aligned}&\frac{\mathrm{Cost}''_P(\hat{z};\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}''_P(z^*;\mathbf{x}^{(n)})}{n}\\ &\le\left|\frac{\mathrm{Cost}''_P(\hat{z};\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}'_P(\hat{z};\mathbf{x}^{(n)})}{n}\right|+\left|\frac{\mathrm{Cost}''_P(z^*;\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}'_P(z^*;\mathbf{x}^{(n)})}{n}\right|\\ &\quad+\frac{\mathrm{Cost}'_P(\hat{z};\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}'_P(z^*;\mathbf{x}^{(n)})}{n}\\ &\le\left|\frac{\mathrm{Cost}''_P(\hat{z};\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}'_P(\hat{z};\mathbf{x}^{(n)})}{n}\right|+\left|\frac{\mathrm{Cost}''_P(z^*;\mathbf{x}^{(n)})}{n}-\frac{\mathrm{Cost}'_P(z^*;\mathbf{x}^{(n)})}{n}\right|=O(n^{-1})\end{aligned}$$

as $n \to \infty$. We conclude that

$$0 \leq \frac{\text{Cost}_P(\hat{z}; \mathbf{x}^{(n)})}{n} - \frac{\text{Cost}_P(z^*; \mathbf{x}^{(n)})}{n} \leq O(n^{-1/2})$$

as $n \to \infty$. ■

The last proof in the paper is that of Lemma 1.3.3, which compares the expected cost of of using a strategy which combines private drivers with in-house vans with the strategy of using only vans.

Proof. Proof of Lemma 1.3.3. To start, use Theorem 8 to obtain that for any z and $\lambda := \lambda(z)$ we have

$$\begin{aligned}
&\text{Cost}_P(z; \mathbf{x}^{(n)}) - \text{Cost}_V(\mathbf{x}^{(n)}) \\
&= \mathcal{C}(T, n, \lambda) \left(\zeta_P + \frac{h_P + z}{v_P}\right) \left(\frac{\overline{r}(n)}{E[B]} + \frac{1}{n} L(\text{TSP}(\mathbf{x}^{(n)}))\right) \\
&\quad + \mathcal{C}(T, n, \lambda) \left((h_P + z)\tau_P - h_V \tau_V\right) \\
&\quad + \left(\zeta_V + \frac{h_V}{v_V}\right) \left\{\mathbb{E}_n \left[L(\text{CVRP}(\mathbf{Y}^{(n)}(\lambda^*)))\right] - L(\text{CVRP}(\mathbf{x}^{(n)}))\right\} \\
&\leq \mathcal{C}(T, n, \lambda) \left(\left(\zeta_P + \frac{h_P + z}{v_P}\right) \frac{r^*}{E[B]} + (h_P + z)\tau_P - h_V \tau_V\right) \\
&\quad + \mathcal{C}(T, n, \lambda) \left(\zeta_P + \frac{h_P + z}{v_P}\right) \left(\frac{\overline{r}(n) - r^*}{E[B]} + \frac{1}{n} L(\text{TSP}(\mathbf{x}^{(n)}))\right) \\
&\quad + \left(\zeta_V + \frac{h_V}{v_V}\right) \left\{\mathbb{E}_n \left[2 \left(\frac{n - \tilde{N}(T, n, \lambda)}{V} + 1\right) \mathscr{R}(n - \tilde{N}(T, n, \lambda)) + L(\text{TSP}(\mathbf{Y}^{(n)}(\lambda)))\right]\right. \\
&\quad \left. - \max\left\{\frac{2n\overline{r}(n)}{V}, \, L(\text{TSP}(\mathbf{x}^{(n)}))\right\}\right\},
\end{aligned}$$

where $\mathscr{R}(n - \tilde{N}(T, n, \lambda)) = (n - \tilde{N}(T, n, \lambda))^{-1} \sum_{i=1}^{n - \tilde{N}(T,n,\lambda)} d(\mathbf{Y}_i, \mathbf{0})$. Next, divide by n, take the limit as $n \to \infty$, and use Theorem 2 to obtain that

$$\begin{aligned}
&\limsup_{n \to \infty} \frac{1}{n} \left(\text{Cost}_P(z; \mathbf{x}^{(n)}) - \text{Cost}_V(\mathbf{x}^{(n)})\right) \\
&\leq \alpha(T, \lambda) \left(\left(\zeta_P + \frac{h_P + z}{v_P}\right) \frac{r^*}{E[B]} + (h_P + z)\tau_P - h_V \tau_V\right) + \alpha(T, \lambda) \limsup_{n \to \infty} \frac{1}{n} L(\text{TSP}(\mathbf{x}^{(n)})) \\
&\quad + \left(\zeta_V + \frac{h_V}{v_V}\right) \left\{\limsup_{n \to \infty} \mathbb{E}_n \left[2 \left(\frac{n - \tilde{N}(T, n, \lambda)}{Vn} + \frac{1}{n}\right) \mathscr{R}(n - N(T, n, \lambda))\right]\right. \\
&\quad \left. + \limsup_{n \to \infty} \mathbb{E}_n \left[\frac{L(\text{TSP}(\mathbf{Y}^{(n)}(\lambda)))}{n}\right] - \max\left\{\frac{2r^*}{V}, \, \liminf_{n \to \infty} \frac{1}{n} L(\text{TSP}(\mathbf{x}^{(n)}))\right\}\right\}.
\end{aligned}$$

Now use Theorem 9 to obtain that

$$0 \le \limsup_{n\to\infty} \mathbb{E}_n \left[\frac{L(\mathrm{TSP}(\mathbf{Y}^{(n)}(\lambda)))}{n} \right] \le \limsup_{n\to\infty} \frac{L(\mathrm{TSP}(\mathbf{x}^{(n)}))}{n} = 0 \qquad \text{a.s.}$$

Moreover, since packages are arranged on a circle at the beginning of the day, they are all equally likely to be undelivered by time T, and their specific location is independent of $\tilde{N}(T,n,\lambda)$ (note that the $\{\mathbf{Y}_i\}$ may not be independent though). Hence, $\mathbb{E}_n[d(\mathbf{Y}_i,\mathbf{0})|\tilde{N}(T,n,\lambda)] = \overline{r}(n)$, which combined with Theorem 3 gives

$$\begin{aligned}
&\limsup_{n\to\infty} \mathbb{E}_n \left[2\left(\frac{n-\tilde{N}(T,n,\lambda)}{Vn} + \frac{1}{n} \right) \mathscr{R}(n-\tilde{N}(T,n,\lambda)) \right] \\
&= \frac{2}{V} \limsup_{n\to\infty} \frac{1}{n}\mathbb{E}_n \left[\sum_{i=1}^{n-\tilde{N}(T,n,\lambda)} d(\mathbf{Y}_i,\mathbf{0}) \right] + 2 \limsup_{n\to\infty} \frac{1}{n}\mathbb{E}_n \left[\mathscr{R}(n-\tilde{N}(T,n,\lambda)) \right] \\
&\le \frac{2}{V} \lim_{n\to\infty} \frac{n-\mathcal{C}(T,n,\lambda)}{n} \cdot \overline{r}(n) + \lim_{n\to\infty} \frac{\rho}{n} \\
&= \frac{2(1-\alpha(T,\lambda))r^*}{V},
\end{aligned}$$

where $\rho := \sup_{\mathbf{x}\in R} d(\mathbf{x},\mathbf{0})$ is the radius of the bounded region $R \subseteq \mathbb{R}^2$.

We conclude that

$$\begin{aligned}
&\limsup_{n\to\infty} \frac{1}{n} \left(\mathrm{Cost}_P(z;\mathbf{x}^{(n)}) - \mathrm{Cost}_V(\mathbf{x}^{(n)}) \right) \\
&\le \alpha(T,\lambda) \left(\left(\zeta_P + \frac{h_P+z}{v_P} \right) \frac{r^*}{E[B]} + (h_P+z)\tau_P - h_V\tau_V \right) + \left(\zeta_V + \frac{h_V}{v_V} \right) \left\{ \frac{2(1-\alpha(T,\lambda))r^*}{V} - \frac{2r^*}{V} \right\} \\
&= -\alpha(T,\lambda) \left(\left(\zeta_V + \frac{h_V}{v_V} \right) \frac{2r^*}{V} - \left(\zeta_P + \frac{h_P}{v_P} \right) \frac{r^*}{E[B]} - (h_P\tau_P - h_V\tau_V) - z \left(\frac{r^*}{v_P E[B]} + \tau_P \right) \right).
\end{aligned}$$

■

1.5 Numerical Experiments.

This section provides numerical examples illustrating the computational effort of implementing our proposed strategy, the calibration of the parameters, as well as a cost-benefit analysis. Specifically, we focus on the following:

- Choosing the parameters: in particular, how to estimate the costs needed for the formulation of the optimization problem in Section 1.3.2, for which we cite empirical studies in the transportation literature.

- Numerical case studies: including a full implementation of our proposed strategy under two different choices for the distribution of package destinations.
- Cost-benefit analysis: a comparison between the proposed strategy that uses a combination of private drivers with an in-house van delivery system, and a more traditional van-only strategy.
- Justification of the mathematical assumptions: in particular, of how using $\hat{z} = \mathrm{argmin}_z \, \mathrm{Cost}'_P(z; \mathbf{x}^{(n)})$ as an approximation for $z^* = \mathrm{argmin}_z \, \mathrm{Cost}_P(z; \mathbf{x}^{(n)})$ has no significant impact on the optimal cost, and how requiring drivers to pick up packages that are adjacent in the TSP route is not too restrictive.

1.5.1 Parameters.

The parameter estimation mostly follows the sources used in [88].

Parameters related to vans: The parameters that need to be estimated are: ζ_V, the per-mile transportation cost; h_V, the opportunity cost for van drivers; v_V, the average speed of a delivery van; V, the van capacity; and τ_V, the end-point delivery time for vans. To compute ζ_V we use the average diesel price in the U.S. for 2017 (taken from the web), and the estimated maintenance and depreciation costs used in [65] and [10], respectively. In particular, the average diesel price in the U.S. for 2017 was \$2.540 per gallon and the fuel efficiency of a UPS van is 10.6 miles per gallon of diesel, which gives a fuel cost for vans of \$0.25 per mile. The maintenance cost for vans was computed in [65] to be \$0.152 per mile in 2009, and the depreciation cost of a van under city driving conditions was computed in [10] to be \$0.081 per mile in 2003. Adjusting for an annual inflation of 2.5%, we compute a van's per mile cost to be:

$$\zeta_V = 0.25 + 0.152 \times (1.025)^8 + 0.081 \times (1.025)^{14} = 0.550.$$

To estimate h_V, we use that a van driver's wage was approximately \$30 per hour [10] in 2003, so the inflation-adjusted wage in 2017 is

$$h_V = 30 \times (1.025)^{14} = 42.389.$$

The average speed of a UPS van was estimated in [65] to be $v_V = 24.1$ miles/hour. We use an average van capacity of $V = 200$ packages as in [88], which corresponds to a van capacity of 2,000 kg and an average package weight of 10 kg. Finally, for the end-point delivery time we use $\tau_V = 97$ seconds as estimated in [88].

Parameters related to private cars: The parameters related to private cars that need to be estimated are: ζ_P, the per-mile transportation cost; v_P, the average speed; and h_P, the opportunity cost for private drivers (Uber/Lyft drivers). For the end-point delivery time for drivers we simply use $\tau_P = \tau_V$. To estimate ζ_P we use the estimated hourly expenses for part-time Uber drivers computed in Table 6 of [49], which after taking the average over vehicle types gives \$3.84 per hour. Also, we use an average speed for cars in U.S. cities of $v_P = 29.9$ miles per hour (http://infinitemonkeycorps.net/projects/cityspeed/), which yields a per mile cost of

$$\zeta_P = \frac{3.84}{29.9} = 0.1284.$$

To estimate h_P note that the average gross income per hour of an Uber driver is given by $h_P + \zeta_P v_P$, which in [49] was estimated to be \$19.35 per hour in 2015. After adjusting for inflation we obtain that $h_P + \zeta_P v_P = 19.35 \times (1.025)^2 = 20.33$, and therefore, the per hour opportunity cost is

$$h_P = 20.33 - 0.1284 \times 29.9 = 16.49.$$

Remark 1.5.1 *Using the parameters for the vans and private drivers discussed above, we have $\zeta_P + h_P/v_P = 0.68$, $h_P = 16.49$, $\zeta_V + \frac{h_V}{v_V} = 2.31$, $h_V = 42.389$, and $\tau_V = \tau_P = 0.0269$ $(= 97s)$. Note that all of these were estimated for an average U.S. city, and are therefore fairly robust to where the strategy is deployed. However, the values of average distance to the depot, r^*, and the mean bundle size, $E[B]$, are strongly related to the specific region, and will significantly impact the amount of possible improvement with the use of the new strategy. In terms of r^* and $E[B]$, the sufficient condition in Lemma 1.3.3 ensuring that our mixed strategy is better than the van-only strategy becomes:*

$$0.023r^* - 0.635\frac{r^*}{E[B]} + 0.698 > 0.$$

Arrival rate and bundle size distribution: Since we were unable to find any references quantifying the relationship between incentives and the supply of private drivers, as well as for the preferences on the number of packages that private drivers may want to deliver, our choices of $\lambda(z)$ and F are somewhat arbitrary. For the first one we assume a linear relationship of the form $\lambda(z) = b + az$, and estimate a and b using data for Uber and Lyft drivers that has been collected for transporting people (not packages). Let $s(w) = b' + a'w$ denote the supply of Uber drivers in the San Francisco area (about 50 square miles) when the average revenue per unit of time is w. To fit the values of (a', b') we found at http://www.govtech.com/question-of-the-day/Question-of-the-Day-for-061917.html that ride-sharing companies make about 170,000 trips on an average weekday in the San Francisco area (about 50 square miles) and Uber's market share was about 3 times that of Lyft in 2017 (https://www.recode.net/2017/8/31/16227670/uber-lyft-market-share-deleteuber-decline-users), so if we assume that the active operation time is 18 hours, then Uber's average number of trips per hour in the San Francisco area is 7083. Using as a base revenue of $w_0 = 20.33$ dollars per hour, we obtain that $s(w_0) = 7083$. According to the work done by Hall et al. [50], the supply of Uber drivers doubles when the surge pricing increases the wage per hour from w_0 to $1.8w_0$, which gives $s(1.8w_0) = 14166$, and we obtain $a' = 435.50$ and $b' = -1770.75$. Since the wage w and the incentive rate z are linearly related $(w = z +$

$h_P + \zeta_P v_P$), it follows from the linearity of $\lambda(z)$ that the total number of Uber drivers willing to deliver packages instead of transporting passengers is a linear function of $s(w)$. Furthermore, we assume that at the base revenue per hour w_0, only 10% of the Uber drivers would be interested in delivering packages[3], which gives $n\lambda(0) = 4000\lambda(0) = 0.1s(w_0)$, while at the surge pricing rate $1.8w_0$ we expect almost all Uber drivers to prefer the delivery of packages, which yields $n\lambda(1.8w_0 - h_P - \zeta_P v_P) = n\lambda(16.26) = s(1.8w_0)$. To estimate n, we used the fact that there are around 356,916 households in San Francisco (https://datausa.io/profile/geo/san-francisco-ca/#housing), and the average online shopping frequency in 2017 was around 21 times per year (https://www.statista.com/statistics/448659/online-shopping-frequency-usa/), so we use $n = 356916 \times 21/365 = 20535$. Hence, we obtain

$$\lambda(z) = 0.03 + 0.04z.$$

For the distribution of bundles, we assume that F is has the following distribution:

$$P(B = k) = \left(\sum_{i=1}^{20} e^{-10}(10)^i/i! \right)^{-1} \frac{e^{-10}(10)^k}{k!}, \qquad k = 1, 2, \ldots, 20,$$

which corresponds to a Poisson distribution with mean 10 packages conditioned on taking values on $\{1, 2, \ldots, 20\}$.

Table 1.1 lists the parameters we use in both of the numerical experiments in the following section. We target a region of size 5 miles $\times$ 5 miles with a total number of packages $n = 2000$.

Table 1.1: Parameters.

Parameter	*Value*
Total number of packages n	2000
Bundle size distribution F	Poisson(10), conditioned in $\{1, 2, \ldots, 20\}$
Requests arrival rate λ	$\lambda = 0.03 + 0.04z$
Delivery time window $[0, T]$	8h
Area of the square region A	25 square miles

1.5.2 Case studies.

To illustrate the impact that the distribution of the package destinations over the service region may have, we consider two different cases, one where the destinations are uniformly distributed throughout a square region, and another one where we can identify three different clusters in the same square region. For both cases we use a total of 2000 packages, and we sample independently their destinations using stochastic simulation. We then solve (1.3.13) and compute the optimal incentive rate $\hat{z}$. To solve the TSP we used Concorde implemented in R, and for the CVRP we used VRPH 1.0.0. Throughout all our computations we use the L_1 distance.

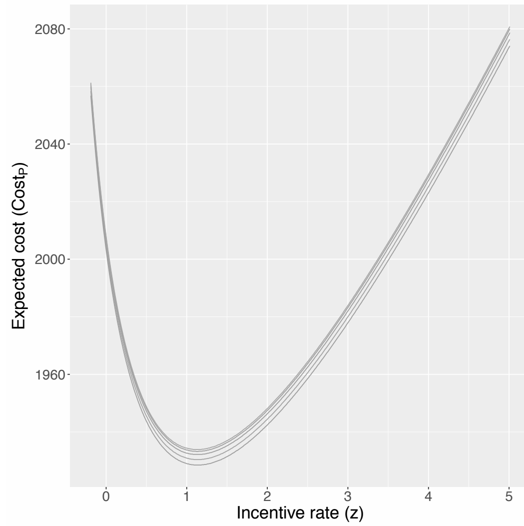

Figure 1.3: Objective function (1.3.13) for 2000 packages uniformly distributed over a square region. Different colors represent different realizations of package destinations.

1.5.2.1 Case 1: Packages distributed uniformly over a square region.

Note that the objective function (1.3.13) depends on the specific package destinations $\mathbf{x}^{(n)} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ only through the length of the TSP route, L(TSP), and the average long-haul distance $\overline{r}(n)$. For $n = 2000$, the differences from one realization to another are very small, as shown in Figure 1.3.

Next, we used stochastic simulation to generate a single realization of the package destinations $\mathbf{x}^{(n)}$ along with its corresponding package pick-up process. More precisely, we simulated the $n = 2000$ independent Poisson processes with rate $\lambda(\hat{z})$, where $\hat{z} = 1.11$ solves (1.3.13), during the time window $[0, T]$, $T = 8$ hours, so the profit rate is 17.6 per hour. We then computed an optimal CVRP route on the leftover packages. To compare the new mixed strategy with the van-only strategy we also computed the length of an optimal CVRP route on the original 2000 packages. Figure 1.4a shows an optimal TSP route of length L(TSP)= 207.81 miles, solved using Concorde, while Figure 1.4b shows an optimal CVRP route of length L(CVRP)= 222.83 miles, solved using VRPH 1.0.0. Both routes are for the original 2000 packages, with the latter needed for the van-only strategy.

Using the simulated CVRP route from Figure 1.4b we computed the cost of the van-only strategy to be:

$$\mathrm{L}(\mathrm{CVRP}(\mathbf{x}^{(n)}))\left(\zeta_V + \frac{h_V}{v_V}\right) + n h_V \tau_V = 222.83(2.309) + 2000(42.389)(97/3600) = \$2798.81.$$

For our proposed mixed strategy, we computed an optimal CVRP route for the leftover packages at time T, which in the simulation run we did had 554 packages and a length

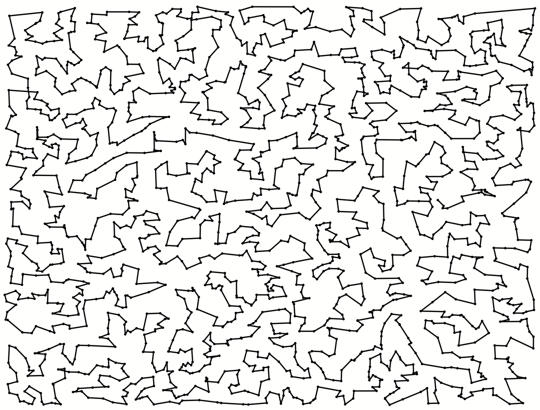

(a) Optimal TSP route solved using Concorde in less than 3 seconds. Length = 207.81 miles. This computation is needed for the mixed strategy.

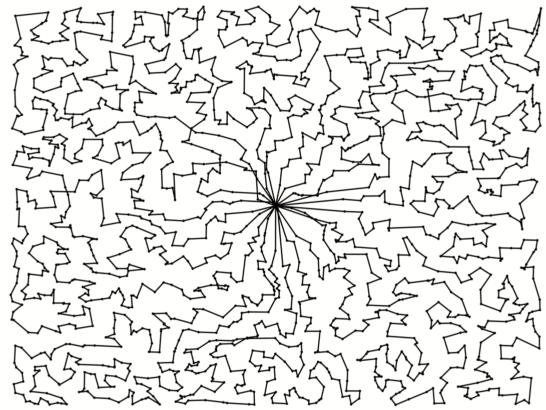

(b) Optimal CVRP route solved using VRPH 1.0.0 in less than 60 seconds. Length = 222.83 miles. This computation is needed for the van-only strategy.

Figure 1.4: Optimal TSP and CVRP routes for 2000 packages uniformly distributed on a 5×5 square miles region.

of 96.07 miles (see Figure 1.5). We then computed the total cost for delivering the 2000 packages, which was computed to be \$1888.26. The improvement compared to the van-only strategy was 32.53%.

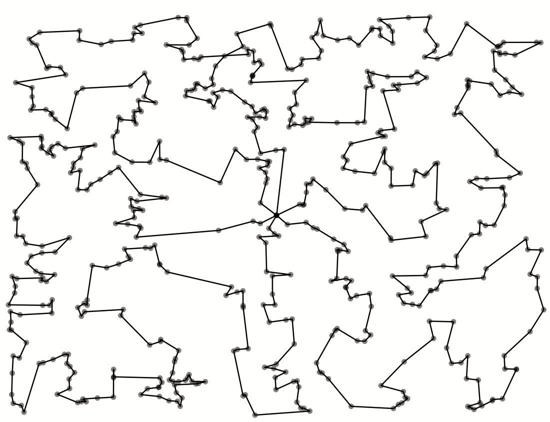

Figure 1.5: Optimal CVRP route for 554 leftover packages solved using VRPH 1.0.0. Length = 96.07 miles.

To show how our methodology scales with the number of packages, we repeated the process described above for different values of the number of packages n. As Table 1.2 shows, in all cases the improvement of the mixed strategy compared to the van-only strategy was between 31% and 33%, indicating the consistency and scalability of our proposed strategy.

Table 1.2: Improvements with different number of packages

Number of packages	*Optimal incentive rate*	*Improvement*
600	1.22	31.93%
1000	1.17	32.01%
1500	1.15	32.15%
2000	1.13	32.40%
3000	1.11	32.70%

1.5.2.2 Case 2: Packages distributed in three clusters over a square region.

For this numerical experiment we consider a more realistic scenario where packages are clustered within the service region. To simulate the package destinations we generated 2000 data points according to the following method:

- 500 packages uniformly distributed in a 5×5 square;
- 700 packages uniformly distributed in the ellipse $\frac{(x-1.5)^2}{1.2^2} + \frac{(y-4)^2}{1^2} = 1$;
- 500 packages uniformly distributed in the ellipse $\frac{(x-3.8)^2}{0.8^2} + \frac{(y-3.3)^2}{1.2^2} = 1$;
- 300 packages uniformly distributed in the ellipse $\frac{(x-2.5)^2}{1.2^2} + \frac{(y-1.4)^2}{1^2} = 1$.

Figure 1.6 plots the objective function (1.3.13) for different realizations of the package destinations $\mathbf{x}^{(n)}$. Note that since the data is clustered around three centers, the typical distance between points is smaller, which yields a smaller expected cost for the same number of packages as in Case 1.

Following the same methodology as in Case 1, we computed optimal TSP and CVRP routes on the same single realization of package destinations, as depicted by Figures 1.7a and 1.7b, respectively. The total lengths for the routes were L(TSP) = 176.61 miles and L(CVRP) = 193.23 miles, respectively. The total cost for the van-only strategy was computed to be

$$\mathrm{L}(\mathrm{CVRP}(\mathbf{x}^{(n)}))\left(\zeta_V + \frac{h_V}{v_V}\right) + nh_V\tau_V = 193.23(2.309) + 2000(42.389)(97/3600) = \$2730.46.$$

We also computed the optimal solution $\hat{z} = 1.21$ to (1.3.13) and the corresponding rate $\lambda(\hat{z})$, and used it to simulate the package pick up process. A single simulation run yielded a total

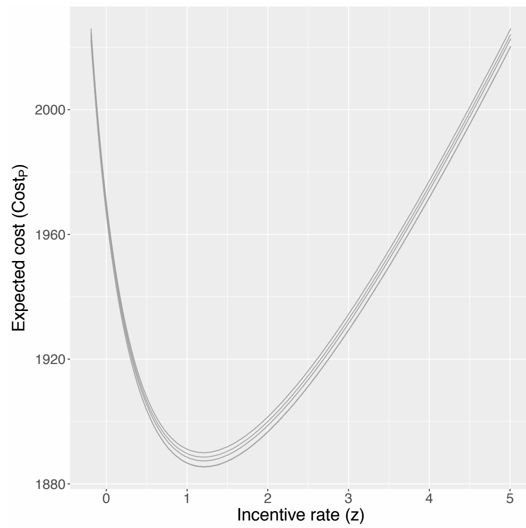

Figure 1.6: Objective function (1.3.13) for 2000 packages distributed in three clusters. Different colors represent different realizations of package destinations.

of 531 leftover packages, an optimal CVRP route of length 81.36 miles, and a total cost of \$1829.31 (see Figure 1.10). Therefore, the improvement of the mixed strategy compared to the van-only strategy was 33.00%.

1.5.2.3 Case 3: Empirical data.

Here we study a real dataset from a well-known delivery company in China. Fig 1.9 shows the distribution of 6312 packages over a 5.30×5.30 square miles region. The optimal incentive rate is 1.07. Following the same methodology as in Case 1, we computed optimal TSP and CVRP routes, as depicted by Figures 1.9a and 1.9b, respectively. The total lengths for the routes were L(TSP) = 265.92 miles and L(CVRP) = 378.48 miles, respectively. The total cost for the van-only strategy was computed to be

$$\mathrm{L}(\mathrm{CVRP}(\mathbf{x}^{(n)}))\left(\zeta_V + \frac{h_V}{v_V}\right) + nh_V\tau_V = 378.48(2.309) + 6312(42.389)(97/3600) = \$8083.15.$$

We also computed the optimal solution $\hat{z} = 1.07$ to (1.3.13) so the profit rate is \$17.56 per hour, and used it to simulate the package pick up process. A single simulation run yielded a total of 1744 leftover packages, an optimal CVRP route of length 81.36 miles, and a total cost of \$5592.19 (see Figure 1.10). Therefore, the improvement of the mixed strategy compared to the van-only strategy was 30.82%.

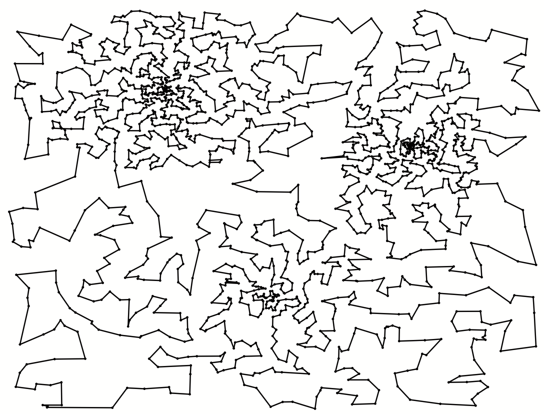

(a) Optimal TSP route solved using Concorde in less than 3 seconds. Length = 176.61 miles. This computation is needed for the mixed strategy.

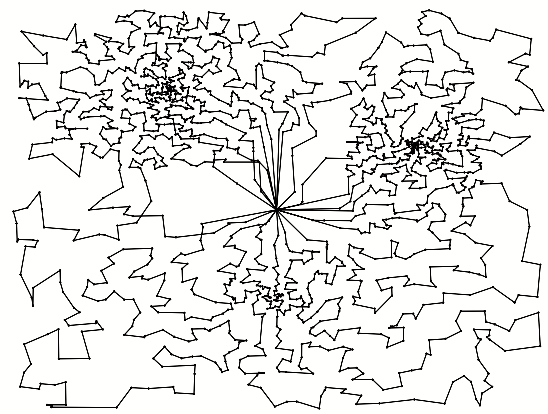

(b) Optimal CVRP route solved using VRPH1.0.0 in less than 60 seconds. Length = 193.23 miles. This computation is needed for the van-only strategy.

Figure 1.7: Optimal TSP and CVRP routes for 2000 packages distributed in three clusters over a 5×5 square miles region.

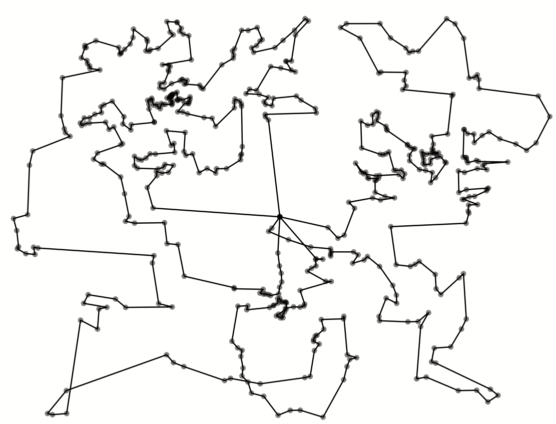

Figure 1.8: Optimal CVRP route for 531 leftover packages solved using VRPH 1.0.0. Length = 81.36 miles.

(a) Optimal TSP route solved using Concorde in less than 30 seconds. Length = 265.92 miles. This computation is needed for the mixed strategy.

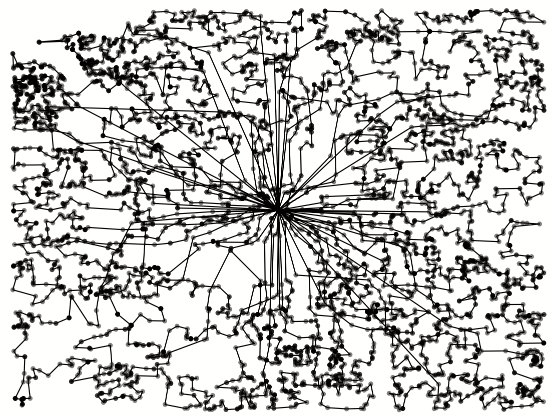

(b) Optimal CVRP route solved using VRPH1.0.0 in less than 6 minutes. Length = 378.48 miles. This computation is needed for the van-only strategy.

Figure 1.9: Optimal TSP and CVRP routes for 6312 packages distributed in three clusters over a 5.30×5.30 square miles region.

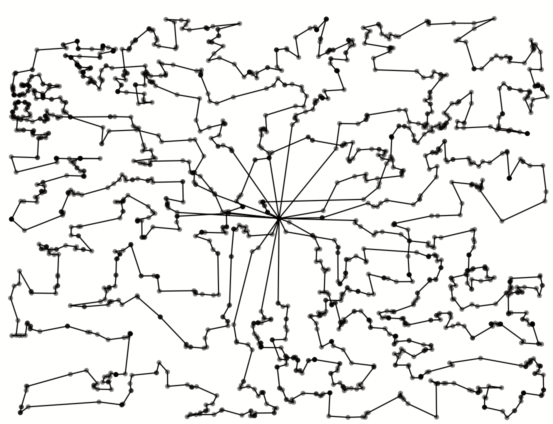

Figure 1.10: Optimal CVRP route for 1744 leftover packages solved using VRPH 1.0.0. Length = 169.70 miles.

Table 1.3: Approximation accuracy and computational cost

Number of packages	$\mathcal{C}(t,n,\lambda)/n$	$\lvert\alpha(t,\lambda)-\mathcal{C}(t,n,\lambda)\rvert/n$	*Computing time* (sec)
100	0.7028	0.0257	< 0.1
300	0.7199	0.0086	0.3396
1000	0.7259	0.0026	6.8489
2000	0.7272	0.0013	44.2113
5000	0.7280	0.0005	629.8755
8000	0.7282	0.0003	2.5436e+03
15000	0.7283	0.0002	1.6490e+04

1.5.3 Approximation accuracy and computational cost.

In Table 1.3, we list values of $\mathcal{C}(t,n,\lambda)/n$ (with $t = 8h$, $\lambda = 0.0728$) and compare them to the limit $\alpha(t,\lambda) = 0.7285$. The corresponding computational time for different values of n, ranging from 100 to 15,000, is also recorded. We notice that computing $\mathcal{C}(t,n,\lambda)/n$ using Proposition 1.3.1 costs more than 4 hours for $n = 15,000$. However, it takes less than 0.1s to compute $\alpha(t,\lambda)$ using Theorem 2, and it provides a very good approximation.

As mentioned in Section 1.3.2, we are approximating the optimal incentive rate z^* that we offer to private drivers by using $\hat{z} = \operatorname{argmin}_z \mathrm{Cost}'_P(z;\mathbf{x}^{(n)})$. To illustrate that this approach not only works asymptotically, but also numerically, we compute here using discrete event simulation the true cost function $\mathrm{Cost}_P(z;\mathbf{x}^{(n)})$ and compare the cost reductions achieved by using $\hat{z}$ vs. z^*. To do this, we generated a set of destinations $\{\mathbf{x}_1,\ldots,\mathbf{x}^{(n)}\}$ and computed the corresponding $\hat{z}$. Then, we chose 30 different values of z in a neighborhood of $\hat{z}$ (including $\hat{z}$) and ran 100 simulation runs for each of them to compute $\mathrm{Cost}_P(z;\mathbf{x}^{(n)})$. We then used these simulated values to determine the minimum cost $\mathrm{Cost}_P(z^*;\mathbf{x}^{(n)})$ and compared it to $\mathrm{Cost}_P(\hat{z};\mathbf{x}^{(n)})$. Our results are included in Table 1.4. As we can see, the values of $\mathrm{Cost}_P(\hat{z};\mathbf{x}^{(n)})$ and $\mathrm{Cost}_P(z^*;\mathbf{x}^{(n)})$ are very close to each other.

Table 1.4: Gap between the expected cost of using the approximately optimal incentive rate $\hat{z}$ vs. the true optimal incentive rate z^*.

Number of packages	$\mathrm{Cost}_P(\hat{z};\mathbf{x}^{(n)})$	$\mathrm{Cost}_P(z^*;\mathbf{x}^{(n)})$
600	653.05	652.54
1000	1032.07	1030.82
1500	1475.73	1473.23
2000	1905.45	1901.30
3000	2755.78	2750.36

1.5.4 Loss due to the TSP assumption.

One of the key features of our framework is that it identifies the package destinations on the 2-dimensional planes with points, first along a circle and then along a line, using an optimal solution to the TSP. Although it is precisely this approach that allows us to analyze the expected number of picked-up packages during $[0, T]$, it is perhaps also the most restrictive assumption we make from a practical point of view since it constraints bundles to be segments of the TSP route. Hence, our definition of bundles leaves open the possibility that a driver's bundle request is rejected even if there are enough nearby packages that could have been picked up but that are not adjacent to each other in the TSP route. To address this possibility we have conducted numerical experiments where we relax our definition of bundle to allow packages to be near each other but not necessarily adjacent in the TSP route, and we count how many times during $[0, T]$ we reject bundle requests because the required segment of the TSP route is no longer available even though there are enough nearby packages for the driver to pick up.

Our experiments consist of two sets of simulations, one where we count the proportion of rejected bundle requests due to our TSP assumption, and one where we relax our definition of bundle and compute the resulting cost. Recall that a bundle of size k at location i in our framework consists of the packages $\{i, i+1, \ldots, i+k-i\}$ along the TSP route, and a request that arrives to find that at least one of the packages in its bundle is no longer available is "rejected". Define r to be the radius of acceptable neighborhoods, where acceptable means that package destinations lying within the circle of radius r centered at location i are considered to be close to each other. A "loss" in the first simulation occurs when a bundle request of size k at location i arrives to find that at least one package in $\{i, i+1, \ldots, i+k-1\}$ is unavailable, say $1 \leq q \leq k$ are unavailable, but the acceptable neighborhood centered at i contains at least q other packages that could have been used to complete the bundle. We then count the "losses" and compute the "loss rate" defined as:

$$\text{Loss rate} = \frac{\text{Number of “losses” in } [0, T]}{\text{Number of “rejections” in } [0, T]}.$$

In the second simulation we consider the same setup as above, however, if a request for a bundle of size k at location i arrives and finds that at least one of the packages in $\{i, i+1, \ldots, i+k-1\}$ is unavailable, say $1 \leq q \leq k$ are unavailable, we attempt to complete the bundle with available packages within its acceptable neighborhood. If it is possible to find q additional packages, we choose the closest ones to location i and allow the driver to pick up the completed bundle; if not, we reject the request. We continue the simulation until time T and then compute the cost of delivering all n packages according to our framework, i.e., using the pricing scheme based on our estimates for the expected number of packages that can be picked up during $[0, T]$ using the TSP route.

Our results for both sets of simulations are given in Table 1.5 for various choices of r. We chose $n = 3000$ packages and $r \in \{0.05, 0.08, 0.09, 0.12, 0.15\}$, which correspond to the 25% quantile, median, mean, 75% quantile, and 85% quantile, respectively, of the distribution of the neighboring distance along the initial TSP route; for each r we ran the simulation 50 independent times. The second column refers to the first experiment, where we only identify the "losses" but keep the "rejections" as in our original framework; the third column refers to the second experiment where we allow the drivers to pick up relaxed bundles; the fourth gives the difference in cost (%) between the relaxed policy and the original policy, whose total cost was 2,755.78. As we can see, the difference in the total cost for delivering n packages between the two approaches is very small, and therefore, our narrow definition of bundles based on the TSP route has minimal influence in the total cost of our proposed approach.

Table 1.5: Simulated results for different choices of r.

r	Loss rate (original)	Total cost (relaxed)	Difference in cost (%)
0.05	0.54%	2754.98	0.03%
0.08	1.20%	2753.29	0.09%
0.09	1.44%	2752.35	0.12%
0.12	2.89%	2751.63	0.15%
0.15	4.46%	2745.19	0.38%

Chapter 2

Sequential choice bandits: Learning with marketing fatigue

2.1 Introduction.

Service providers and retailers routinely rely on emails and app notifications to interact with their users. When done well, these messages act as digital reminders that increase customer engagement, raise brand awareness and conversion. However, frequent messaging can easily backfire. Marketing fatigue, which refers to an overexposure to unwanted marketing messages, could aggravate users and prompt them to forgo receipt of future messages by unsubscribing or deleting the app.

Motivated by this dilemma, we consider a setting where a platform needs to dynamically determine a sequence of messages for its users. The messages are presented to a user sequentially. Upon reviewing a message, a user can either accept or reject the message. The platform earns a reward when a message is accepted. If the user rejects the message, she then decides to either receive the next message unless the sequence runs out, or abandon the platform. If a user abandons, the platform incurs a penalty cost from losing that user. The objective of the platform is to maximize its expected payoff which is the reward after subtracting the penalty cost due to abandonment.

To draw a connection between this problem and the earlier motivating example, messages can represent digital marketing content such as an email or app notification regarding a product or service that a marketer wishes to recommend. The marketer earns revenue whenever a user clicks on the content. It is an indication that the content is of interest to that user. On the other hand, when the user is not interested in the content, there is a possibility that she will unsubscribe from the email list or delete the app. In fact, 66% of responders of a recent survey[4] cited frequent and irrelevant messaging as the main reasons behind unsubscribing from mailing lists. The abandonment cost in our problem can be viewed as the cost of user acquisition as the marketer replenishes his user base. Based on a Harvard Business Review article[5], the cost of customer acquisition is estimated to be 5 to

25 times higher than keeping an existing customer. Therefore, fatigue control is a critical component of digital marketing content dissemination.

In this work, we investigate the platform's problem in an online learning framework. More specifically, the platform optimizes its cumulative payoff which includes the abandonment cost over a horizon of length T by offering sequences of messages to its users. By observing users' feedback, the platform learns two pieces of information to improve its decisions, i.e., messages' attractiveness and users' abandonment behavior. The former measures the relevance of a message to a user, while the latter quantifies users' tolerance towards irrelevant messages in terms of how likely one will abandon the platform upon receiving unsatisfactory messages. We refer to the online learning task that the platform faces as the *sequential choice bandit* (*SC-Bandit*) problem. We first analyze the *non-contextual SC-Bandit* setting, where the optimal strategy is a static sequence for all users. Next, we move onto the *contextual SC-Bandit* setting where, message features and user features are incorporated to enable personalized recommendation. For both settings, we design a learning algorithm, and evaluate its performance by analyzing its regret, which measures the difference between the maximum payoff the platform could have obtained had it known all the underlying parameters and the accumulated reward under the proposed algorithm.

A common assumption used in the academic literature for online learning is that users' interactions with a platform are instantaneous. Algorithms then apply learning from *preceding* users to *subsequent* users as decisions for different users do not overlap. The resulting recommendation for an individual user is hence static. In reality, a user may engage with a platform for some period of time. As new users continue to arrive, there could be multiple users interacting with the platform at each moment. This is the setting that we are analyzing in this work, which makes a significant departure from the existing literature. For each user, our learning algorithm dynamically updates and adjusts her sequence of recommendations, taking into account of what has been previously shown to avoid duplication. The interwoven events (e.g., arrival of new users and existing users on the platform) along with the sequential and adaptive nature of decisions significantly complicate the analysis of the SC-Bandit problem. In addition, incorporating user's abandonment behavior leads to incomplete feedback that further increases the complexity of the analysis, as users' responses are not guaranteed due to abandonment. The main contribution of our work is summarized as follows:

1. *Novel formulation* We propose a novel sequential choice model which extends the popular cascading model[6] to capture a user's interactions with a platform, and model the effect of marketing fatigue with a very general abandonment distribution. To the best of our knowledge, our paper is the first to analyze adaptive algorithms for sequential recommendation in an online setting.

2. *Efficient offline algorithm* Even in the offline setting where attraction probabilities of messages and users' abandonment behavior are known, the optimization problem to determine an optimal sequence to recommend is combinatorial in nature without an

obvious efficient algorithm. We provide a polynomial-time algorithm that determines the optimal sequence.

3. *Analysis on non-contextual SC-Bandit* In the online setting where message attractiveness and abandonment behavior are unknown, we propose a learning algorithm which simultaneously explores and exploits. The algorithm yields a sequence of messages for every user and is updated at every time step. We show that its regret is bounded above by $O(\sqrt{NMT\log T})$, where N is the number of available messages and M is the maximum number of messages that can be sent to a user.

4. *Analysis on contextual SC-Bandit* We extend the non-contextual SC-Bandit problem to enable personalized recommendation by incorporating features. As we observe binary feedback (i.e., click, or abandon), we model message attractiveness and abandonment behavior as a generalized linear function of the message feature and user feature respectively. To the best of our knowledge, we are the first to analyze contextual cascading bandits with generalized linear reward function. We propose a contextual SC-bandit algorithm and prove that its regret is bounded above by $M\left(\sqrt{d_X T\log T\log(T/d_X)} + \sqrt{d_Z T\log T\log(NT/d_Z)}\right)$ where d_X and d_Z refer to the dimensions of user feature and message feature respectively.

5. *Numerical studies* In addition to using synthetic data to study the robustness of our algorithms, we also perform experiments using real-world data from a large e-commerce website. The dataset reveals different abandonment behavior across users, highlighting the benefit of contextual learning. Using the data, we evaluate the performance of our algorithms for both the contextual and non-contextual settings.

The paper is organized as follows. We review the related literature in Section 2.2. In Section 2.3, we introduce our model, and derive the optimal solution to the offline problem. We begin the analysis of the online problem in Section 2.4, where we propose a learning algorithm to the non-contextual SC-Bandit problem and characterize its regret bound. Contextual SC-Bandit problem is analyzed in Section 2.5. We evaluate the performance of our algorithms with both synthetic and real-world data in Section 2.6. Finally, we conclude the paper in Section 2.7. For the key technical results in the paper, we include proof sketches in the main body of the paper. Detailed proofs can be found in the Appendix.

2.2 Literature Review.

Our work is closely related to four streams of academic literature, namely, multi-armed bandits, learning to rank applications, assortment optimization, and studies on marketing fatigue.

Multi-armed bandits (MAB) There has been a plethora of work on MAB and its variants (e.g., [96, 108, 90, 107]). The trade-off between exploration and exploitation is a central problem in many different communities. Our problem can be viewed as a combinatorial bandit problem where a platform chooses a set of messages to be displayed in a certain order. A naive approach is to treat each possible combination with a specific order as an arm. However, the number of arms increases exponentially with the number of messages under this approach. Other combinatorial bandit work assuming linear reward [9, 99] or independent rewards [20] cannot be directly applied to our model, as our rewards are dependent on the order of the actions (i.e., the order in which the messages are being sent).

A popular variant of MAB is the contextual bandit problem, which has been applied in a deluge of applications such as recommender systems, search engines and dynamic pricing (e.g., [4, 69, 70, 89]). A common assumption for contextual bandit is that the expected reward is a linear function of features [24, 2, 83]. In our setup, as the platform only observes binary feedback (i.e., whether a user accepts a message, and whether a user abandons the platform), we assume the expected reward follows a logit function of features, as in generalized linear bandit problems [70, 41]. In addition, we also consider sequential interactions with adaptive recommendation and users' abandonment behavior.

Learning to rank Our work adds to the large body of literature on learning to rank algorithms, motivated by applications in recommendation systems, web search and information retrieval. Depending on the choice of the underlying click models which describe users' feedback through clicks (see [25] for an overview), different algorithms have been proposed. Our setting is based on the cascade model, which is arguably one of the most popular click models [33]. It assumes that users examine the results sequentially and click on the first relevant recommendation. The probability of a click depends on the relevance of a result, as well as the irrelevance of all previous results; hence the name. Recently, [61, 62, 31, 72, 23] investigate this model in an online setting, which they refer to as the cascading bandits. Their task is to learn the attraction probabilities of messages and select m messages, where m is an exogenous parameter. [72, 116, 71, 22] further study contextual cascading bandits. Compared to the previous works on cascading bandits, there are three key differences in our setting: 1) Instead of providing a fixed sequence of messages, our content recommendation is dynamic and adaptive based on the feedback; 2) Besides learning attraction probabilities of messages, we also need to learn users' abandonment behavior. By allowing abandonment, the order of the messages becomes critical, as a user might leave early before she sees the message that she would otherwise have clicked. This addresses one key criticism of the cascade model, where the order of the messages does not matter, since a user views the recommendations sequentially until she finds something she likes or the list runs out [64, 58]; 3) As abandonment incurs costs, the total length of the sequence m is also a decision variable in our setting, as opposed to a fixed parameter in the existing cascading bandits. Compared to the contextual cascading bandits, besides the aforementioned differences, our reward function is generalized linear function instead of linear relationship, which is a more general function class.

Assortment optimization Assortment optimization refers to the problem of selecting a set of products to offer to a group of customers so as to maximize the revenue that is realized when customers make purchases according to their preferences. It is a central topic in the operations management literature. We refer the reader to [57] for a comprehensive review. [109] formulate the assortment planning problem by using a multinomial logit model [112, 73] to describe user behavior.

More recent literature such as [17, 98, 5, 6, 22] combine learning with the assortment problem, as customer preferences are unknown a priori and need to be learned. Our work can be viewed as a dynamic assortment problem to determine a set of messages with a specific order. Existing dynamic assortment problems typically model a single interaction between the platform and a user, who can either choose an item from the assortment or leave without a purchase. In contrast, our model captures multiple interactions between a user and a platform, where the sequential nature of the decision-making plays a crucial role in the analysis.

Marketing fatigue It is a well-documented phenomenon in marketing that "more" is not necessarily "better", as the benefits of a marketing campaign are not in fact purely increasing with the number of messages sent [106]. Some work contributes customer dissatisfaction to information overload which occurs when individuals receive more information than they can process, and proposes countermeasure to control the flow of information (e.g., [21]). [16] and [3] study the relationship between overexposure of marketing content and customer dissatisfaction to explain the "Groupon effect", in which viral marketing via Groupon coupons leads to lower Yelp ratings.

One of the consequences of marketing fatigue is user abandonment. Some recent work such as [103] combines both online learning with abandonment behavior. In their setting, a user has a threshold drawn from an unknown distribution and she abandons if the platform's action x exceeds that threshold. The platform needs to learn the distribution while optimizing x to maximize its discounted reward. One of our key differentiators is how we model abandonment in the presence of sequential behavior. The decision to abandon is an interplay of the relevance of the messages, the order of these message in a sequence and a user's tolerance towards unsatisfactory content. [11] focuses on the recommendation task in a setting where little is known about incoming customers. They provide empirical evidences of customer disengagement based on ad campaign data and propose a modified learning algorithm which constrains the search space to avoid over-exploration. In contrast to providing a single recommendation at every time step in their setting, multiple interactions could take place for several users in our work. Moreover, while the parameters describing abandonment behavior are given in [11], they need to be learned in our setting.

2.3 Problem Setup.

In this section, we formally introduce our model. Next, we characterize the optimal strategy of the platform in the offline setting where all the problem parameters are known.

2.3.1 Model.

Assume there are N different messages for the platform to choose from. Define the set of these N messages as X. We use u_i to denote the attractiveness of message i to a user, which is also known as the attraction probability [61], where $0 \leq u_i < 1$. It is the probability that a user finds the content of message i relevant and interesting, reflecting her preferences.

Users arrive at time $t = 1, \cdots, T$. For each user, the platform determines a sequence of messages $\mathbf{S} = (S_1, S_2, \cdots, S_m)$, where S_i denotes the message ranked in the i^{th} position and m denotes the total length of the sequence. We assume that at most M messages will be sent to a single user, i.e., $m \leq M$, potentially due to operational or budget constraints. The maximum value that M can take is N. We use $I(\cdot)$ to denote the index function that maps the position in a sequence to the message content, i.e., $I(i) = k$ if and only if $S_i = k$.

Messages are displayed sequentially to a user. When a user accepts [7] the first message that she is satisfied with, no further messages will be shown to her and the platform earns a reward of 1. Whenever a message is rejected (non-click), we consider its content unsatisfactory or not attractive, as they are not of sufficient interest to the user. When that happens, the user can either choose to abandon the platform (e.g. unsubscribe the promotion email), or see the next message unless the sequence has run out. The platform incurs a penalty cost of c when a user abandons, which can be viewed as the cost of user acquisition as the marketer replenishes his user base.

Abandonment behavior under marketing fatigue We model users' abandonment behavior by a random variable W, drawn from a distribution F. W can be viewed as a proxy for a user's tolerance towards irrelevant recommendation, which measures the maximum number of unsatisfactory messages before triggering abandonment. That is, the probability of abandonment upon receiving the k^{th} unsatisfactory message is $P(W = k)$. Similarly, the probability that a user has not abandoned after viewing k unsatisfactory messages is $P(W > k)$.

Given a sequence of messages $\mathbf{S} = (S_1, S_2, \cdots, S_m)$, define $\mathbb{P}_a(\mathbf{S})$ as the total abandonment probability, which can be calculated as

$$\mathbb{P}_a(\mathbf{S}) = \sum_{k=1}^{m} P(W = k) \prod_{j=1}^{k} (1 - u_{I(j)}).$$

It sums over the joint probabilities of not finding the first k messages attractive and user's tolerance level is k.

We use f_i to denote the abandonment probability conditioned on prior rejected messages. More specifically, f_i is the probability that a user abandons the platform upon receiving the i^{th} unsatisfactory message, conditional on the user's tolerance is larger than or equal to i, i.e.,

$$f_i = P(W = i | W \geq i) = \frac{P(W = i)}{P(W \geq i)}.$$

Sequential choice model When message i is part of a sequence $\mathbf{S} = (S_1, S_2, \cdots, S_m)$, the probability of that message being clicked, which is denoted as $\mathbb{P}_i(\mathbf{S})$, depends on its position in the sequence, its own content, as well as the content of other messages shown earlier. Formally,

$$\mathbb{P}_i(\mathbf{S}) = \begin{cases} u_i, & \text{if } i = S_1 \\ P(W \geq l) \prod_{k=1}^{l-1} (1 - u_{I(k)}) u_i, & \text{if } i = S_l, l \geq 2 \\ 0, & \text{if } i \notin \mathbf{S}. \end{cases} \tag{2.3.1}$$

When message i is the first message of the sequence, the probability of accepting i is simply u_i, which is the attraction probability of message i. For the remainder of the sequence, the probability of accepting i as the l^{th} message is the joint probability that 1) she has rejected the first $l-1$ messages, $\prod_{k=1}^{l-1}(1-u_{I(k)})$; 2) the user has not yet abandoned after receiving $l-1$ unsatisfactory messages, $P(W > l-1) = P(W \geq l)$; 3) the probability that she finds message i attractive, u_i. The model which we have adopted in this work to describe users' browsing behavior (without abandonment) is known as the cascade model [33]. The model assumes that users examine the results sequentially and click on the first relevant recommendation. The probability of a click in a cascade model depends on the relevance of a result, as well as the irrelevance of all previous results; hence the name. In this work, we extend the cascade model by incorporating user abandonment behavior. The probability of abandonment also depends on the quality of messages.

Payoff optimization problem Let $U(\mathbf{S}, \mathbf{u}, F)$ denote the total payoff that the platform receives from a given sequence of messages $\mathbf{S}$ when the attraction probabilities are $\mathbf{u}$ and the abandonment behavior follows a general distribution F. For ease of notation, we use $U(\mathbf{S})$ to denote $U(\mathbf{S}, \mathbf{u}, F)$. The expected payoff which the platform is trying to optimize is defined as

$$E[U(\mathbf{S})] = \sum_{i \in X} \mathbb{P}_i(\mathbf{S}) - c\mathbb{P}_a(\mathbf{S}),$$

where c is the cost of losing a customer due to abandonment. In contrast to the traditional assortment problems which only focus on revenue maximization, the objective in our model also includes a penalty of losing users.

The platform's optimization problem is defined as follows,

$$\begin{aligned} \max_{\mathbf{S}} \quad & E[U(\mathbf{S})] \\ s.t. \quad & S_i \neq S_j, \forall i \neq j \\ & |\mathbf{S}| \leq M \end{aligned} \tag{2.3.2}$$

The constraint specifies that the sequence cannot contain duplicated messages. It is included to avoid unrealistic solutions where the optimal sequence consisting of identical messages. The decision variables for the platform contain both the order of messages and the total length of messages m. We define the optimal sequence of messages, $\mathbf{S}^*$, as the sequence that maximizes the expected payoff, i.e., $\mathbf{S}^* = \text{argmax}_{\mathbf{S}}\, E[U(\mathbf{S})]$ under the constraint.

2.3.2 Characterization of the optimal sequence $\mathbf{S}^*$.

In this section, we first investigate properties of the optimal sequence of messages, and then describe an algorithm to solve the optimal payoff optimization problem when attraction probabilities $\mathbf{u}$ and abandonment behavior characterized by F are both known to the platform.

In Proposition 2.3.2, we show that we can compare the expected payoff generated under different abandonment distributions if they follow a stochastic order. We state the definition of stochastic order below for completeness.

Definition 2.3.1 (Stochastic order) *Real random variable W_1 is stochastically larger than or equal to W_2, denoted as $W_1 \gtrsim_{s.t.} W_2$, if*

$$P(W_1 > x) \geq P(W_2 > x) \text{ for all } x \in \mathbb{R}.$$

Proposition 2.3.2 *Assume $\mathbf{S}'$ and $\mathbf{S}''$ are the optimal sequences generated under abandonment distributions W_1 and W_2 respectively. If $W_1 \gtrsim_{s.t.} W_2$, we have*

$$E[U(\mathbf{S}', \mathbf{u}, F_{W_1})] \geq E[U(\mathbf{S}'', \mathbf{u}, F_{W_2})],$$

where $U(\mathbf{S}, \mathbf{u}, F_W)$ denotes the payoff under strategy $\mathbf{S}$ when the valuation and abandonment distribution are $\mathbf{u}$ and F_W, respectively.

The definition of $W_1 \gtrsim_{s.t.} W_2$ implies that users under F_{W_1} are less likely to abandon the platform upon receiving the same number of unsatisfactory messages than users under F_{W_2}. Thus, intuitively, Proposition 2.3.2 states that the expected payoff is higher when users are more tolerant.

We now turn to solve the platform's optimization problem. It is combinatorial in nature as it needs to choose a subset with a specific order from all available messages, giving rise to $\sum_{k=1}^{N} \binom{N}{k} k!$ possible combinations. The problem closely related to is the project selection and sequencing problem [47, 63]. To see this, we first expand the platform's expected payoff function, i.e.,

$$E[U(\mathbf{S})] = \sum_{i=1}^{|\mathbf{S}|} \prod_{j=1}^{i-1} (1 - f_j)(1 - u_{I(j)}) \left(u_{I(i)} - c f_i (1 - u_{I(i)}) \right).$$

Let w_{ji} denote the reward for completing job i as order j, where $w_{ji} := u_i - c f_j (1 - u_i)$. The reward collected after job i will be discounted at rate $p_{ti} := (1 - f_t)(1 - u_i)$. Thus, the objective of the platform problem in (2.3.2) can be viewed as to choose a processing order with an optimal job length that maximizes the expected discounted reward. Fig B.1 in the appendix shows an example of a processing order. Finally, we can rewrite the platform's optimization problem in (2.3.2) as follows.

$$\begin{aligned}
\max_{x} \quad & \sum_{t=1}^{N}\sum_{i=1}^{N} x_{ti} w_{ti} \prod_{j=1}^{t-1}\prod_{k=1}^{N} p_{jk}^{x_{jk}} \\
s.t. & \sum_{j=1}^{N} x_{jk} \leq 1, \forall k \\
& \sum_{k=1}^{N} x_{jk} \leq 1, \forall j \\
& \sum_{j=M+1}^{N}\sum_{k=1}^{N} x_{jk} = 0 \\
& x_{jk} = \{0,1\}, \forall j,k
\end{aligned} \tag{2.3.3}$$

where the decision variable $x_{jk} = 1$ if message k is placed at the j^{th} position in the sequence.

In the following result, we characterize a key property of the optimal sequence $\mathbf{S}^*$, which is critical for determining an efficient offline algorithm.

Theorem 4 *For any pair of messages i, j such that $u_i \geq u_j$, we have $\sum_{t=1}^{k} x_{ti} \geq \sum_{t=1}^{k} x_{tj}$ for any $k = 1, \cdots, N$ in the optimal solution to the optimization problem* (2.3.3).

An important implication of Theorem 4 is that: Assume the index function of the optimal sequence $\mathbf{S}^*$ is $I(\cdot)$, then we have $u_{I(i)} \geq u_{I(j)}$ for any $1 \leq i < j \leq |\mathbf{S}^*|$. Moreover, $u_i \geq u_j$ for all $i \in \mathbf{S}^*$ and $j \notin \mathbf{S}^*$.

Theorem 4 states that the optimal sequence should satisfy the condition that $u_{I(i)} \geq u_{I(i+1)}$ for $i < |\mathbf{S}|$. Therefore, the problem is reduced to select m messages where m is also a decision variable. However, if we naively enumerate m from 1 to N and calculate the corresponding reward, the complexity is $O(N^2)$ since it needs to perform k calculations in the k^{th} loop. We propose Algorithm 1 to determine the optimal sequence which reduces the complexity from $O(N^2)$ by doing a brute-force search to $O(NlogN)$. As many recommenders have a huge selection of content to choose from, the reduction by a factor of N indicates a significant improvement in terms of the efficiency of the algorithm.

In Algorithm 1, we use $U((S_j, \cdots, S_m), \mathbf{u}, (f_j, \cdots, f_m))$ to denote the payoff for the sequence of messages $(S_j, \cdots, S_m)$ with corresponding abandonment probabilities $(f_j, \cdots, f_m)$. In line 5-7, we evaluate the expected payoff of $U((S_j, \cdots, S_m), \mathbf{u}, (f_j, \cdots, f_m))$. If it is negative, then we remove the subsequence $(S_j, \cdots, S_m)$ from the current optimal sequence. In the next iteration, we shorten the sequence and re-evaluate the payoff. We prove the optimality of the sequence found by Algorithm 1 in Theorem 5.

Theorem 5 *Algorithm 1 finds the optimal sequence.*

Algorithm 1: Determine the optimal sequence $\mathbf{S}^*$.

Sort messages based on their attraction probabilities in a descending order
$S' = (S_1, S_2, \cdots, S_N)$; $m = M$;
for $j = M : 1$ **do**
  Calculate $E[U((S_j, \cdots, S_m), \mathbf{u}, (f_j, \cdots, f_m))]$;
  if $E[U((S_j, \cdots, S_m), \mathbf{u}, (f_j, \cdots, f_m))] < 0$ **then**
    $m = j - 1$;
  end
end
$\mathbf{S}^* = (S_1, \cdots, S_m)$

We prove Theorem 5 by induction. The detailed proof can be found in Appendix B.3.

Intuitively, to avoid incurring the abandonment cost, the platform may prefer sending shorter sequences before it is confident of how users perceive the content. On the other hand, when $f_j = 0$ for all j, i.e., there is no risk of user abandonment, the platform can send out the entire sequence. As long as the sequence contains one attractive message, the user will eventually select it and the platform earns a reward. Thus, the order of the messages does not influence the reward, which is significantly different when we incorporate abandonment behavior.

2.4 Online Learning.

In the previous section, we assumed that both the attraction probabilities of the messages $\mathbf{u}$ and the user abandonment behavior F are known to the platform. It is natural to ask what the platform should do in the absence of such knowledge. Beginning from this section, we start investigating the online setting where these two pieces of information need to be learned based on users' feedback. We will start by analyzing the non-contextual SC-Bandit problem, and defer the discussion on the contextual setting which utilizes message and user features in the following section.

2.4.1 Setting for non-contextual SC-Bandit.

For each user arriving at time $t \in [1, T]$, the recommendation process begins at t, as the platform determines a list of messages and sends the individual content sequentially at t, $t+D$, $t+2D$, where D denotes the fixed and pre-determined time interval. The recommendation process terminates and no further message will be sent when one of the following events occurs: 1) the user clicks on a message; 2) she abandons the platform; 3) the sequence of the recommendations has run out. We call a user *active* at time t if her recommendation process has not yet terminated, and *inactive* otherwise. We want to point out that during the time period when a user is active, her list of messages may change at every time step as the platform re-solves the optimization problem using the updated estimates of $\mathbf{u}$ and F. Thus, at every time step, the platform needs to make multiple decisions for both the incoming user, as well as active users who arrived earlier. The adaptive nature of recommendation at an individual user level marks the key difference of our setting from the existing work on cascading bandits (e.g., [61, 62, 31]). We use the following example to illustrate the interwoven events and decisions in our online setting.

Example 2.4.1 *Figure 2.1 shows the behavior of 4 users who arrive at time* $t = 1, \cdots, 4$*, and* $D = 1$*. Actions are color-coded, i.e., blue (red) action indicates that the user rejects (clicks) the message, while the cross indicates that the user abandons the platform after rejecting a message. Take user 1 as an example. When she arrives at* $t = 1$*, the initial sequence of recommendation includes messages M3, M5, M10 and M15. M3 is sent and User 1 rejects it. At* $t = 2$*, the sequence is unchanged and the platform sends the next message M5. User 1 again rejects this message, but continues to stay on the platform. At* $t = 3$*, the sequence is updated to M15 and M20, and the platform sends M15. User 1 accepts this message and the recommendation process terminates.*

At every time stamp t*, the platform needs to update the recommendation for all active users. For instance, at* $t = 4$*, the platform sends messages to both User 3 and 4, as the recommendation processes for User 1 and 2 have terminated (i.e., User 1 accepted the message and User 2 abandoned the platform at* $t = 3$*).*

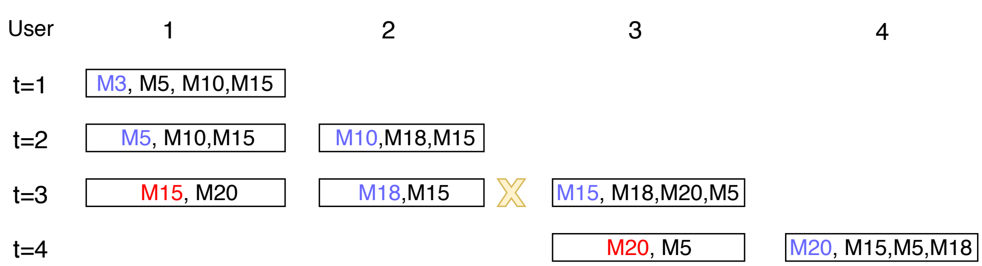

Figure 2.1: An illustrative example on users' interactions with the platform.

We want to point out that the platform is not restricted to sending only one message per time. The platform can send at most $\pounds$ messages at the same time where $\pounds$ is a predetermined variable and $\pounds \leq M$. Figure 2.1 shows a special example where each single user only receives one message per time.

2.4.2 Algorithm for non-contextual SC-Bandit.

In this section, we will present an online algorithm that simultaneously explores and exploits for the SC-Bandit problem to learn the attraction probabilities of message $\mathbf{u}$, as well as the distribution F which describes users' abandonment behavior. We want to emphasize several challenges in our analysis: 1) We only observe partial feedback as a user may exit the platform before examining all the messages which are being recommended. 2) Both $\mathbf{u}$ and F which jointly determine the feedback are latent and needed to be learnt. In particular, f_i is not merely a fixed position-based quantity (recall f_i is the probability of abandonment upon receiving the i^{th} *unsatisfactory* message). Hence, the feedback on abandonment to learn f_i also depends on the content which has been previously recommended to a user. 3) Because the recommendation process for an individual user is dynamic, the decision at time t for an active user is constrained by what has been previously offered to her, as well as the feedback collected from all users up to that time instance.

Our proposed algorithm is optimistic in the face of uncertainty, coming in the form of the upper confidence bounds. We first need to identify the unbiased estimators $\hat{u}_i(t)$ and $\hat{f}_j(t)$ for $1 \le i \le N$ and $1 \le j \le M$, respectively. Let $T_i(t)$ denote the total number of users who observe message i by time t, and $Q_i(t)$ denote the total number of users accepting message i. Note that a user does not necessarily observe message i even if i is included in the offered sequence $\mathbf{S}$ due to abandonment. Let $n_j^a(t)$ denote the number of users who abandon the platform after receiving j unsatisfactory messages by time t. We use $n_j^e(t)$ to denote the number of users that reject a message without abandonment after receiving j unsatisfactory messages by time t. In Example 2.4.1 (see Fig 2.1), $n_1^e(1) = 1$, $n_1^e(2) = 2$, $n_2^e(2) = 1$, $n_1^e(3) = 3$, $n_2^e(3) = 1$, $n_3^e(3) = 0$. $n_1^a(1) = n_1^a(2) = n_1^a(3) = n_1^a(4) = 0$, $n_2^a(3) = 1$. Let $\tilde{T}_j(t) = n_j^e(t) + n_j^a(t)$, which denotes the total number of times users that reject j unsatisfactory messages by time t. Formally,

$$T_i(t) = \sum_{l=1}^{t} 1(i \in \mathbf{S}^l) 1(\text{user } l \text{ observes message } i \text{ before time } t) \quad \text{and}$$

$$\tilde{T}_j(t) = \sum_{l=1}^{t} 1(\text{user } l \text{ rejects the } j^{th} \text{ unsatisfactory message before time } t),$$

where $\mathbf{S}^l$ is the sequence of messages sent to user l.

Lemma 2.4.2 (Unbiased estimator) *$\hat{u}_i(t) = Q_i(t)/T_i(t)$ is an unbiased estimator for u_i. Moreover, $\hat{f}_j(t) = n_j^a(t)/\tilde{T}_j(t)$ is an unbiased estimator for f_j.*

With the unbiased estimators shown in Lemma 2.4.2, we define an optimistic estimator for the attraction probabilities $\mathbf{u}$ and the abandonment probability f as follows,

$$u_{i,t}^{OP} = \min\left\{\hat{u}_i(t) + \sqrt{2\log t / T_i(t)}, 1\right\} \quad \text{and} \tag{2.4.1}$$

$$\tilde{f}_{j,t}^{OP} = \max\left\{\hat{f}_j(t) - \sqrt{2\log t / \tilde{T}_j(t)}, 0\right\}, \quad f_{j,t}^{OP} = \max_{k \le j} \quad \tilde{f}_{k,t}^{OP}. \tag{2.4.2}$$

We propose Algorithm 2 as an exploration-exploitation algorithm for the SC-Bandit problem. At a high-level, for a user arriving at time t, we use $u_{i,t-1}^{OP}$ and $f_{j,t-1}^{OP}$ to calculate the current optimal sequence of messages and offer the first message in the sequence to her. For every remaining active user, we use $u^{OP}_{i,t-1}$ and $f^{OP}_{j,t-1}$ to determine the optimal sequence from a subset of available messages after excluding content which has been previously shown to the user, and offer the first message in the updated sequence. We update the optimistic estimator to $u^{OP}_{i,t}$ and $f^{OP}_{j,t}$ each time when the feedback is received.

Algorithm 2: An online learning algorithm for SC-Bandit under marketing fatigue

---

1 **Initialization:** Available messages X; set $u^{OP}_{i,0} = 1$ for all $i \in X$ and $f^{OP}_{j,0} = 0$ for all
2 $1 \leq j \leq N$; $n^e_j(0) = 0$; $n^a_j(0) = 0$; $t = 1$;
3 **while** $t < T$ **do**
4   **for** *Any active user l with message scheduled to be sent at time t* **do**
5   Compute $\mathbf{S} = \arg\max_{\mathbf{S} \subset X \setminus O_l} \quad U(\mathbf{S}, \mathbf{u}^{OP}_{t-1}, (f^{OP}_{|O_l|+1,t-1}, \cdots, f^{OP}_{M,t-1}))$ according to
6   Algorithm 1;
7   **if** $|\mathbf{S}| > 0$ *and* $|O_l| < M$ **then**
8   **for** $i = 1 : min(£, |\mathbf{S}|, M - |O_l|)$ **do**
9                 $I_l(k) = S_i$ where $k = |O_l| + i$;
10   **end**
11   Send message $\mathbf{S}' = (I_l(|O_l| + 1), \cdots, I_l(min(£, |\mathbf{S}|, M - |O_l|)))$ to user l; $O_l = O_l \cup \mathbf{S}'$;
12   **else**
13   Label user l as inactive;
14   **end**
15   **end**
16   Compute $\mathbf{S}^t = \arg\max_{\mathbf{S},m} \quad E[U(\mathbf{S}, \mathbf{u}^{OP}_{t-1}, F^{OP}_{t-1})]$ according to Algorithm 1 ;
17   Offer $\mathbf{S}^t$ to user t; $O_t = \mathbf{S}^t$;
18   **for** *Any feedback for some message i as the j^{th} message* **do**
19   Update $u^{OP}_{i,t}$ and $f^{OP}_{j,t}$ according to Equation (2.4.1) and (2.4.2);
20   **end**
21   **for** *Any user l abandons the platform* **do**
22   Label user l as inactive;
23   **end**
24     $t = t + 1$;
25 **end**

In Algorithm 2, we define $I_l(j)$ as the index of message sent to user l at time t as the j^{th} message. O_l denotes the set of messages that has been sent to user l. In line 3-11, for any active user at time t, the platform re-solves the optimization problem using Algorithm 1 and determines an updated sequence selected from the pool of messages that excludes O_l. If the updated optimal sequence is empty or if we have extinguished the message budget M, we terminate the recommendation process, and label this user as inactive. No further message will be sent to her. In line 12-13, we calculate the optimal sequence using Algorithm 1 for the incoming user and send the first message. In line 14-16, we update the optimistic estimators. Lastly, in line 17-19, if any user abandons the platform, we label user l as inactive.

2.4.3 Regret analysis on Algorithm 2.

In this section, we analyze the performance of Algorithm 2. The regret for a policy π is defined as

$$Regret_\pi(T; \mathbf{u}, F) = E_\pi \left[\sum_{l=1}^{T} U(\mathbf{S}^*, \mathbf{u}, F) - U(\mathbf{S}^l, \mathbf{u}, F) \right],$$

where $\mathbf{S}^*$ is the optimal sequence when $\mathbf{u}$ and F are known to the platform, while $\mathbf{S}^l$ is the sequence offered to user l. E_π denotes the expectation under the policy π. Without loss of generality, we assume the optimal sequence is $\mathbf{S}^* = (1, 2, \cdots, m^*)$.

For the regret analysis, we make the following assumption on the abandonment distribution.

Assumption 2.4.3 *$f_j \geq f_{j-1}$ for $1 < j \leq M$.*

It states that the abandonment probability is non-increasing in i. In other words, as a user receives more unsatisfactory messages, she is more likely to abandon the platform. One example that satisfies this assumption is when W follows the geometric distribution with parameter p. In this case, $f_i = p$ for all i, i.e., the abandonment probability is independent of the number of unsatisfactory messages. We give more examples such as truncated Poisson distributions in Section 2.6.

To analyze the regret, we need to establish the following results which provide the concentration bound on $u_{i,t}^{OP}$ and $f_{j,t}^{OP}$.

Define the epoch for user l, denoted as $\mathfrak{M}_l$, as timestamps that user l receives messages from the platform. That is, $t \in \mathfrak{M}_l$ if a message is sent to user l at time t. Define $t \in \mathfrak{M}_l^j$, if the j^{th} message is sent to user l at time t. Lemma 2.4.4 below states a critical result for the regret analysis, which quantifies the difference in the expected payoff between the j^{th} message in the offered sequence under our proposed algorithm and the j^{th} message in the optimal sequence.

Lemma 2.4.4 *Define ρ_l^j as the timestamp of sending the j^{th} message to user l. For $t \in \mathfrak{M}_l^j$ with $j \geq 1$,*

$$E_\pi[(U(j, \mathbf{u}, f_j) - U(I_l(j), \mathbf{u}, f_j))1(\mathbf{u}_t^{OP} \geq \mathbf{u})] \leq (1 + c)E_\pi \left[\left(u_{I_l(j),t}^{OP} - u_{I_l(j)} \right) 1(\mathbf{u}_t^{OP} \geq \mathbf{u}) \right],$$

where $I_l(j)$ is the index of the j^{th} message sent to user l at time ρ_l^j.

In Theorem 6 below, we characterize a regret bound of our online learning algorithm for the non-contextual SC-Bandit problem. As mentioned earlier, our setting differs from the existing work on cascading bandits as the offered sequence for an individual user may change before the recommendation process terminates. To facilitate the regret analysis, we utilize a novel approach - we couple the two recommendation processes which offer sequences based on our algorithm and the optimal sequence respectively to quantify the differences in the expected payoff. For more information on the coupling technique, we refer the reader to Section 2.2 in [114].

Theorem 6 (Performance for Algorithm 2) *Given message attraction probabilities* $\mathbf{u}$ *and abandonment distribution* F*, the regret of the policy during time* T *is bounded by*

$$Regret_\pi(T; \mathbf{u}, F) \leq C\sqrt{NMT \log T}$$

for some constant C*, where* N *is the total number of available messages and* M *is the maximum possible number of messages that can be sent to a single user.*

A special case to our setting is without the abandonment behavior which we call cascading bandits. [23] establishes a lower bound for cascading bandits as $O(\sqrt{NT/M})$. Therefore the regret of our algorithm matches with the lower bound in terms of N and T up to a multiplicative logarithmic factor.

2.5 Personalization with Contextual SC-Bandit.

Thus far, we have developed a learning algorithm for the non-contextual SC-Bandit problem, under a setting where the abandonment distribution F is homogeneous across users and the attraction probability u_i is independent of message features. As a result, the optimal sequence of recommendations is identical for all users. In this section, we consider a more realistic setting where the platform needs to offer personalized recommendations to individuals users.

2.5.1 Setting for contextual SC-Bandit.

We assume that the abandonment behavior differs across users and depends on some user features $\mathbf{x} \in \mathbb{B}^{d_X}$, where $\mathbb{B} \subseteq \mathbb{R}$ denotes a compact set in $\mathbb{R}$ and d_X denotes the dimension of the user's feature space. Meanwhile, the attraction probability u_i depends on the features of message i, $\mathbf{z}_i \in \mathbb{B}^{d_Z}$, where $\mathbb{B} \subseteq \mathbb{R}$ and d_Z denotes the dimension of the message feature space.

In the existing work on cascading bandits (e.g., [72, 83]), linear reward which is a function of message contexts has been used. However, in our setting, since the feedback we observe on users' click and abandonment behavior is binary, we use a logit function to model the relationship between contexts and feedback. To the best of our knowledge, we are the first to analyze contextual cascading bandits with generalized linear function. Define the logit function as $\mu(y) = \exp(y)/(1 + \exp(y))$, and $\phi_j(\mathbf{x}_l)$ and $h(\mathbf{z}_i)$ as

$$\phi_j(\mathbf{x}_l) = \mathbf{x}_l'\alpha_j^* \qquad \text{and} \qquad h(\mathbf{z}_i) = \mathbf{z}_i'\beta^*.$$

The probability that a user with context $\mathbf{x}_l$ abandons the platform upon receiving the j^{th} unsatisfactory message is

$$f_j(\mathbf{x}_l) = P(\text{abandon at the } j^{th} \text{ unsatisfactory message}|\mathbf{x}_l) = e^{\mathbf{x}_l'\alpha_j^*}/(1+e^{\mathbf{x}_l'\alpha_j^*}) = \mu(\phi_j(\mathbf{x}_l)). \tag{2.5.1}$$

Similarly, the attraction probability of message i with feature $\mathbf{z}_i$ is

$$u(\mathbf{z}_i) = P(\text{message } i \text{ is attractive}|\mathbf{z}_i) = e^{\mathbf{z}_i'\beta^*}/(1+e^{\mathbf{z}_i'\beta^*}) = \mu(\mathbf{z}_i'\beta^*) = \mu(h(\mathbf{z}_i)), \tag{2.5.2}$$

where α_j^* and β^* are the unknown parameters to be learned for $j = 1, \cdots, M$.

Remark 2.5.1 *Our model can be extended to the setting where the attraction probability u_i depends on both the message and user contexts. Suppose the contexts of user t and message i are $\mathbf{x}_t$ and $\mathbf{z}_i$ respectively. Similar to the approach introduced in [72], we construct a new feature vector as the outer product of the two contexts, i.e., $\mathbf{v}_{t,i} = \mathbf{x}_t\mathbf{z}_i'$, and model the attraction probability u_i as a logit function of $\mathbf{v}_{t,i}$. By doing so, the model is capable of capturing the interactions between user and message contexts. However, it also increases the feature space and the number of unknown parameters which need to be learned. To enhance the clarity of our analysis, we focus on the setting where u_i only depends on the message contexts in this paper.*

2.5.2 Algorithm for contextual SC-Bandit.

In this section, we first characterize the maximum likelihood estimator for α_j and β respectively, then describe the concentration property of the estimator, and propose an optimistic algorithm inspired by the generalized linear bandit problem studied in [68]. In addition to our novel setting with abandonment behavior and adaptive user-level recommendation, to the best of our knowledge, we are the first to analyze the contextual learning-to-rank problems with a generalized linear reward function. The techniques we develop here have the potential to be lent to other learning-to-rank applications due to the prevalence of binary feedback (e.g., click, like, download).

Suppose $Y_{l,i}$ is the feedback of user l when message $I_l(i)$ is shown. $Y_{l,i} = 1$ indicates that user l accepts message $I_l(i)$ upon examining the i^{th} message, and $Y_{l,i} = 0$ otherwise. Then $Y_{l,i}$ is a Bernoulli random variable with mean $\mu(\mathbf{z}'_{I_l(i)}\beta^*)$, i.e., $P(Y_{l,i} = 1) = \mu(\mathbf{z}'_{I_l(i)}\beta^*)$ and $P(Y_{l,i} = 0) = 1 - \mu(\mathbf{z}'_{I_l(i)}\beta^*)$. Equivalently, $Y_{l,i} = \mu(\mathbf{z}'_{I_l(i)}\beta^*) + \epsilon_{l,i}$. Since we have a bounded reward, then the noise $\epsilon_{l,i}$ is bounded and hence $\epsilon_{l,i}$ is sub-gaussian with parameter σ (in our case, $\sigma = 1$ where we give detailed proof in Lemma B.5.1).

The log-likelihood function of β at time t can be written as follows,

$$\log L_t(\beta) = \sum_{l=1}^{t}\sum_{i=1}^{|\mathbf{S}^l|} \log\left(\frac{e^{\mathbf{z}'_{I_l(i)}\beta}}{1+e^{\mathbf{z}'_{I_l(i)}\beta}}1(Y_{l,i}=1) + \frac{1}{1+e^{\mathbf{z}'_{I_l(i)}\beta}}1(Y_{l,i}=0)\right)1(\text{user } l \text{ examines message } I_l(i)),$$

where $1(\cdot)$ is the indicator function, which indicates that only users who view messages provide information on the estimation of parameter β. The maximum likelihood estimator for β at time t is defined as $\hat{\beta}_t = \text{argmax}_\beta \log L_t(\beta)$, which can be re-written as,

$$\hat{\beta}_t = \underset{\beta}{\text{argmax}} \sum_{l=1}^{t} \sum_{i=1}^{|\mathbf{S}^l|} \left(Y_{l,i} \mathbf{z}'_{I_l(i)} \beta - \log \left(1 + e^{\mathbf{z}'_{I_l(i)}\beta} \right) \right) 1(\text{user } l \text{ examines message } I_l(i)).$$

Similarly, define $\hat{\alpha}_{j,t}$ as the maximum likelihood estimator for α_j at time t. Only users who face the choice to abandon the platform when receiving the j^{th} unsatisfactory message provide information on this estimator. In our model, the user faces the choice of abandonment each time when she finds the message unattractive. Define $\hat{Y}_{l,j}$ as the decision on abandonment made by user l who rejects the j^{th} message, i.e., $\hat{Y}_{l,j} = 0$ indicates that user l does not abandon the platform after finding j^{th} message unattractive, while $\hat{Y}_{l,j} = 1$ means the user abandons. Then $\hat{Y}_{l,j}$ is a Bernoulli random variable with mean $\mu(\mathbf{x}'_l \alpha^*_j)$, i.e., $P(\hat{Y}_{l,j} = 1) = \mu(\mathbf{x}'_l \alpha^*_j)$ and $P(\hat{Y}_{l,j} = 0) = 1 - \mu(\mathbf{x}'_l \alpha^*_j)$.

Therefore, the maximum likelihood estimator at time t, $\hat{\alpha}_{j,t}$, can be obtained as

$$\hat{\alpha}_{j,t} = \arg\max_{\alpha} \sum_{l=1}^{t} \left(\hat{Y}_{l,j} \mathbf{x}'_l \alpha - \log(1 + e^{\mathbf{x}'_l \alpha}) \right) 1(\text{user } l \text{ rejects the } j^{th} \text{ message before } t).$$

We define the positive semidefinite matrix $M_{j,t}$ and V_t as follows, where

$$M_{j,t} = \sum_{l=1}^{t} \mathbf{x}_l \mathbf{x}'_l 1(\text{user } l \text{ rejects the } j^{th} \text{ message before } t) \quad \text{and} \tag{2.5.3}$$

$$V_t = \sum_{l=1}^{t} \sum_{i=1}^{|\mathbf{S}^l|} \mathbf{z}_{I_l(i)} \mathbf{z}'_{I_l(i)} 1(\text{user } l \text{ examines message } I_l(i) \text{ before } t). \tag{2.5.4}$$

The "exploration bonus" term for the estimator of $\mu(\mathbf{z}'_i\beta)$ is defined as $\omega_{i,t} := \rho_{Z,t} \|\mathbf{z}_i\|_{V_t^{-1}}$ and that for the estimator of $\mu(\alpha'_j \mathbf{x}_l)$ is $\tilde{\omega}_{j,t}(\mathbf{x}_l) := \rho_{X,t} \|\mathbf{x}_l\|_{M_{j,t}^{-1}}$, respectively, where $\|\mathbf{x}\|_A$ denotes the matrix norm and it equals $\sqrt{\mathbf{x}' A \mathbf{x}}$ for any matrix A. ρ_Z and ρ_X are two constants. Finally, define the optimistic estimator for $\phi_j(\mathbf{x})$ as $\phi^{OP}_{j,t}(\mathbf{x})$, and that for $h(\mathbf{z})$ as $h^{OP}_{i,t}(\mathbf{z}_i)$, where

$$\phi^{OP}_{j,t}(\mathbf{x}) = \mathbf{x}' \hat{\alpha}_{j,t} - \rho_{X,t} \|\mathbf{x}\|_{M_{j,t}^{-1}} \quad \text{and} \quad h^{OP}_{i,t}(\mathbf{z}_i) = \mathbf{z}'_i \hat{\beta}_t + \rho_{Z,t} \|\mathbf{z}_i\|_{V_t^{-1}}, \tag{2.5.5}$$

and $\rho_{X,t} = \frac{3\sigma}{\eta_X} \sqrt{\frac{1}{2} \log t}$, $\rho_{Z,t} = \frac{3\sigma}{\eta_Z} \sqrt{\frac{1}{2} \log(Nt)}$.

We propose Algorithm 3 for the contextual SC-Bandit problem. In line 5-7, the platform explores to gather information for β. In the main loop between line 3-32, the algorithm seeks the optimal message for active user l. Inside of the loop, at stage s, we set the confidence interval at stage s to be 2^{-s}. The "exploration bonus" term for the estimator of $\mu(\mathbf{z}'_i\beta)$

at stage s is defined as $\omega_{i,t}^{(s)} := \rho_{Z,t}\|\mathbf{z}_i\|_{V_{s,t}^{-1}}$ where $V_{s,t} = \sum_{t\in\Psi_s(T)}\sum_{i=1}^{N} 1(i \in \mathcal{A}_t)\mathbf{z}_{i,t}\mathbf{z}'_{i,t}$ and $\mathcal{A}_t$ is the index of messages sent out at time t. If the width $\omega_{i,t}^{(s)}$ is larger than 2^{-s} for some $i \in X\backslash O_l$, we need to do more exploration on $\mathbf{z}_i$. Therefore, we send message i to user. Otherwise messages are filtered in line 27, and only those messages whose attraction probabilities are close enough to the highest attraction probability are passed to the next stage. If we have already obtained accurate estimate of all $\mathbf{z}'_i\beta$ up to the $1/\sqrt{T}$ level, exploration is not required and we calculate the optimal message sequence based on the current optimistic estimators in line 15. If the updated sequence obtained after solving the constrained optimization which excludes previously shown messages is not empty, we send its first message to user l. Otherwise, we label user l as *inactive* and the recommendation process for this user terminates. Lastly, in line 33-36, the optimistic estimators are updated when user's feedback is received.

2.5.3 Regret analysis on Algorithm 3.

In this section, we analyze the performance of Algorithm 3 in terms of its regret. We first make the following assumptions which are fairly standard in the contextual bandit literature (see [68]).

Assumption 2.5.2 *There exist* $0 < \eta_X < 1$ *and* $0 < \eta_Z < 1$ *such that*

$$\eta_X \le \mu(\mathbf{x}'\alpha_j^*) \le 1-\eta_X, \text{ for all } \mathbf{x} \in \mathbb{B}^{d_X} \quad \text{and} \quad \eta_Z \le \mu(\mathbf{z}'\beta^*) \le 1-\eta_Z, \text{ for all } \mathbf{z} \in \mathbb{B}^{d_Z}, \text{ and}$$

$$\inf_{\mathbf{x}\in\mathbb{B}^{d_X}, \|\alpha_j-\alpha_j^*\|^2\le 1} \dot{\mu}(\mathbf{x}'\alpha_j) > \eta_X \quad \text{and} \quad \inf_{\mathbf{z}\in\mathbb{B}^{d_Z}, \|\beta-\beta^*\|^2\le 1} \dot{\mu}(\mathbf{z}'\beta) > \eta_Z,$$

for $1 \le j \le M$, *where* $\dot{\mu}(\cdot)$ *is the first derivative.*

Assumption 2.5.3 *Define* $W(\mathbf{x})$ *as the random variable of the user's tolerance with context* $\mathbf{x}$. *The following inequalities hold*

$$0 < \lambda_0 \le \lambda_{min}(E[\mathbf{x}\mathbf{x}'1(W(\mathbf{x}) \ge j)]) \le \lambda_{max}(E[\mathbf{x}\mathbf{x}'1(W(\mathbf{x}) \ge j)]) \le R, \text{ and}$$

$$0 < \lambda_0 \le \lambda_{min}(E[\mathbf{z}\mathbf{z}']) \le \lambda_{max}(E[\mathbf{z}\mathbf{z}']) \le R,$$

for any $1 \le j \le M$ *where* λ_{min} *is the minimum eigenvalue and* λ_{max} *is the maximum eigenvalue.*

Assumption 2.5.4 *The feature space is restricted within a unit ball, i.e.,*

$$\|\mathbf{x}\|_2^2 \le 1, \quad \text{for all } \mathbf{x} \in \mathbb{B}^{d_X} \quad \text{and} \quad \|\mathbf{z}\|_2^2 \le 1, \quad \text{for all } \mathbf{z} \in \mathbb{B}^{d_Z}.$$

Lemma 2.5.5 is a property of the logit function, which can be easily verified.

Lemma 2.5.5 μ *is twice differentiable. Its first and second order derivatives are upper-bounded by* $L_\mu = 1/4$ *and* $M_\mu = 1/4$, *respectively.*

Algorithm 3: An online learning algorithm for contextual SC-Bandit under marketing fatigue

```
Initialization: tuning parameter ξ_Z, ξ_X; t = 1; ν = log T;
Set Ψ_0 = Ψ_1 = ··· = Ψ_ν = ∅, and Ψ̃^j_ν = ∅ for j = 1, ··· , M;
while t < T do
    for Any active user l with message scheduled to be sent at time t and |O_l| < M do
        if l ≤ ξ_Z then
            Randomly choose message i ∈ X\O_l for user l; O_l = O_l ∪ {i};
        else
            A_1 = X\O_l; s = 1;
            while no message is selected for user l yet and user l is active do
                Calculate ω^(s)_{i,t}(z_i) for all i ∈ (X\O_l) ∩ A_s;
                if ω^(s)_{i,t}(z_i) > 2^{-s} for some i then
                    Send message i to user l; Ψ_s = Ψ_s ∪ {i};
                else
                    if ω^(s)_{i,t}(z_i) ≤ 1/√T for all i ∈ X\O_l then
                        Compute
                         S = arg max_{S⊂X\O_l}   E[U(S, u^{OP}_{t-1}, (f^{OP}_{|O_l|+1,t-1}, ··· , f^{OP}_{M,t-1}))]
                        according to Algorithm 1; Ψ_0 = Ψ_0 ∪ {t};
                        if |S| > 0 and |O_l| < M then
                            for i = 1 : min(£, |S|, M − |O_l|) do
                                I_l(k) = S_i where k = |O_l| + i;
                            end
                            Send message
                             S′ = (I_l(|O_l| + 1), ··· , I_l(min(£, |S|, M − |O_l|))) to
                            user l; O_l = O_l ∪ S′;
                        else
                            Label user l as inactive;
                        end
                    else
                        A_{s+1} = {i ∈ A_s, μ(z′_{i,t} β̂_t) > max_{k∈A_s} μ(z′_{k,t} β̂_t) − 2^{-s+1}};
                         s = s + 1;
                    end
                end
            end
        end
    end
    for Any feedback for some message i as the j^th message do
        Update h^{OP}_{i,t} and φ^{OP}_{j,t} according to Equation (2.5.5);
        Update f^{OP}_{j,t} = μ(φ^{OP}_{j,t}) and u^{OP}_{i,t} = μ(h^{OP}_{i,t});
    end
    t = t + 1;
end
```

The next result, Lemma 2.5.6, is a critical step for the analysis of the regret bound, where we prove that the estimated quantities on the attraction probability and abandonment distribution are close enough to their true values with high probability. In addition, we also show that the minimum eigenvalue of the empirical covariance matrix is large enough for precise estimation with high probability.

Lemma 2.5.6 *Set* $\xi_X = \sqrt{d_X T}$, $\xi_Z = \sqrt{d_Z T}$, $\rho_{X,t} = \frac{3\sigma}{\eta_X}\sqrt{\frac{1}{2}\log t}$ *and* $\rho_{Z,t} = \frac{3\sigma}{\eta_Z}\sqrt{\frac{1}{2}\log(Nt)}$ *where* σ *is the sub-gaussian parameter of the noise. Suppose* t *satisfies condition that* $t \geq T_0$ *where*

$$T_0 = \min_s \left\{ s : \frac{1}{2}\lambda_0\sqrt{d_X s'} \geq \frac{512 M_\mu^2 \sigma^2}{\eta_X^4}\left(d_X^2 + \frac{1}{2}\log s'\right) \right.$$
$$\left. \text{and } \frac{1}{2}\lambda_0\sqrt{d_Z s'} \geq \frac{512 M_\mu^2 \sigma^2}{\eta_Z^4}\left(d_Z^2 + \frac{1}{2}\log s'\right), \forall s' \geq s \right\}.$$

Define the following events:

$$\mathcal{E}_{Z,t} := \{|\mathbf{z}_i'\hat{\beta}_t - \mathbf{z}_i'\beta^*| \leq \omega_{i,t}(\mathbf{z}), \forall i \in [N]\}, \forall t \in [\xi_Z + 1, T], \text{ and}$$
$$\mathcal{E}_{X,t}^j := \{|\mathbf{x}'\hat{\alpha}_{j,t} - \mathbf{x}'\alpha_j^*| \leq \tilde{\omega}_{j,t}(\mathbf{x})\}, \forall \mathbf{x} \in \mathbb{B}^{d_X}, \forall t \text{ s.t. } \tilde{T}_j(t) \geq \xi_X + 1;$$
$$\mathcal{P}_{Z,t} = \left\{\lambda_{min}(V_t) > \frac{1}{2}\lambda_0\xi_Z\right\}, \forall t \in [\xi_Z + 1, T], \text{ and}$$
$$\mathcal{P}_{X,t}^j = \left\{\lambda_{min}(M_{j,t}) > \frac{1}{2}\lambda_0\xi_X\right\}, \forall t \text{ s.t. } \tilde{T}_j(t) \geq \xi_X + 1.$$

Then, event $\mathcal{E}_{Z,t}$ *holds with probability at least* $1 - 3/\sqrt{t} - d_Z\left(\frac{e}{2}\right)^{-\lambda_0\sqrt{d_Z T}/(2R)}$, *event* $\mathcal{E}_{X,t}$ *holds with probability at least* $1 - 3/\sqrt{t} - d_X\left(\frac{e}{2}\right)^{-\lambda_0\sqrt{d_X T}/(2R)}$, *event* $\mathcal{P}_{Z,t}$ *holds with probability at least* $1 - d_Z\left(\frac{e}{2}\right)^{-\lambda_0\sqrt{d_Z T}/(2R)}$, *and event* $\mathcal{P}_{X,t}^j$ *holds with probability at least* $1 - d_X\left(\frac{e}{2}\right)^{-\lambda_0\sqrt{d_X T}/(2R)}$.

The complete proof can be found in Appendix B.5.

The following lemma provides an upper bound for the exploration terms. This result is in a similar vein to Lemma 2 in [68]. However, we are dealing with a more general setting where we allow the platform to make multiple actions at the same time (i.e., update recommendations for both the incoming and existing active users).

Lemma 2.5.7 *Let* $\{X_{i,t}\}_{t=1}^\infty$ *be a sequence in* $\mathbb{R}^d$ *satisfying* $\|X_{i,t}\| \leq 1$. *Define* $X_0 = \mathbf{0}$ *and* $V_t = \sum_{s=1}^{t-1}\sum_{i\in\mathcal{A}_s} X_{i,s}X_{i,s}'$. *Suppose there is an integer* r *such that* $\lambda_{min}(V_r) \geq 1$, *then for all* $n > 0$,

$$\sum_{t=r+1}^{r+n}\sum_{i\in\mathcal{A}_t}\|X_{i,t}\|_{V_t^{-1}} \leq \sqrt{2nM^2 d\log\left(\frac{(n+r)M}{d}\right)},$$

where $|\mathcal{A}_t| \leq M$ *for any* t.

We now present the regret bound for the contextual SC-Bandit problem in the following theorem.

Theorem 7 (Performance for Algorithm 3) *Under Assumptions 1-4, the regret of our policy during time T is bounded by*

$$Regret_\pi(T;\mathbf{u},F) \le CM\left(\sqrt{\log(T/d_X)}\sqrt{d_X T\log T} + \sqrt{\log(T/d_Z)}\sqrt{d_Z T\log(NT)}\right),$$

for some constant C.

Compared to the non-contextual SC-Bandit problem, the contextual setting offers significantly more flexibility. For instance, the pool of available messages may change with time. With each new message, the platform only needs to use the current estimator and the contexts of the new message to estimate its attraction probability. Meanwhile, the platform is also capable of recommending personalized content which takes into account of a user's abandonment behavior estimated from her features.

When the model does not incorporate user's abandonment behavior, the problem is reduced to contextual cascading bandits. In the existing literature studying the contextual cascading bandits [72, 116, 71, 23], a common assumption is that the click probability is a linear function of message feature. However, this assumption is not very appropriate because the click probability ranges from 0 to 1 while linear model cannot guarantee that. We relax this assumption that the click probability is a generalized linear function of message feature. The following table shows the regret comparison between different works. Note that the regret dependence on d in our work is $\sqrt{d}$.

Reference	Model	Bound
[72]	linear	$O(d\sqrt{TM}\log(T))$ [α-regret]
[116]	linear	$O(Md\sqrt{T}\log(MT))$
[71]	linear	$O(d\sqrt{TM}\log(T))$
[23]	linear	$O(\min\{\sqrt{d},\sqrt{\log N}\}Md\sqrt{T}(\log T)^{3/2})$
this work	generalized linear	$O(M\sqrt{dT}\log(NT)\log(T/d))$

2.6 Computational Studies.

In this section, we perform several numerical studies to evaluate our proposed algorithms. We first use synthetic data to study the robustness of our algorithms. Next, we investigate a real-world dataset which reveals different abandonment behavior among users. We then perform several experiments with this data as the ground truth in both the non-contextual and contextual settings, and compare the performance with benchmarks. Due to the space limit, we move the robustness study and numerical experiments on non-contextual SC-Bandit (real data) to Appendix B.7.

2.6.1 Sequential update versus batch update.

One novelty of our learning problem is that it can dynamically adjust the recommendation for a single user when more feedback is obtained. Meanwhile, in the existing learning-to-rank literature, recommendations are static for an individual user and are only updated across users. We will use the following experiment to compare our proposed sequential update strategy with the batch update method (in the sense that the latter only updates its strategy after the entire "batch" or sequence of a user's feedback is received).

We consider a setting with $N = 50$ and the cost of abandonment $c = 0.5$. The attraction probability $\mathbf{u}$ is uniformly generated from [0,0.2]. The abandonment distribution is drawn from the truncated Poisson distribution with $\lambda = 10$. In contrast to Algorithm 2 which updates the sequence for all active users whenever a feedback is received, a batch update algorithm determines the message sequence for a new user arriving at time t and will not make further changes to the list.

Figure 2.2 shows the comparison between two algorithms. The average regret from 50 simulations in the first $T = 50000$ rounds is 418.13 for the sequential update strategy (Algorithm 2) and 500.30 for the batch update algorithm. It shows that the regret achieved by the sequential update strategy is about 20% lower than that of the batch update approach, demonstrating the benefits of frequent update on the estimators.

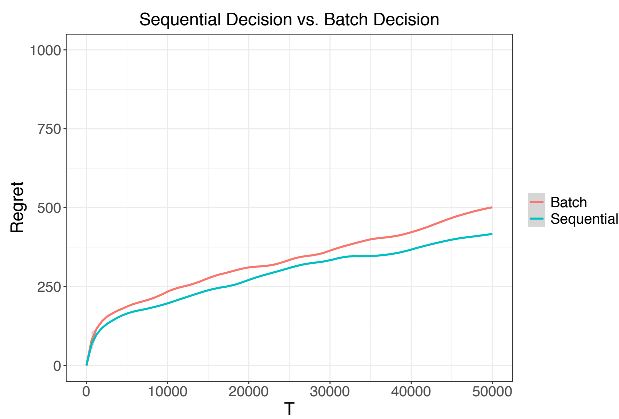

Figure 2.2: Comparison between sequential update and batch update.

2.6.2 Experiments with real data.

2.6.2.1 Data overview.

For this experiment, we utilize a dataset from Taobao[8], one of the world's biggest e-commerce websites owned by Alibaba. The dataset consists of multiple files, one of which contains 26 million ad display and click logs from 1,140,000 randomly sampled users from the website of

User ID	Time stamp	Ad ID	Click
449818	1494638778	3	0
914836	1494650879	4	1

Table 2.1: Sample of user behavior log data. The time stamp “1494137644” represents “2017-05-07 02:14:04”.

Taobao for 8 days (5/6/2017-5/13/2017) . The raw behavior log data includes time stamped record of ads shown to a user and the user’s response in terms of clicks (see Table 2.1 for some examples).

In addition, the dataset contains a file on user features, which includes gender, age group and shopping level. The summary statistics of the user features are shown in Table 2.2.

Feature	Summary statistics
Gender	Male (35.6%), Female (64.4%)
Age level	0 (0.05%),1 (6.17%),2 (17.86%),3 (28.95%),4 (24.65%),5 (20.20%),6 (2.12%)
Shopping level	Shallow (7.03%), Moderate (14.26%), Deep (78.71%)

Table 2.2: User features overview.

The dataset also includes information on ads in terms of its category and the price of the product (in Chinese Yuan, ¥) being advertised. For our experiments, we focus on one of the categories with the highest number of products, i.e., category 1665 which includes over 15,000 products. We use price as the item feature. Figure B.5a (refer to Appendix B.7) shows the distribution of prices for products in this category, where the mean is ¥274 and more than 95% products are priced below ¥700.

2.6.2.2 Measure abandonment.

We use inactivity from the last interaction with the website as a proxy for abandonment. Data reveals that users interacted frequently with the website, i.e., the 50th and 95th percentile of the time interval between two consecutive visits to the website are 4 hours 24 minutes and 2 days 13 hours and 8 minutes respectively. We believe that such a phenomenon is driven by the fact that 78.8% of users in the dataset are “deep shoppers”, i.e., frequent shoppers at Taobao whose purchase frequency and total expenditure have reached the highest status. Thus, for our experiment, we set the threshold on inactivity for abandonment as 3 days. That is, we consider a user “abandoned” the website if it has been more than 3 days of inactivity since her last visit.

To fit the data into our framework, for each user, we construct the entire ad response as a sequence. Depending on when the last interaction takes place, we label whether a user abandoned the website. For instance, a sequence "001100A" indicates 6 interactions a user had with the website, where she clicked on the 3rd and 4th ads and abandoned eventually. To handle the situation that a user can select multiple items as the example shown earlier, we separate that response into multiple sequences once a user clicked on an ad, i.e., "001100A" is broken into "001", "1", "00A"[9].

The abandonment rate is defined as the total number number of abandonments divided by the total number of ads that are not being clicked. We found that the average abandonment rate is 0.64% in this dataset. Figure 2.3 depicts how the abandonment rate varies across different sub-populations characterized by the user features. Interestingly, it shows that male users are less tolerant with ads compared to female users, i.e., the abandonment rate for male is 14.87% higher than female. Figure 2.3 also reveals that the abandonment rate decreases when a user shops more. The conjecture is that for a frequent shopper, the website has acquired a better understanding of the user's preferences and is able to recommend more relevant ads which then lead to a lower abandonment rate. The abandonment rate distribution for age follows a U-shape, indicating that on average the youngest and the oldest users are more likely to abandon the website, while the middle-aged users are least likely to do so.

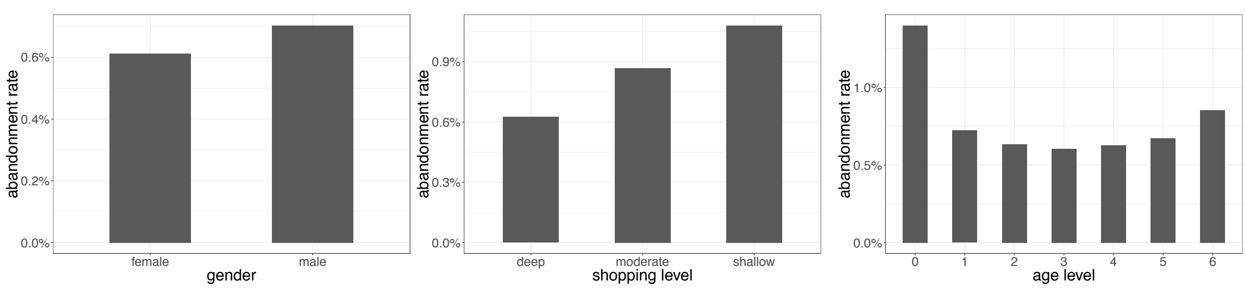

Figure 2.3: Abandonment rate versus user features.

2.6.2.3 Model calibration.

Estimate α: For simplicity, we use the geometric distribution to approximate user's tolerance. We include Gender, Age level and Shopping level as user features $\mathbf{x}$ to estimate α, and the abandonment behavior parameter $\mu(\phi(\mathbf{x}))$ is defined according to Equation (2.5.1). As the user features are categorical (see Table 2.2), we first convert them into dummy variables. The regression results are shown in Table B.1 in Appendix B.7. Note that most estimated coefficients are statistically significant, and their signs agree with the behavior observed in Figure 2.3.

Estimate β: The relationship between users' product valuation $u_i(\mathbf{z}_i)$ and β is defined in Equation (2.5.2). We use Price as the product feature, $\mathbf{z}_i$. The regression output and the fitted attraction probability can be found in Appendix B.7 as Table B.2 and Figure B.5b respectively.

2.6.2.4 Experiments on contextual SC-Bandit.

Experiment setup We use the same 100 products from category 1665 and compute the corresponding $h(\mathbf{z})$ as the ground truth as in the previous experiments. Unlike the previous experiments where we use the aggregate abandonment rate as the ground truth, we compute $\mu(\phi(\mathbf{x}))$ using the estimated α for user with feature $\mathbf{x}$. Each user's feature is sampled according to the uniform distribution in the feature space. We set $c = 6$ in two experiments.
Experiment result Figure 2.4 shows the regrets for Algorithm 3, compared with a benchmark algorithm which separates the exploration and exploitation periods. The average regret is 76.16 for Algorithm 3 and 102.69 for the benchmark algorithm. The learned parameters $\hat{\alpha}$ and $\hat{\beta}$ are very close to the true parameters α and β by the first 1000 iterations based on Algorithm 3.

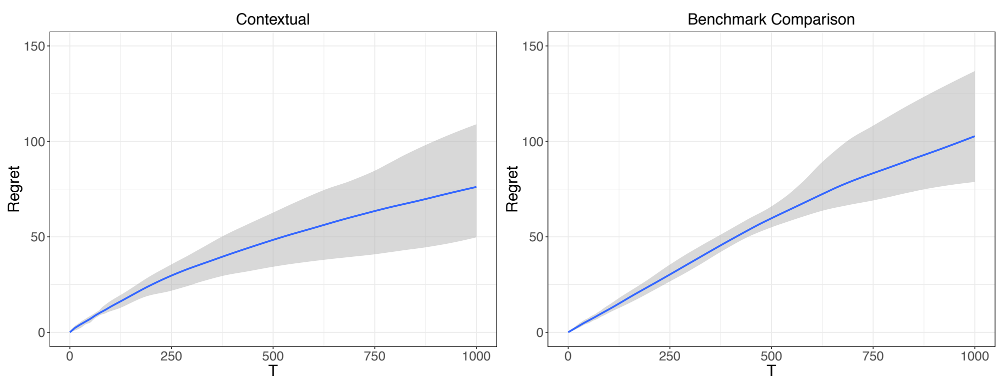

Figure 2.4: Comparison of Algorithm 3 with the benchmark.

2.7 Conclusion.

In this work, we proposed a novel model to capture sequential interactions between users and a platform. The platform earns a reward when a user clicks on the recommended content and incurs a cost when a user abandons due to marketing fatigue. As a user browses the recommended content sequentially, her list of recommendation is dynamic and adaptive to her feedback as well as learning acquired from other users.

For the offline optimization problem which is combinatorial in nature, we showed a polynomial-time algorithm to determine an optimal sequence of messages. For the online learning task which we refer to as the SC-Bandit problem, we first studied the non-contextual version. An exploration-exploitation algorithm was proposed and the regret was shown to be $O(\sqrt{NMT\log T})$. Next we studied the contextual SC-bandit problem, where a user's abandonment probability with respect to prior rejected messages depends on her features and attractiveness of a message depends on its content features. This setting allows significantly more flexibility as the pool of available messages may change with time. In addition, the platform is capable of providing personalized recommendations. An optimistic algorithm that simultaneously explores and exploits for the contextual SC-bandit problem was proposed. We showed that the regret was $O\left(M\left(\sqrt{\log(T/d_X)}\sqrt{d_XT\log T}+\sqrt{\log(T/d_Z)}\sqrt{d_ZT\log(NT)}\right)\right)$. Lastly, we evaluated the algorithms' performance with both synthetic and real-world datasets. In particular, we found empirical evidences that the abandonment behavior varies across different population groups. We investigated the robustness of our proposed algorithms and also performed benchmark comparison.

There are several future directions of this work. Firstly, as users' preferences may vary over time, it is interesting to incorporate the temporal dimension into the setting. Secondly, different user actions could reveal different levels of interest (e.g., the amount of time a user spent on a message, a user clicked on a message but did not complete a purchase etc.). One question is how to construct and analyze a more accurate user behavior model by utilizing such fine-grained behavior data. Thirdly, Thompson Sampling would be another natural algorithm to solve the problem we have proposed, especially for the personalized recommendation version. However, analyzing this setting and providing theoretical results remain a challenging problem.

Chapter 3

Connectivity of a general class of inhomogeneous random digraphs

3.1 Introduction.

Complex networks appear in essentially all branches of science and engineering, and since the pioneering work of Erdős and Rényi in the early 1960s [39, 38], people from various fields have used random graphs to model, explain and predict some of the properties commonly observed in real-world networks. Until the last decade or so, most of the work had been mainly focused on the study of undirected graphs, however, some important networks, such as the World Wide Web, Twitter, and ResearchGate, to name a few, are directed. The present paper describes a framework for analyzing a large class of directed random graphs, which includes as special cases the directed versions of some of the most popular undirected random graph models.

Specifically, we study directed random graphs where the presence or absence of an arc is independent of all other arcs. This independence among arcs is the basis of the classical Erdős-Rényi model [39, 38], where the presence of an edge is determined by the flip of coin, with all possible edges having the same probability of being present. However, it is well-known that the Erdős-Rényi model tends to produce very homogeneous graphs, that is, where all the vertices have close to the same number of neighbors, a property that is almost never observed in real-world networks. In the undirected setting, a number of models have been proposed to address this problem while preserving the independence among edges. Some of the best known models include the Chung-Lu model [27, 26, 28, 78], the generalized random graph [15, 14, 42], and the Norros-Reittu model or Poissonian random graph [82, 114, 14]. In the undirected case, all of these models were simultaneously studied in [14] under a broader class of graphs, which we will refer to as kernel-based models. In all of these models the inhomogeneity of the degrees is accomplished by assigning to each vertex a *type*, which is used to make the edge probabilities different for each pair of vertices. From a modeling perspective, the types correspond to vertex attributes that influence how likely a vertex is to have neighbors, and inhomogeneity among the types translates into inhomogeneous degrees.

Our proposed family of directed random graphs, which we will refer to as *inhomogeneous random digraphs*, provides a uniform treatment of essentially any model where arcs are present independently of each other, in the same spirit as the work in [14] written for the undirected case. The main results in this paper establish some of the basic properties studied on random graphs, including the expected number of arcs, the joint distribution of the in-degree and out-degree, and the phase transition for the size of the largest strongly connected component. We pay special attention to the so-called *scale-free* property, which states that the tail degree distribution(s) decay according to a power law. Since many real-world directed complex networks exhibit the scale-free property in either their in-degrees, their out-degrees, or both, we provide a theorem stating how the family of random directed graphs studied here can be used to model such networks. Our main result on the connectivity properties of the graphs produced by our model shows that there exists a phase transition, determined by the types, after which the largest strongly connected component contains (with high probability) a positive fraction of all the vertices in the graph, i.e., the graph contains a "giant" strongly connected component.

That the undirected models mentioned above satisfy these basic properties (e.g., scale-free degree distribution, existence of a giant connected component, etc.) constitutes a series of classical results within the random graph literature. Closely related to the results presented here for directed graphs, are the existence of a giant strongly connected component and giant weak-component in the directed configuration model [**Cooper2004size**, 59, 60], the existence of a giant strongly-connected component in the deterministic directed kernel model with a finite number of types [12], the scale-free property on a directed preferential attachment model [100, 93], and the limiting degree distributions in the directed configuration model [19][10]. From a computational point of view, the work in [110] provides numerical algorithms to identify secondary structures on directed graphs. Our present work includes as a special case the main theorem in [12] and extends it to a larger family of directed random graphs, and it also compiles several results for the number of arcs and the joint distribution of the degrees. It is also worth pointing out that the directed nature of our framework introduces some non-trivial challenges that are not present in the undirected setting, which is the reason we chose to provide a different approach from the one used in [14] for establishing some of our main results. We refer the reader to Section 3.3.3 for more details on these challenges and what they imply.

The paper is organized as follows. In Section 3.2 we specify a class of directed random graphs via their arc probabilities, and explain how the models mentioned above fit into this framework. In Section 3.3 we provide our main results on the basic properties of the graphs produced by our model, and in Section 3.4 we give all the proofs.

3.2 The Model.

As mentioned in the introduction, we study directed random graphs with independent arcs. Since we are particularly interested in graphs with inhomogeneous degrees, each vertex in the graph will be assigned a *type*, which will determine how large its in-degree and out-degree are likely to be. In applications, the type of a vertex can also be used to model other vertex attributes not directly related to its degrees. We will assume that the types take values in a separable metric space $\mathcal{S}$, which we will refer to as the "type space".

In order to describe our family of directed random graphs, we start by defining the type sequence $\{\mathbf{x}_1^{(n)}, \dots, \mathbf{x}_n^{(n)}\}$, where $\mathbf{x}_i^{(n)}$ denotes the type of vertex i in a graph on the vertex set $[n]$. Note that, depending on how we construct the type sequence, it is possible for $\mathbf{x}_i^{(n)}$ to be different from $\mathbf{x}_i^{(m)}$ for $n \neq m$. Define $G_n(\kappa(1+\varphi_n))$ to be the graph on the vertex set $[n]$ whose arc probabilities are given by

$$p_{ij}^{(n)} = \left(\frac{\kappa(\mathbf{x}_i^{(n)}, \mathbf{x}_j^{(n)})}{n} (1 + \varphi_n(\mathbf{x}_i^{(n)}, \mathbf{x}_j^{(n)})) \right) \wedge 1, \qquad 1 \leq i \neq j \leq n, \tag{3.2.1}$$

where κ is a nonnegative function on $\mathcal{S} \times \mathcal{S}$,

$$\varphi_n(\mathbf{x}, \mathbf{y}) = \varphi\left(n, \{\mathbf{x}_k^{(n)} : 1 \leq k \leq n\}, \mathbf{x}, \mathbf{y}\right) > -1 \qquad \text{for all } \mathbf{x}, \mathbf{y} \in \mathcal{S},$$

and $x \wedge y = \min\{x, y\}$ ($x \vee y = \max\{x, y\}$). In other words, $p_{ij}^{(n)}$ denotes the probability that there is an arc from vertex i to vertex j in $G_n(\kappa(1+\varphi_n))$. The presence or absence of arc (i, j) is assumed to be independent of all other arcs. Note that the function $\varphi_n(\mathbf{x}, \mathbf{y})$ may depend on n, on the types of the two vertices involved, or on the entire type sequence; however, to simplify the notation, we emphasize only the arguments $(\mathbf{x}, \mathbf{y})$ of the two types involved. Following the terminology used in [14] and [12], we will refer to κ as the kernel of the graph. Note that we have decoupled the dependence on n and on the type sequence by including it in the term $\varphi_n(\mathbf{x}, \mathbf{y})$, which implies that with respect to the notation used in [14], $\kappa_n(\mathbf{x}, \mathbf{y})$ there corresponds to $\kappa(\mathbf{x}, \mathbf{y})(1 + \varphi_n(\mathbf{x}, \mathbf{y}))$ here.

Throughout the paper, we will refer to any directed random graph generated through our model as an *inhomogeneous random digraph* (IRD).

We end this section by explaining how the directed versions of the Erdős-Rényi graph [38, 39, 40, 13], the Chung-Lu (or "given expected degrees") model [27, 26, 28, 78], the generalized random graph [15, 14, 42], and the Norros-Reittu model (or "Poissonian random graph") [82, 114, 14], as well as the directed deterministic kernel model in [12], fit into our framework. The first four examples fall into the category of so-called rank-1 kernels, where the graph kernel is of the form $\kappa(\mathbf{x}, \mathbf{y}) = \kappa_+(\mathbf{x})\kappa_-(\mathbf{y})$ for some nonnegative continuous functions κ_- and κ_+ on $\mathcal{S}$.

Example 3.2.1 *Directed versions of some well-known inhomogeneous random graph models. All of them, with the exception of the last one, are defined on the space* $\mathcal{S} = \mathbb{R}_-$ *for a* *type of the form* $\mathbf{x} = (x^-, x^+)$*, and correspond to rank-1 kernels with* $\kappa_-(\mathbf{x}) = x^-/\sqrt{\theta}$ *and* $\kappa_+(\mathbf{x}) = x^+/\sqrt{\theta}$*, with* $\theta > 0$ *a constant. For convenience, we have dropped the superscript* $^{(n)}$ *from the type sequence, i.e.,* $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\} = \{\mathbf{x}_1^{(n)}, \ldots, \mathbf{x}_n^{(n)}\}$.

1. Directed Erdős-Rényi Model: *the arc probabilities are given by*

$$p_{ij}^{(n)} = \lambda/n$$

where λ *is a given constant and* n *is the total number of vertices;* $\varphi_n(\mathbf{x}_i, \mathbf{x}_j) = 0$.

2. Directed Given Expected Degree Model (Chung-Lu): *the arc probabilities are given by*

$$p_{ij}^{(n)} = \frac{x_i^+ x_j^-}{l_n} \wedge 1,$$

where $l_n = \sum_{i=1}^n (x_i^- + x_i^+)$*. In terms of* (3.2.1)*, it satisfies* $\varphi_n(\mathbf{x}_i, \mathbf{x}_j) = \frac{\theta n - l_n}{l_n}$*, where* $\theta = \lim_{n\to\infty} l_n/n$.

3. Generalized Directed Random Graph: *the arc probabilities are given by*

$$p_{ij}^{(n)} = \frac{x_i^+ x_j^-}{l_n + x_i^+ x_j^-},$$

which implies that $\varphi_n(\mathbf{x}_i, \mathbf{x}_j) = \frac{\theta n - l_n - x_i^+ x_j^-}{l_n + x_i^+ x_j^-}$*, with* l_n *and* θ *defined as above.*

4. Directed Poissonian Random Graph (Norros-Reittu): *the arc probabilities are given by*

$$p_{ij}^{(n)} = 1 - e^{-x_i^+ x_j^-/l_n},$$

which implies that $\varphi_n(\mathbf{x}_i, \mathbf{x}_j) = \left(n\theta(1 - e^{-x_i^+ x_j^-/l_n}) - x_i^+ x_j^-\right)/(x_i^+ x_j^-)$*, with* l_n *and* θ *defined as above.*

5. Deterministic Kernel Model: *the arc probabilities are given by*

$$p_{ij}^{(n)} = \frac{\kappa(\mathbf{x}_i, \mathbf{x}_j)}{n} \wedge 1,$$

for a finite type space $\mathcal{S} = \{\mathbf{s}_1, \ldots, \mathbf{s}_M\}$*, and a strictly positive function* κ *on* $\mathcal{S} \times \mathcal{S}$*; in terms of* (3.2.1)*,* $\varphi_n(\mathbf{x}_i, \mathbf{x}_j) = 0$*. This model is also known as the stochastic block model.*

3.3 Main Results.

We now present our main results for the family of inhomogeneous random digraphs defined through (3.2.1). As mentioned in the introduction, we focus on establishing some of the basic properties of this family, including the distribution of the degrees, the mean number of arcs, and the size of the largest strongly connected component. When analyzing the degree distributions, we specifically explain how to obtain the scale-free property under degree-degree correlations.

As mentioned in the previous section, we assume throughout the paper that the nth graph in the sequence is constructed using the types $\{\mathbf{X}_1, \ldots, \mathbf{X}_n\} = \{\mathbf{X}_1^{(n)}, \ldots, \mathbf{X}_n^{(n)}\}$, where we will often drop the superscript $^{(n)}$ to simplify the notation. From now on we will use upper case letters to emphasize the possibility that the $\{\mathbf{X}_i\}$ may themselves be generated through a random process. To distinguish between these two levels of randomness, let P be a probability measure on a space large enough to construct all the type sequences $\{\{\mathbf{X}_i^{(n)}, 1 \leq i \leq n\} : n \geq 1\}$, as well as the random graphs $G_n(\kappa(1 + \varphi_n))$, simultaneously. Define $\mathscr{F} = \sigma(\mathbf{X}_i^{(n)}, 1 \leq i \leq n)$ and the corresponding conditional probability and expectation $\mathbb{P}(\cdot) = P(\cdot|\mathscr{F})$ and $\mathbb{E}[\cdot] = E[\cdot|\mathscr{F}]$, respectively.

Our first assumption will be to ensure that the $\{\mathbf{X}_i^{(n)}\}$ converge in distribution under the unconditional probability P. As is to be expected from the work in [14] for the undirected case, we will also need to impose some regularity conditions on the kernel κ, as well as on the function φ_n. Our main assumptions are summarized below.

Assumption 3.3.1 *1. There exists a Borel probability measure μ on $\mathcal{S}$ such that for any μ-continuity set $A \subseteq \mathcal{S}$,*

$$\mu_n(A) := \frac{1}{n}\sum_{i=1}^n 1(\mathbf{X}_i^{(n)} \in A) \xrightarrow{P} \mu(A) \qquad n \to \infty,$$

where $\xrightarrow{P}$ denotes convergence in probability. Note that μ_n is a random probability measure, whereas μ is not random.

2. *κ is nonnegative and continuous a.e. on $\mathcal{S} \times \mathcal{S}$.*

3. *For any sequences $\{\mathbf{x}_n\}, \{\mathbf{y}_n\} \subseteq \mathcal{S}$ such that $\mathbf{x}_n \to \mathbf{x}$ and $\mathbf{y}_n \to \mathbf{y}$ as $n \to \infty$, we have $\varphi_n(\mathbf{x}_n, \mathbf{y}_n) \xrightarrow{P} 0$ as $n \to \infty$.*

4. *The following limits hold :*

$$\lim_{n\to\infty} \frac{1}{n^2} E\left[\sum_{i=1}^n \sum_{j=1}^n \kappa(\mathbf{X}_i^{(n)}, \mathbf{X}_j^{(n)})\right] = \lim_{n\to\infty} \frac{1}{n} E\left[\sum_{i=1}^n \sum_{j\neq i} p_{ij}^{(n)}\right] = \iint_{\mathcal{S}^2} \kappa(\mathbf{x}, \mathbf{y})\, \mu(d\mathbf{x})\mu(d\mathbf{y}) < \infty.$$

Remark 3.3.2 *The pair $(\mathcal{S}, \mu)$, where $\mathcal{S}$ is a separable metric space and μ is a Borel probability measure, is referred to in [14] as a* generalized ground space*. For convenience, we* *will adopt the same terminology throughout the paper. Throughout the paper, we use "a.e." to mean "almost everywhere with respect to the (non-random) measure μ".*

3.3.1 Number of arcs.

Our assumption that the types $\{\mathbf{X}_i\}$ converge in distribution as the size of the graph grows implies that the graphs produced by our model are sparse, in the sense that the mean number of arcs is of the same order as the number of vertices. Our first result provides an expression for the exact ratio between the number of arcs and the number of vertices.

Proposition 3.3.3 *Define $e(G_n(\kappa(1+\varphi_n)))$ to be the number of arcs in $G_n(\kappa(1+\varphi_n))$. Then, under Assumption 3.3.1(a)-(d) we have*

$$\frac{1}{n}e(G_n(\kappa(1+\varphi_n))) \longrightarrow \iint_{\mathcal{S}^2} \kappa(\mathbf{x},\mathbf{y})\,\mu(d\mathbf{x})\mu(d\mathbf{y}) \qquad \text{in } L^1(P)$$

as $n \to \infty$.

3.3.2 Distribution of vertex degrees.

We now move on to describing the vertex degree distribution, which is best accomplished by looking at the properties of a typical vertex, i.e., one chosen uniformly at random. In particular, if $D^-_{n,i}$ and $D^+_{n,i}$ denote the in-degree and out-degree, respectively, of vertex $i \in [n] \triangleq \{1, \cdots, n\}$, and we let ξ be a uniform random variable in $\{1, 2, \ldots, n\}$, then we study the distribution of $(D^-_{n,\xi}, D^+_{n,\xi})$. We point out that the distribution of $(D^-_{n,\xi}, D^+_{n,\xi})$ also allows us to compute the proportion of vertices in the graph having in-degree k and out-degree l for any $k, l \geq 0$. In the sequel, $\Rightarrow$ denotes weak convergence with respect to P.

Theorem 3.3.4 *Under Assumption 3.3.1 we have*

$$\left(D^-_{n,\xi}, D^+_{n,\xi}\right) \Rightarrow (Z^-, Z^+), \quad E[D^{\pm}_{n,\xi}] \to E[Z^{\pm}], \qquad \text{as } n \to \infty,$$

where Z^- and Z^+ are conditionally independent (given $\mathbf{X}$) mixed Poisson random variables with mixing distributions

$$\lambda_-(\mathbf{X}) := \int_{\mathcal{S}} \kappa(\mathbf{y},\mathbf{X})\,\mu(d\mathbf{y}) \quad \text{and} \quad \lambda_+(\mathbf{X}) := \int_{\mathcal{S}} \kappa(\mathbf{X},\mathbf{y})\,\mu(d\mathbf{y}),$$

respectively, and $\mathbf{X}$ is distributed according to μ.

As mentioned earlier, we are particularly interested in models capable of creating scale-free graphs, perhaps with a significant correlation between the in-degree and out-degree of the same vertex. To see that our family of inhomogeneous random digraphs can accomplish this, we first introduce the notion of non-standard regular variation (see [93, 100]), which extends the notion of regular variation in the real line to multiple dimensions, with each dimension having potentially different tail indexes. In our setting we only need to consider two dimensions, so we only give the bivariate version of the definition.

Definition 3.3.5 *A nonnegative random vector $(X, Y) \in \mathbb{R}^2$ has a distribution that is non-standard regularly varying if there exist scaling functions $a(t) \uparrow \infty$ and $b(t) \uparrow \infty$ and a non-zero limit measure $\nu(\cdot)$, called the limit or tail measure, such that*

$$tP\left((X/a(t), Y/b(t)) \in \cdot\right) \xrightarrow{v} \nu(\cdot), \qquad t \to \infty,$$

where $\xrightarrow{v}$ denotes vague convergence of measures in $M_-([0,\infty]^2 \backslash \{\mathbf{0}\})$, the space of Radon measures on $[0,\infty]^2 \setminus \{\mathbf{0}\}$.

In particular, if the scaling functions $a(t)$ and $b(t)$ are regularly varying at infinity with indexes $1/\alpha$ and $1/\beta$, respectively, that is $a(t) = t^{1/\alpha}L_a(t)$ and $b(t) = t^{1/\beta}L_b(t)$ for some $\alpha, \beta > 0$ and slowly varying functions L_a and L_b, then the marginal distributions $P(X > t)$ and $P(Y > t)$ are regularly varying with tail indexes $-\alpha$ and $-\beta$, respectively (see Theorem 6.5 in [92]). Throughout the paper we use the notation $\mathcal{R}_\alpha$ to denote the family of regularly varying functions with index α.

To see how our family of IRDs can be used to model complex networks where both the in-degrees and the out-degrees possess the scale-free property, perhaps with different tail indexes, we give a theorem stating that the non-standard regular variation of the limiting degrees (Z^-, Z^+) follows from that of the vector $(\lambda_-(\mathbf{X}), \lambda_+(\mathbf{X}))$. Moreover, for the models (a)-(d) in Example 3.2.1, we have

$$(\lambda_-(\mathbf{X}), \lambda_+(\mathbf{X})) = \left(\kappa_-(\mathbf{X}) \int_{\mathcal{S}} \kappa_+(\mathbf{y})\mu(d\mathbf{y}),\, \kappa_+(\mathbf{X}) \int_{\mathcal{S}} \kappa_-(\mathbf{y})\mu(d\mathbf{y})\right) = \left(cX^-,\, (1-c)X^+\right),$$

where $c = E[X^+]/\theta$ and $\theta = E[X^- + X^+]$, so the non-standard regular variation of (Z^-, Z^+) can be easily obtained by choosing a non-standard regularly varying type distribution μ.

Theorem 3.3.6 *Let $\mathbf{X}$ denote a random vector in the type space $\mathcal{S}$ distributed according to μ. Suppose that μ is such that $(\lambda_-(\mathbf{X}), \lambda_+(\mathbf{X}))$ is non-standard regularly varying with scaling functions $a(t) \in \mathcal{R}_{1/\alpha}$ and $b(t) \in \mathcal{R}_{1/\beta}$ and limiting measure $\nu(\cdot)$. Then, (Z^-, Z^+) is non-standard regularly varying with scaling functions $a(t)$ and $b(t)$ and limiting measure $\nu(\cdot)$ as well.*

To illustrate our result, we give below an example that illustrates how our family of random digraphs along with Theorem 3.3.6 can be used to model real-world networks.

Example 3.3.7 *As discussed in [115], many real-world networks exhibit both heavy-tailed in-degrees and heavy-tailed out-degrees. In many of those cases there also appears to be a relationship between the vertices with very high in-degrees and those with very high out-degrees, as is shown in [115] for portions of the Web graph and the English Wikipedia graph (this dependence was computed using the angular measure in [115]). Suppose we want to model such graphs using an inhomogeneous random digraph. Interesting levels of dependence ranging from the case where the in-degree and out-degree are independent to where they are* *essentially the same can be obtained by choosing* $\mathbf{X} = (X^-, X^+)$*,* $P(X^- > x) \sim k_- x^{-\alpha}$ *as* $x \to \infty$ *and* $X^+ = r(X^-)^\gamma + (1-r)Y$*, where* Y *is independent of* X^- *and satisfies* $P(Y > y) \sim k' y^{-\beta}$*,* $\alpha, \beta, k_-, k' > 0$*,* $r \in [0,1]$ *and* $0 \leq \gamma \leq \alpha/\beta$*. This choice leads to* $P(X^+ > x) \sim k_+ x^{-\beta}$ *for some other constant* $k'_+ > 0$*, and covers the independent case when* $r = 0$*, and the perfectly dependent case when* $r = 1$ *and* $\gamma = \alpha/\beta$*. Now choose* $\kappa(\mathbf{x}, \mathbf{y}) = x^+ y^-$ *and note that* $(\lambda_-(\mathbf{X}), \lambda_+(\mathbf{X})) = (cX^-, (1-c)X^+)$*, where* $c = E[X^+]/E[X^- + X^+]$*. It follows from Theorems 3.3.4 and 3.3.6 that* $(D^-_{n,\xi}, D^+_{n,\xi}) \Rightarrow (Z^-, Z^+)$ *as* $n \to \infty$*, where* (Z^+, Z^-) *is non-standard regularly varying. In particular,* $P(Z^- > z) \sim k_- c^\alpha z^{-\alpha}$ *and* $P(Z^+ > x) \sim k_+(1-c)^\beta z^{-\beta}$ *as* $z \to \infty$*, and the angular measure between* Z^- *and* Z^+ *will mimic that of* X^- *and* X^+*.*

3.3.3 Phase transition for the largest strongly connected component.

Our last result in the paper establishes a phase transition for the existence of a giant strongly connected component in $G_n(\kappa(1+\varphi_n))$. That is, we provide a critical threshold for a functional of the kernel κ and the type distribution μ, such that above this threshold the graph will have a giant strongly connected component with high probability, and below it will not. Before stating the corresponding theorem, we give a brief overview of some basic definitions.

For any two vertices i, j in the graph, we say that there is a directed path from i to j if the graph contains a set of arcs $\{(i, k_1), (k_1, k_2), \ldots, (k_t, j)\}$ for some $t \geq 0$. A set of vertices $V \subseteq [n]$ is *strongly connected*, if for any two vertices $i, j \in V$ we have that there exists a directed path from i to j and one from j to i. Moreover, we say that a *giant* strongly connected component exists for our family of random digraphs if $\liminf_{n\to\infty} |\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))|/n > \epsilon$ for some $\epsilon > 0$, where $\mathcal{C}_1(G_n(\kappa(1+\varphi_n))$ is the largest strongly connected component of $G_n(\kappa(1+\varphi_n))$ and $|A|$ denotes the cardinality of set A.

For undirected graphs, the phase transition for the Erdős-Rényi model ($p_{ij}^{(n)} = \lambda/n$ for some $\lambda > 0$) dates back to the classical work of Erdős and Rényi in [38], where the threshold for the existence of a giant connected component is $\lambda = 1$. The critical case, i.e., $\lambda = 1$, was studied in [76] using edge probabilities of the form $p_{ij}^{(n)} = (1 + cn^{-1/3})/n$ for some $c > 0$, in which case the size of the largest connected component was shown to be of order $n^{2/3}$. Somewhat unrelated, the corresponding phase transition was established for the (undirected) configuration model in [80], where the threshold was shown to be $E[D(D-1)]/E[D] = 1$, with D distributed according to the limiting degree distribution (as the number of vertices grows to infinity). Back to the (undirected) inhomogeneous random graph setting, i.e., $p_{ij}^{(n)} = \kappa(\mathbf{x}_i, \mathbf{x}_j)(1 + \varphi_n(\mathbf{x}_i, \mathbf{x}_j))/n$ with κ symmetric, the phase transition was first proven for various forms of rank-1 kernels. In particular, Chung and Lu established in [26] the phase transition for the existence of a giant connected component in the so-called "given expected degree" model. The same authors also give in [27] a phase transition for the average distance between vertices when the type distribution μ follows a power-law. Norros and Reittu proved the phase transition for the existence of a giant connected component for the Poissonian random graph in [82], along with a characterization of the distance between two randomly chosen vertices, and Riordan proved it in [95] for the $c/\sqrt{ij}$ model, which is equivalent to the rank-1 kernel $\kappa(\mathbf{x},\mathbf{y}) = \psi(\mathbf{x})\psi(\mathbf{y})$ with $\psi(\mathbf{x}) = \sqrt{c\mathbf{x}}$ and μ the distribution of a Pareto(2,1). More generally, the work in [14] gives the phase transition for the giant connected component for the general kernel case, along with some other properties (e.g., second largest connected component, distances between vertices, and stability). The threshold for the existence of a giant connected component is $\|T_\kappa\|_{op} = 1$, with $\|\cdot\|_{op}$ the operator norm[11], where T_κ is a linear operator induced by κ, which in the rank-1 case becomes $\|T_\kappa\|_{op}^2 = E[\psi(\mathbf{X})^2] = 1$, with $\mathbf{X}$ distributed according to μ.

For the directed case, the phase transition for the existence of a giant strongly connected component was proven for the directed Erdős-Rényi model ($p_{ij}^{(n)} = \lambda/n$ for some $\lambda > 0$) in [56] and for the "given number of arcs" version of the Erdős-Rényi model (number of arcs $= \lambda n$ for some $\lambda > 0$) in [74], with the threshold being $\lambda = 1$. The work in [75] studies a related model where each vertex i can have three types of arcs: up arcs for $j > i$, down arcs for $j < i$, and bidirectional arcs, and proved the corresponding phase transition for the appearance of a giant strongly connected component. For the directed configuration model the phase transition for the existence of a giant strongly connected component was given in [42] under the assumption that the limiting degrees have finite variance and satisfy some additional conditions on the growth of the maximum degree, and can also be indirectly obtained from the results in [53] under only finite covariance between the in-degree and out-degree. The threshold for the directed configuration model is $E[D^-D^+]/E[D^-+D^+] = 1$, where (D^-, D^+) are the limiting in-degree and out-degree. A hybrid model where the out-degree has a general distribution with finite mean and the destinations of the arcs are selected uniformly at random among the vertices (which gives Poisson in-degrees) was studied in [85] and was shown to have a phase transition at $E[D^+] = 1$. Finally, for general inhomogeneous random digraphs such as those studied here, the main theorem in [12] establishes the phase transition for the deterministic kernel in Example 3.2.1(d) with finite type space $\mathcal{S} = \{1, 2, \ldots, M\}$, without characterizing the strict positivity of the survival probability. The authors in [12] also suggest that the general case can be obtained using the same techniques used in [14] to go from a finite type space to the general one, however, the proof in [14] requires a critical step that does not hold for directed graphs; see Section 3.4.3 for more details.

Our Theorem 3.3.10 provides the full equivalent of the main theorem in [14] (Theorem 3.1) for the directed case, and its proof is based on a coupling argument between the exploration of both the inbound and outbound components of a randomly chosen vertex and a double multi-type branching process with a finite number of types. Our approach differs from that of [14], done for undirected graphs, in the order in which the couplings are done, and it leverages on the main theorem in [12] to obtain a lower bound for the size of the strongly connected component. We give more details on how our proof technique compares to that used in [14] in Section 3.4.3.

As in the undirected case, the size of the largest strongly connected component is related to the survival probability of a suitably constructed double multi-type branching process. To define it, let $\mathcal{T}_\mu^-(\kappa)$ and $\mathcal{T}_\mu^+(\kappa)$ denote two conditionally independent (given their common root) multi-type branching processes defined on the type space $\mathcal{S}$ whose roots are chosen according to μ and such that the number of offspring having types in a subset $A \subseteq \mathcal{S}$ that an individual of type $\mathbf{x} \in \mathcal{S}$ can have, is Poisson distributed with means

$$\int_A \kappa(\mathbf{y}, \mathbf{x})\mu(d\mathbf{y}) \quad \text{for } \mathcal{T}_\mu^-(\kappa) \qquad \text{and} \qquad \int_A \kappa(\mathbf{x}, \mathbf{y})\mu(d\mathbf{y}) \quad \text{for } \mathcal{T}_\mu^+(\kappa), \tag{3.3.1}$$

respectively. Next, let $\rho_-(\kappa; \mathbf{x})$ and $\rho_+(\kappa; \mathbf{x})$ denote the survival probabilities of $\mathcal{T}_\mu^-(\kappa; \mathbf{x})$ and $\mathcal{T}_\mu^+(\kappa; \mathbf{x})$, respectively, where $\mathcal{T}_\mu^-(\kappa; \mathbf{x})$ and $\mathcal{T}_\mu^+(\kappa; \mathbf{x})$ denote the trees whose root has type $\mathbf{x}$. We recall that a branching process is said to survive if its total population is infinite. We refer the reader to [79, 8] for more details on multi-type branching processes, including those with uncountable type spaces as the ones defined above.

In order to state our result for the phase transition in IRDs we first need to introduce the following definitions.

Definition 3.3.8 *A kernel κ defined on a separable metric space $\mathcal{S}$ with respect to a Borel probability measure μ is said to be* irreducible *if for any subset $A \subseteq \mathcal{S}$ satisfying $\kappa = 0$ a.e. on $A \times A^c$, we have either $\mu(A) = 0$ or $\mu(A^c) = 0$. We say that κ is* quasi-irreducible *if there is a μ-continuity set $\mathcal{S}' \subseteq \mathcal{S}$ with $\mu(\mathcal{S}') > 0$ such that the restriction of κ to $\mathcal{S}' \times \mathcal{S}'$ is irreducible, and $\kappa(\mathbf{x}, \mathbf{y}) = 0$ if $\mathbf{x} \notin \mathcal{S}'$ or $\mathbf{y} \notin \mathcal{S}'$.*

Definition 3.3.9 *A kernel κ on a separable metric space $\mathcal{S}$ with respect to a Borel probability measure μ is* regular finitary *if $\mathcal{S}$ has a finite partition into sets $\mathcal{J}_1, ..., \mathcal{J}_r$ such that κ is constant on each $\mathcal{J}_i \times \mathcal{J}_j$, and each $\mathcal{J}_i$ is a μ-continuity set, i.e., it is measurable and has $\mu(\partial \mathcal{J}_i) = 0$.*

To give the condition under which a giant strongly connected component exists we also need to define the operators induced by kernel κ, i.e.,

$$T_\kappa^+ f(\mathbf{x}) = \int_{\mathcal{S}} \kappa(\mathbf{x}, \mathbf{y}) f(\mathbf{y})\mu(d\mathbf{y}) \qquad \text{and} \qquad T_\kappa^- f(\mathbf{x}) = \int_{\mathcal{S}} \kappa(\mathbf{y}, \mathbf{x}) f(\mathbf{y})\mu(d\mathbf{y}).$$

Note that T_κ^+ and T_κ^- are integral linear operators on $(\mathcal{S}, \mu)$ equipped with the norm

$$\left\|T_\kappa^\pm\right\|_{op} = \sup\{\left\|T_\kappa^\pm f\right\|_2 : f \geq 0, \|f\|_2 \leq 1\} \leq \infty,$$

which makes them (potentially) unbounded operators in $L^2(\mathcal{S}, \mu)$. We also define their corresponding spectral radii $r(T_\kappa^+)$ and $r(T_\kappa^-)$, where the spectral radius of operator T in $L^2(\mathcal{S}, \mu)$ is defined as

$$r(T) = \sup\{|\lambda| : \lambda \in \sigma(T)\},$$

where $\sigma(T) = \{\lambda \in \mathbb{C} : T - \lambda I \text{ is not boundedly invertible}\}$ is the spectrum of T and I is the operator that maps f onto itself.[12]

The phase transition result for the largest strongly connected component is given below.

Theorem 3.3.10 *Suppose Assumption 3.3.1 is satisfied and κ is irreducible. Let $\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))$ denote the largest strongly connected component of $G_n(\kappa(1+\varphi_n))$. Then,*

$$\frac{|\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))|}{n} \xrightarrow{P} \rho(\kappa) \qquad n \to \infty,$$

where

$$\rho(\kappa) = \int_{\mathcal{S}} \rho_-(\kappa; \mathbf{x})\rho_+(\kappa; \mathbf{x})\mu(d\mathbf{x}).$$

Furthermore, if $\rho(\kappa) > 0$ then $r(T_\kappa^-) > 1$ and $r(T_\kappa^+) > 1$, and if there exists a regular finitary quasi-irreducible kernel $\tilde{\kappa}$ such that $\tilde{\kappa} \leq \kappa$ a.e. and $r(T_{\tilde{\kappa}}^-) > 1$ (equivalently, $r(T_{\tilde{\kappa}}^+) > 1$), then $\rho(\kappa) > 0$.

Moreover, when $\rho(\kappa) > 0$ we can characterize the "bow-tie" structure defined by the giant strongly connected component, $\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))$, the set of vertices that can reach it (its fan-in), and the set of vertices that can be reached from it (its fan-out). The following result makes this precise.

Theorem 3.3.11 *Suppose Assumption 3.3.1 is satisfied and κ is irreducible. For each vertex $v \in [n]$ define its in-component and out-component as:*

$$R^-(v) = \{i \in [n] : v \text{ is reachable from } i \text{ by a directed path } (i, v) \text{ in } G_n(\kappa(1+\varphi_n))\}$$
$$R^+(v) = \{i \in [n] : i \text{ is reachable from } v \text{ by a directed path } (v, i) \text{ in } G_n(\kappa(1+\varphi_n))\}.$$

Define $L_n^- = \{v \in [n] : |R^-(v)| \geq (\log n)/n\}$ and $L_n^+ = \{v \in [n] : |R^+(v)| \geq (\log n)/n\}$. Then, if $\rho(\kappa) > 0$,

$$\lim_{n\to\infty} P\left(\mathcal{C}_1(G_n(\kappa(1+\varphi_n))) = L_n^+ \cap L_n^-\right) = 1,$$

and

$$\frac{|L_n^+|}{n} \xrightarrow{P} \int_{\mathcal{S}} \rho_+(\kappa; \mathbf{x})\mu(d\mathbf{x}) \qquad \text{and} \qquad \frac{|L_n^-|}{n} \xrightarrow{P} \int_{\mathcal{S}} \rho_-(\kappa; \mathbf{x})\mu(d\mathbf{x})$$

as $n \to \infty$.

Remark 3.3.12 *We point out that we do not have a full if and only if condition for the strict positivity of $\rho(\kappa)$, since our operators T_κ^- and T_κ^+ may be unbounded, in which case the continuity of the spectral radius is not guaranteed. However, when κ satisfies*

$$\int_{\mathcal{S}}\int_{\mathcal{S}} \kappa(x, y)^2 \mu(d\mathbf{x})\mu(d\mathbf{y}) < \infty,$$

then the operators T_κ^- and T_κ^+ are compact (see Lemma 5.15 in [14]), and Theorem 2.1(a) in [36] gives the continuity of the spectral radius for a sequence of quasi-irreducible kernels $\kappa_m \nearrow \kappa$ as $m \to \infty$, ensuring the existence of $\tilde{\kappa}$ in Theorem 3.3.10. Interestingly, for the rank-1 case we can indeed provide a full characterization even when the operators T_κ^- and T_κ^+ are unbounded, as Proposition 3.3.13 shows.

We end the expository part of the paper with a compilation of all our results for the rank-1 case, which includes the first four models in Example 3.2.1.

Proposition 3.3.13 (IRDs with rank-1 kernel) *Suppose that Assumption 3.3.1 is satisfied with κ irreducible and of the form $\kappa(\mathbf{x}, \mathbf{y}) = \kappa_+(\mathbf{x})\kappa_-(\mathbf{y})$. Let $\mathbf{X}$ denote a random variable distributed according to μ. Then, the following properties hold:*

1. ***Number of arcs:*** *let $e(G_n(\kappa(1+\varphi_n)))$ denote the number of arcs in $G_n(\kappa(1+\varphi_n))$, then*

$$\frac{e(G_n(\kappa(1+\varphi_n)))}{n} \to E[\kappa_-(\mathbf{X})]E[\kappa_+(\mathbf{X})] \quad \text{in } L^1(P) \qquad \text{as } n \to \infty.$$

2. ***Distribution of vertex degrees:*** *let $(D^-_{n,\xi}, D^+_{n,\xi})$ denote the in-degree and out-degree of a randomly chosen vertex in $G_n(\kappa(1+\varphi_n))$. Set $\lambda_+(\mathbf{x}) = \kappa_+(\mathbf{x})E[\kappa_-(\mathbf{X})]$ and $\lambda_-(\mathbf{x}) = \kappa_-(\mathbf{x})E[\kappa_+(\mathbf{X})]$. Then,*

$$(D^-_{n,\xi}, D^+_{n,\xi}) \Rightarrow (Z^-, Z^+), \qquad E[D^\pm_{n,\xi}] \to E[Z^\pm],$$

as $n \to \infty$, where Z^- and Z^+ are conditionally independent (given $\mathbf{X}$) mixed Poisson random variables with mixing distributions $\lambda_-(\mathbf{X})$ and $\lambda_+(\mathbf{X})$.

3. ***Scale-free degrees:*** *suppose that $(\kappa_-(\mathbf{X}), \kappa_+(\mathbf{X}))$ is non-standard regularly varying with scaling functions $a(t) \in \mathcal{RV}(1/\alpha)$ and $b(t) \in \mathcal{RV}(1/\beta)$ and limiting measure $\tilde{\nu}(\cdot)$. Then, (Z^-, Z^+) is non-standard regularly varying with scaling functions $a(t)$ and $b(t)$ and limiting measure $\nu(\cdot)$ satisfying*

$$\nu((x,\infty] \times (y,\infty]) = \tilde{\nu}\left(\left(\frac{x}{E[\kappa_-(\mathbf{X})]}, \infty\right] \times \left(\frac{y}{E[\kappa_+(\mathbf{X})]}, \infty\right]\right).$$

4. ***Phase transition for the largest strongly connected component:*** *suppose κ is irreducible and let $\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))$ denote the largest strongly connected component of $G_n(\kappa(1+\varphi_n))$. Then,*

$$\frac{|\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))|}{n} \xrightarrow{P} \rho(\kappa), \qquad n \to \infty,$$

with $\rho(\kappa) > 0$ if and only if $E[\kappa_+(\mathbf{X})\kappa_-(\mathbf{X})] > 1$.

The remainder of the paper is devoted to the proofs of all the results mentioned above.

3.4 Proofs.

This section contains all the proofs of the theorems in Section 3.3. They are organized according to the same order in which their corresponding statements appear. Throughout this section we use the notation

$$q_{ij}^{(n)} = \frac{\kappa(\mathbf{X}_i, \mathbf{X}_j)}{n} \qquad 1 \le i, j \le n,$$

to denote the asymptotic limit of the arc probabilities in the graph, and to avoid having to explicitly exclude possible self-loops, we define $p_{ii}^{(n)} = 0$ for all $1 \le i \le n$. We also use $f(x) = O(g(x))$ as $x \to \infty$ to mean that $\limsup_{x\to\infty} |f(x)/g(x)| < \infty$.

3.4.1 Number of arcs.

The first result we prove corresponds to Proposition 3.3.3, which gives the asymptotic number of edges in $G_n(\kappa(1 + \varphi_n))$. Before we do so, we state and prove two preliminary technical lemmas that will be used several times throughout the paper.

Lemma 3.4.1 *Assume Assumption 3.3.1 holds and define for any* $0 < \epsilon < 1/2$ *the events*

$$B_{ij} = \left\{ (1-\epsilon) q_{ij}^{(n)} \le p_{ij}^{(n)} \le (1+\epsilon) q_{ij}^{(n)},\, q_{ij}^{(n)} \le \epsilon \right\}. \tag{3.4.1}$$

Then,

$$\lim_{n\to\infty} E\left[\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \left(p_{ij}^{(n)} + q_{ij}^{(n)} \right) 1(B_{ij}^c) \right] = 0.$$

Proof. We start by defining $A_{ij} = \{q_{ij}^{(n)} \le \epsilon\}$ and noting that the expression inside the expectation is bounded from above by

$$\frac{2}{n} \sum_{i=1}^n \sum_{j=1}^n q_{ij}^{(n)} 1\left(p_{ij}^{(n)} < (1-\epsilon) q_{ij}^{(n)}, A_{ij} \right) + \frac{2}{n} \sum_{i=1}^n \sum_{j=1}^n p_{ij}^{(n)} 1\left(p_{ij}^{(n)} > (1+\epsilon) q_{ij}^{(n)},\, A_{ij} \right) \tag{3.4.2}$$

$$+ \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n (1 + q_{ij}^{(n)}) 1(A_{ij}^c). \tag{3.4.3}$$

To show that (3.4.3) converges to zero, let $\mathbf{X}^{(n)} = \mathbf{X}_I$ and $\mathbf{Y}^{(n)} = \mathbf{Y}_J$ where I and J are mutually independent and uniformly distributed in $\{1, \cdots, n\}$, and independent of everything else. Note that

$$\begin{aligned} \frac{1}{n} E\left[\sum_{i=1}^n \sum_{j=1}^n (1 + q_{ij}^{(n)}) 1(A_{ij}^c) \right] &\le \frac{1}{n} E\left[\sum_{i=1}^n \sum_{j=1}^n (\epsilon^{-1} + 1) q_{ij}^{(n)} 1(A_{ij}^c) \right] \\ &= (\epsilon^{-1} + 1) E\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) 1(\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) > \epsilon n) \right]. \end{aligned}$$

Note that Assumption 3.3.1(a)-(b) imply that $\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) \Rightarrow \kappa(\mathbf{X}, \mathbf{Y})$ as $n \to \infty$, where $\mathbf{X}$ and $\mathbf{Y}$ are i.i.d. with distribution μ. Moreover, Assumption 3.3.1(d) gives $E[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)})] \to E[\kappa(\mathbf{X}, \mathbf{Y})]$ as $n \to \infty$. Hence, we can construct $(\{\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}\}_{n\geq 1}, \mathbf{X}, \mathbf{Y})$ on a common probability space such that $(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) \to (\mathbf{X}, \mathbf{Y})$ P-a.s. and $\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) \to \kappa(\mathbf{X}, \mathbf{Y})$ P-a.s. Fatou's lemma then gives

$$\begin{aligned}
&\limsup_{n\to\infty} E\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) 1(\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) > \epsilon n)\right] \\
&= \lim_{n\to\infty} E\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)})\right] - \liminf_{n\to\infty} E\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) 1(\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) \leq \epsilon n)\right] \\
&\leq E\left[\kappa(\mathbf{X}, \mathbf{Y})\right] - E\left[\kappa(\mathbf{X}, \mathbf{Y})\right] = 0.
\end{aligned}$$

To analyze the expectation of the first sum in (3.4.2), note that

$$\begin{aligned}
&\frac{1}{n} E\left[\sum_{i=1}^{n}\sum_{j=1}^{n} q_{ij}^{(n)} 1\left(p_{ij}^{(n)} < (1-\epsilon) q_{ij}^{(n)}, A_{ij}\right)\right] \\
&= \frac{1}{n} E\left[\sum_{i=1}^{n}\sum_{j=1}^{n} q_{ij}^{(n)} 1\left(q_{ij}^{(n)}(1+\varphi_n(\mathbf{X}_i, \mathbf{X}_j)) < (1-\epsilon) q_{ij}^{(n)} \leq \epsilon(1-\epsilon)\right)\right] \\
&\leq \frac{1}{n^2} E\left[\sum_{i=1}^{n}\sum_{j=1}^{n} \kappa(\mathbf{X}_i, \mathbf{X}_j) 1\left(\varphi_n(\mathbf{X}_i, \mathbf{X}_j) < -\epsilon\right)\right] \\
&= E\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)})\right] - E\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) 1(\varphi_n(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) \geq -\epsilon)\right]. \qquad (3.4.4)
\end{aligned}$$

Similarly, the expectation of the second sum in (3.4.2) can be bounded as follows

$$\begin{aligned}
&\frac{1}{n} E\left[\sum_{i=1}^{n}\sum_{j=1}^{n} p_{ij}^{(n)} 1\left(p_{ij}^{(n)} > (1+\epsilon) q_{ij}^{(n)},\ A_{ij}\right)\right] \\
&= \frac{1}{n} E\left[\sum_{i=1}^{n}\sum_{j=1}^{n} p_{ij}^{(n)} 1\left(q_{ij}^{(n)}(1+\varphi_n(\mathbf{X}_i, \mathbf{X}_j)) > (1+\epsilon) q_{ij}^{(n)},\ q_{ij}^{(n)} \leq \epsilon\right)\right] \\
&\leq \frac{1}{n} \sum_{i=1}^{n}\sum_{j=1}^{n} p_{ij}^{(n)} 1(\varphi_n(\mathbf{X}_i, \mathbf{X}_j) > \epsilon) \\
&= \frac{1}{n} E\left[\sum_{i=1}^{n}\sum_{j=1}^{n} p_{ij}^{(n)}\right] - E\left[\left(\left\{\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)})(1+\varphi_n(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}))\right\} \wedge n\right) 1(\varphi_n(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) \leq \epsilon)\right]. \\
&\qquad (3.4.5)
\end{aligned}$$

Using Fatou's lemma again and Assumption 3.3.1(c) (which implies that $\varphi_n(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) \xrightarrow{P} 0$ as $n \to \infty$), we have that

$$\liminf_{n\to\infty} E\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) 1(\varphi_n(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) \geq -\epsilon)\right] \geq E\left[\kappa(\mathbf{X}, \mathbf{Y})\right]$$

and

$$\liminf_{n\to\infty} E\left[\left(\left\{\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})(1+\varphi_n(\mathbf{X}^{(n)},\mathbf{Y}^{(n)}))\right\}\wedge n\right)1(\varphi_n(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})\leq\epsilon)\right]\geq E\left[\kappa(\mathbf{X},\mathbf{Y})\right].$$

It follows then from Assumption 3.3.1(d) that both (3.4.4) and (3.4.5) converge to zero. This completes the proof. ■

The next result establishes the convergence in probability of the expected number of edges in the graph.

Lemma 3.4.2 *Under Assumption 3.3.1 we have*

$$\frac{1}{n^2}\sum_{i=1}^n\sum_{j=1}^n \kappa(\mathbf{X}_i,\mathbf{X}_j)\to\iint_{\mathcal{S}^2}\kappa((\mathbf{x},\mathbf{y})\mu(d\mathbf{x})\mu(d\mathbf{y}) \quad \text{and} \quad \frac{1}{n}\sum_{i=1}^n\sum_{j\neq i} p_{ij}^{(n)}\to\iint_{\mathcal{S}^2}\kappa((\mathbf{x},\mathbf{y})\mu(d\mathbf{x})\mu(d\mathbf{y})$$

in $L^1(P)$ *as* $n\to\infty$.

Proof. As in the proof of Lemma 3.4.1, note that

$$\frac{1}{n^2}\sum_{i=1}^n\sum_{j=1}^n \kappa(\mathbf{X}_i,\mathbf{X}_j)=\mathbb{E}\left[\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})\right],$$

where $\mathbf{X}^{(n)}$ and $\mathbf{Y}^{(n)}$ are conditionally i.i.d. given $\mathscr{F}$ with distribution μ_n (constructed as in Lemma 3.4.1). Let $\mathbf{X}$ and $\mathbf{Y}$ be i.i.d. with distribution μ and note that

$$\iint_{\mathcal{S}^2}\kappa((\mathbf{x},\mathbf{y})\mu(d\mathbf{x})\mu(d\mathbf{y})=E[\kappa(\mathbf{X},\mathbf{Y})].$$

Next, note that for any fixed $M>0$ we have that $\kappa(\mathbf{x},\mathbf{y})\wedge M$ is bounded and continuous, so by Lemma A.2 in [14] we have that

$$\mathbb{E}[\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})\wedge M]\xrightarrow{P}E[\kappa(\mathbf{X},\mathbf{Y})\wedge M]$$

as $n\to\infty$. Next, fix $\epsilon>0$ and choose $M>0$ such that $E[(\kappa(\mathbf{X},\mathbf{Y})-M)^+]\leq\epsilon/2$. Then,

$$\begin{aligned}
&P\left(\left|\mathbb{E}\left[\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})\right]-E[\kappa(\mathbf{X},\mathbf{Y})]\right|>\epsilon\right)\\
&=P\left(\left|\mathbb{E}\left[\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})\wedge M+(\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})-M)^+\right]-E[\kappa(\mathbf{X},\mathbf{Y})\wedge M+(\kappa(\mathbf{X},\mathbf{Y})-M)^+]\right|>\epsilon\right)\\
&\leq P\left(\left|\mathbb{E}\left[\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})\wedge M\right]-E[\kappa(\mathbf{X},\mathbf{Y})\wedge M]\right|+\mathbb{E}\left[(\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})-M)^+\right]>\epsilon/2\right)\\
&\leq P\left(\left|\mathbb{E}\left[\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})\wedge M\right]-E[\kappa(\mathbf{X},\mathbf{Y})\wedge M]\right|>\epsilon/4\right)+P\left(\mathbb{E}\left[(\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})-M)^+\right]>\epsilon/4\right)\\
&\leq P\left(\left|\mathbb{E}\left[\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})\wedge M\right]-E[\kappa(\mathbf{X},\mathbf{Y})\wedge M]\right|>\epsilon/4\right)+\frac{4}{\epsilon}E\left[(\kappa(\mathbf{X}^{(n)},\mathbf{Y}^{(n)})-M)^+\right].
\end{aligned}$$

Furthermore, the same arguments used in the proof of Lemma 3.4.1 give that

$$\begin{aligned}\limsup_{n\to\infty} E\left[(\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) - M)^+\right] &= E\left[\kappa(\mathbf{X}, \mathbf{Y})\right] - \liminf_{n\to\infty} E\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) \wedge M\right] \\ &\leq E\left[(\kappa(\mathbf{X}, \mathbf{Y}) - M)^+\right].\end{aligned}$$

Therefore,

$$\limsup_{n\to\infty} P\left(\left|\mathbb{E}\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)})\right] - E[\kappa(\mathbf{X}, \mathbf{Y})]\right| > \epsilon\right) \leq \frac{4}{\epsilon} E\left[(\kappa(\mathbf{X}, \mathbf{Y}) - M)^+\right],$$

and taking $M \to \infty$ gives $\mathbb{E}\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)})\right] \xrightarrow{P} E[\kappa(\mathbf{X}, \mathbf{Y})]$ as $n \to \infty$. Since by Assumption 3.3.1(d) we have $E\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)})\right] \to E[\kappa(\mathbf{X}, \mathbf{Y})]$, then

$$\mathbb{E}\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)})\right] \to E[\kappa(\mathbf{X}, \mathbf{Y})] \quad \text{in } L^1(P) \qquad n \to \infty. \tag{3.4.6}$$

For the second result recall that $p_{ii}^{(n)} = 0$ and $q_{ij}^{(n)} = \kappa(\mathbf{X}_i, \mathbf{X}_j)/n$, so it suffices to show that

$$\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \left(p_{ij}^{(n)} - q_{ij}^{(n)}\right) \to 0 \quad \text{in } L^1(P) \qquad n \to \infty. \tag{3.4.7}$$

To see that this is the case fix $0 < \epsilon < 1/2$ and define B_{ij} according to Lemma 3.4.1. Next, note that by (3.4.6) and Lemma 3.4.1 we have

$$\begin{aligned}E\left[\left|\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \left(p_{ij}^{(n)} - q_{ij}^{(n)}\right)\right|\right] &\leq \frac{1}{n} E\left[\sum_{i=1}^n \sum_{j=1}^n \epsilon q_{ij}^{(n)}\right] + \frac{1}{n} E\left[\sum_{i=1}^n \sum_{j=1}^n \left(p_{ij}^{(n)} + q_{ij}^{(n)}\right) 1(B_{ij}^c)\right] \\ &= \epsilon E\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)})\right] + \frac{1}{n} E\left[\sum_{i=1}^n \sum_{j=1}^n \left(p_{ij}^{(n)} + q_{ij}^{(n)}\right) 1(B_{ij}^c)\right] \\ &\to \epsilon E[\kappa(\mathbf{X}, \mathbf{Y})]\end{aligned}$$

as $n \to \infty$. Taking $\epsilon \to 0$ establishes (3.4.7), which completes the proof. ■

We are now ready to prove Proposition 3.3.3.

Proof of Proposition 3.3.3. We start by defining W_n to be the average number of arcs in the graph $G_n(\kappa(1 + \varphi_n))$ given the types, that is, $W_n := \mathbb{E}[e(G_n(\kappa(1 + \varphi_n)))]/n = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n p_{ij}^{(n)}$. Note that by Lemma 3.4.2 we have that $W_n \to E\left[\kappa(\mathbf{X}, \mathbf{Y})\right] < \infty$ in $L^1(P)$ as $n \to \infty$, where $\mathbf{X}$ and $\mathbf{Y}$ are i.i.d. with common distribution μ. Therefore, it suffices to show that $e(G_n(\kappa(1 + \varphi_n)))/n - W_n \to 0$ in $L^1(P)$ as $n \to \infty$.

To do this, let Y_{ij} denote the indicator of whether arc (i, j) is present in $G_n(\kappa(1 + \varphi_n))$ and note that

$$e(G_n(\kappa(1 + \varphi_n))) = \sum_{i=1}^n \sum_{j\neq i} Y_{ij},$$

where the $\{Y_{ij}\}$ are Bernoulli random variables with means $\{p_{ij}^{(n)}\}$, conditionally independent given $\mathscr{F}$. It follows that

$$\mathrm{V}ar(e(G_n(\kappa(1+\varphi_n)))|\mathscr{F}) = \sum_{i=1}^{n}\sum_{j=1}^{n} \mathrm{V}ar(Y_{ij}|\mathscr{F}) \leq \sum_{i=1}^{n}\sum_{j=1}^{n} \mathbb{E}[Y_{ij}] = \sum_{i=1}^{n}\sum_{j=1}^{n} p_{ij}^{(n)} = nW_n.$$

Therefore,

$$\begin{aligned} E\left[(e(G_n(\kappa(1+\varphi_n)))/n - W_n)^2\right] &= E\left[\mathbb{E}\left[(e(G_n(\kappa(1+\varphi_n)))/n - W_n)^2\right]\right] \\ &= E\left[n^{-2}\mathrm{V}ar(e(G_n(\kappa(1+\varphi_n)))|\mathscr{F})\right] \\ &\leq n^{-2}E\left[nW_n\right] \xrightarrow{P} 0, \end{aligned}$$

as $n \to \infty$. Hence, $e(G_n(\kappa(1+\varphi_n)))/n - W_n \to 0$ in $L^2(P)$, which completes the proof. ■

3.4.2 Distribution of vertex degrees.

We now move on to the proof for Theorem 3.3.6. The proof of Theorem 3.3.4 is given in Section 3.4.3, since it can be obtained as a corollary to Theorem 3.4.6. We will show that (Z^-, Z^+) has a non-standard regularly varying distribution whenever their conditional means $(\lambda^-(\mathbf{X}), \lambda^+(\mathbf{X}))$ have a non-standard regularly varying distribution. Throughout the proof we use the notation $[a, b] = \{\mathbf{x} \in \mathbb{R}^2 : \mathbf{a} \leq \mathbf{x} \leq \mathbf{b}\}$ to denote the rectangles in $\mathbb{R}^2$.

Proof of Theorem 3.3.6. To simplify the notation, let $\mathbf{W} = (W^-, W^+) = (\lambda^-(\mathbf{X}), \lambda^+(\mathbf{X}))$, and recall that we need to show that $\tilde{\nu}_t(\cdot) = tP((Z^-/a(t) \in du,\, Z^+/b(t)) \in \cdot)$ converges vaguely to $\nu(\cdot)$ in $M_-([0,\infty]^2 \setminus \{\mathbf{0}\})$ as $t \to \infty$. Note that by Lemma 6.1 in [92], it suffices to show that $\tilde{\nu}_t([\mathbf{0}, \mathbf{x}]^c) \to \nu([\mathbf{0}, \mathbf{x}]^c)$ as $t \to \infty$ for any continuity point $\mathbf{x} \in [\mathbf{0}, \boldsymbol{\infty}) \setminus \{\mathbf{0}\}$ of $\nu([\mathbf{0}, \cdot]^c)$.

To start, fix $(p, q) \in [\mathbf{0}, \boldsymbol{\infty}) \setminus \{\mathbf{0}\}$ to be a continuity point of $\nu([\mathbf{0}, \cdot]^c)$ and note that

$$\begin{aligned} \tilde{\nu}_t((p,\infty] \times (q,\infty]) &= \int_p^\infty \int_q^\infty tP\left(\frac{Z^-}{a(t)} \in du,\, \frac{Z^+}{b(t)} \in dv\right) \\ &= tP\left(\frac{Z^-}{a(t)} > p,\, \frac{Z^+}{b(t)} > q\right) \\ &= tE\left[P\left(\frac{Z^-}{a(t)} > p,\, \frac{Z^+}{b(t)} > q \middle| \mathbf{W}\right)\right] \\ &= tE\left[P\left(Z^- > pa(t)\middle| \mathbf{W}\right) P\left(Z^+ > qb(t)\middle| \mathbf{W}\right)\right]. \end{aligned}$$

It follows that we need to show that

$$\lim_{t\to\infty} tE\left[P\left(Z^- > pa(t)\middle| \mathbf{W}\right) P\left(Z^+ > qb(t)\middle| \mathbf{W}\right)\right] = \nu((p,\infty] \times (q,\infty]).$$

To this end, define $e(t) = \sqrt{\gamma a(t) \log a(t)}$ and $d(t) = \sqrt{\eta b(t) \log b(t)}$ with $\gamma > 2q\beta$, $\eta > 2p\alpha$, and use them to define the events

$$A_t = \{W^- > pa(t) - e(t)\} \qquad \text{and} \qquad B_t = \{W^+ > qb(t) - d(t)\}.$$

Now note that

$$\begin{aligned} & tE\left[P\left(Z^- > pa(t)\middle|\,\mathbf{W}\right)P\left(Z^+ > qb(t)\middle|\,\mathbf{W}\right)\right] \\ &= tE\left[P\left(Z^- > pa(t)\middle|\,\mathbf{W}\right)P\left(Z^+ > qb(t)\middle|\,\mathbf{W}\right)1(A_t \cap B_t)\right] \qquad (3.4.8)\\ &\quad + tE\left[P\left(Z^- > pa(t)\middle|\,\mathbf{W}\right)P\left(Z^+ > qb(t)\middle|\,\mathbf{W}\right)1(A_t^c \cup B_t^c)\right]. \qquad (3.4.9) \end{aligned}$$

To see that (3.4.9) vanishes in the limit, use the bound $P(\mathrm{Po}i(\lambda) \geq p) \leq e^{-\lambda}(e\lambda/p)^p$ for $p > \lambda$, where $\mathrm{Poi}(\lambda)$ is Poisson random variable with mean λ, to obtain that

$$\begin{aligned} & tE\left[P\left(Z^- > pa(t)\middle|\,\mathbf{W}\right)P\left(Z^+ > qb(t)\middle|\,\mathbf{W}\right)1(A_t^c)\right] \\ &\leq tE\left[P\left(Z^- > pa(t)\middle|\,\mathbf{W}\right)1(A_t^c)\right] \\ &\leq tE\left[\exp\left\{-W^- + pa(t)\left(1 + \log(W^-) - \log(pa(t))\right)\right\}1(A_t^c)\right] \\ &\leq t\exp\{-(pa(t) - e(t)) + pa(t)(1 + \log(pa(t) - e(t)) - \log(pa(t)))\} \\ &= t\exp\left\{e(t) + pa(t)\log\left(1 - \frac{e(t)}{pa(t)}\right)\right\} \\ &= t\exp\left(-\frac{e(t)^2}{2pa(t)} + O\left(\frac{e(t)^3}{(pa(t))^2}\right)\right) = ta(t)^{-\frac{\gamma}{2p}}\left(1 + O\left(\frac{(\log a(t))^{3/2}}{a(t)^{1/2}}\right)\right), \end{aligned}$$

where in the third inequality we used the observation that $g(u) = -u + pa(t)\log u$ is concave with a unique maximizer at $u^* = pa(t)$. Similarly,

$$\begin{aligned} & tE\left[P\left(Z^- > pa(t)\middle|\,\mathbf{W}\right)P\left(Z^+ > qb(t)\middle|\,\mathbf{W}\right)1(B_t^c)\right] \\ &\leq tb(t)^{-\frac{\eta}{2q}}\left(1 + O\left(\frac{(\log b(t))^{3/2}}{b(t)^{1/2}}\right)\right). \end{aligned}$$

Our choice of γ, η guarantees that both terms converge to zero as $t \to \infty$, hence showing that (3.4.9) does so as well.

It remains to show that (3.4.8) converges to $\nu((p,\infty] \times (q,\infty])$ as $t \to \infty$. To do this, we first note that (3.4.8) is equal to

$$tP(A_t \cap B_t) - tE\left[\left(1 - P\left(Z^- > pa(t)\middle|\,\mathbf{W}\right)P\left(Z^+ > qb(t)\middle|\,\mathbf{W}\right)\right)1(A_t \cap B_t)\right],$$

where

$$\begin{aligned} & tE\left[\left(1 - P\left(Z^- > pa(t)\middle|\,\mathbf{W}\right)P\left(Z^+ > qb(t)\middle|\,\mathbf{W}\right)\right)1(A_t \cap B_t)\right] \\ &\leq tE\left[P\left(Z^- \leq pa(t)\middle|\,\mathbf{W}\right)1(A_t \cap B_t)\right] + tE\left[P\left(Z^+ \leq qb(t)\middle|\,\mathbf{W}\right)1(A_t \cap B_t)\right] \\ &\leq tE\left[P\left(Z^- \leq pa(t)\middle|\,\mathbf{W}\right)1(\tilde{A}_t \cap B_t)\right] + tE\left[P\left(Z^+ \leq qb(t)\middle|\,\mathbf{W}\right)1(A_t \cap \tilde{B}_t)\right] \\ &\quad + tP(\tilde{A}_t^c \cap A_t \cap B_t) + tP(A_t \cap B_t \cap \tilde{B}_t^c) \end{aligned}$$

with

$$\tilde{A}_t = \{W^- > pa(t) + e(t)\} \subseteq A_t \qquad \text{and} \qquad \tilde{B}_t = \{W^+ > qb(t) + d(t)\} \subseteq B_t.$$

Now note that the inequality $P(\text{Poi}(\lambda) \leq p) \leq e^{-\lambda}(e\lambda/p)^p$ for $0 \leq p < \lambda$ gives that

$$\begin{aligned}
& tE\left[P\left(Z^- \leq pa(t) \middle| \mathbf{W}\right) 1(\tilde{A}_t \cap B_t)\right] \\
& \leq tE\left[\exp\left\{-W^- + pa(t)\left(1 + \log(W^-) - \log(pa(t))\right)\right\} 1(\tilde{A}_t)\right] \\
& \leq t\exp\left\{-(pa(t) + e(t)) + pa(t)\left(1 + \log(pa(t) + e(t)) - \log(pa(t))\right)\right\} \\
& = t\exp\left\{-e(t) + pa(t)\log\left(1 + \frac{e(t)}{pa(t)}\right)\right\} \\
& = t\exp\left(-\frac{e(t)^2}{2pa(t)} + O\left(\frac{e(t)^3}{(pa(t))^2}\right)\right) = ta(t)^{-\frac{\gamma}{2p}}\left(1 + O\left(\frac{(\log a(t))^{3/2}}{a(t)^{1/2}}\right)\right),
\end{aligned}$$

where we used again the concavity of $g(u) = -u + pa(t)\log u$. Similarly,

$$tE\left[P\left(Z^+ \leq qb(t) \middle| \mathbf{W}\right) 1(A_t \cap \tilde{B}_t)\right] \leq tb(t)^{-\frac{\eta}{2q}}\left(1 + O\left(\frac{(\log b(t))^{3/2}}{b(t)^{1/2}}\right)\right),$$

and our choice of γ, η give again that

$$\lim_{t\to\infty}\left\{tE\left[P\left(Z^- \leq pa(t) \middle| \mathbf{W}\right) 1(\tilde{A}_t \cap B_t)\right] + tE\left[P\left(Z^+ \leq qb(t) \middle| \mathbf{W}\right) 1(A_t \cap \tilde{B}_t)\right]\right\} = 0. \tag{3.4.10}$$

Next, let $\nu_t(du, dv) = tP(W^-/a(t) \in du, W^+/b(t) \in dv)$ and note that for any $0 < \epsilon < p \wedge q$, we have that

$$\begin{aligned}
& \limsup_{t\to\infty}\left\{tP(\tilde{A}_t^c \cap A_t \cap B_t) + tP(A_t \cap B_t \cap \tilde{B}_t^c)\right\} \\
& = \limsup_{t\to\infty}\left\{\nu_t\left((p - e(t)/a(t), p + e(t)/a(t)] \times (q - d(t)/b(t), \infty]\right)\right. \\
& \qquad \left. + \nu_t\left((p - e(t)/a(t), \infty] \times (q - d(t)/b(t), q + d(t)/b(t)]\right)\right\} \\
& \leq \limsup_{t\to\infty}\left\{\nu_t\left((p-\epsilon, p+\epsilon] \times (q-\epsilon, \infty]\right) + \nu_t\left((p-\epsilon, \infty] \times (q-\epsilon, q+\epsilon]\right)\right\} \\
& = \nu\left((p-\epsilon, p+\epsilon] \times (q-\epsilon, \infty]\right) + \nu\left((p-\epsilon, \infty] \times (q-\epsilon, q+\epsilon]\right).
\end{aligned}$$

Moreover, since (p, q) is a continuity point of ν, then

$$\lim_{\epsilon\downarrow 0}\left\{\nu\left((p-\epsilon, p+\epsilon] \times (q-\epsilon, \infty]\right) + \nu\left((p-\epsilon, \infty] \times (q-\epsilon, q+\epsilon]\right)\right\} = 0.$$

It follows that

$$\lim_{t\to\infty}\left\{tP(\tilde{A}_t^c \cap A_t \cap B_t) + tP(A_t \cap B_t \cap \tilde{B}_t^c)\right\} = 0,$$

which combined with (3.4.10) gives that

$$\lim_{t\to\infty} tE\left[\left(1 - P\left(Z^- > pa(t)\middle|\, \mathbf{W}\right) P\left(Z^+ > qb(t)\middle|\, \mathbf{W}\right)\right) 1(A_t \cap B_t)\right] = 0.$$

Finally, the continuity of ν at (p, q) also yields that

$$\lim_{t\to\infty} tP(A_t \cap B_t) = \lim_{t\to\infty} \nu_t\left((p - e(t)/a(t), \infty] \times (q - d(t)/b(t), \infty]\right) = \nu\left((p, \infty] \times (q, \infty]\right).$$

■

3.4.3 Phase transition for the largest strongly connected component.

The last part of the paper considers the connectivity properties of the graph, in particular, the size of the largest strongly connected component. As mentioned in Section 3.3.3, our Theorem 3.3.10 provides the directed version of Theorem 3.1 in [14]. However, our proof approach differs from the one used in [14] in the order in which we construct the different couplings involved. Specifically, in [14] the authors first couple the graph $G_n(\kappa(1 + \varphi_n))$ with another graph $G_n(\kappa_m)$, where κ_m is a piecewise constant kernel taking at most a finite number of different values and such that $\kappa_m \nearrow \kappa$ as $m \to \infty$. Then, they provide a coupling between the exploration of the component of a randomly chosen vertex in $G_n(\kappa_m)$ and that of a multi-type branching process, $\mathcal{T}_\mu(\kappa_m)$, whose offspring distribution is determined by κ_m. The phase transition result is then obtained by relating the survival probability of $\mathcal{T}_\mu(\kappa_m)$ with the survival probability of its limiting tree $\mathcal{T}_\mu(\kappa)$. Our proof leverages on the work done in [12], which applies to a related graph $G_{n'}(\kappa_m)$, to establish a lower bound for the size of the largest strongly connected component. For the upper bound, we give a new direct coupling between the exploration of the in-component and out-component of a randomly chosen vertex in $G_n(\kappa(1 + \varphi_n))$ and a double tree $(\mathcal{T}_\mu^-(\kappa_m), \mathcal{T}_\mu^+(\kappa_m))$, where $\kappa_m \nearrow \kappa$ as $m \to \infty$. We then relate the survival probabilities of $(\mathcal{T}_\mu^-(\kappa_m), \mathcal{T}_\mu^+(\kappa_m))$ with those of their limiting trees $(\mathcal{T}_\mu^-(\kappa), \mathcal{T}_\mu^+(\kappa))$ as $m \to \infty$.

Interestingly, trying to adapt the approach used in [14] to the directed case leads to a phenomenon that does not occur when analyzing undirected graphs. Namely, if we consider two coupled undirected graphs $G_n(\kappa(1 + \varphi_n))$ and $G_n(\kappa'(1 + \varphi_n'))$ such that every edge in the first graph is also present in the second one but not the other way around (e.g., when $\kappa(\mathbf{x}, \mathbf{y})(1 + \varphi_n(\mathbf{x}, \mathbf{y})) \leq \kappa'(\mathbf{x}, \mathbf{y})(1 + \varphi_n'(\mathbf{x}, \mathbf{y}))$ for all $\mathbf{x}, \mathbf{y} \in \mathcal{S}$), then, the difference in the sizes of the components of a vertex present in both graphs can be bounded by the difference in their number of edges (see Lemma 9.4 in [14]). However, in the directed case, this is no longer true, as Figure 3.1 illustrates. In other words, the existence of a (giant) strongly connected component can be determined by a single arc. For this reason, a coupling of the graphs $G_n(\kappa(1 + \varphi_n))$ and $G_n(\kappa_m)$, such as the one used in [14], does not provide an upper bound for the size of the strongly connected component in the directed case. This may be a notable observation considering the folklore that exists around the equivalence of undirected and directed networks.

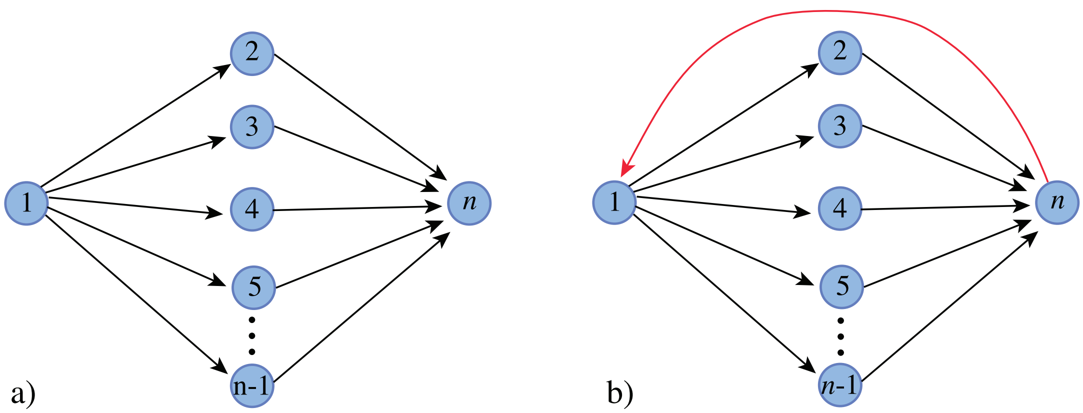

Figure 3.1: Directed graph with n vertices. a) There is no strongly connected component. b) The same graph with one additional arc; the largest strongly connected component is *giant* of size n.

With respect to how this section is organized, we have subdivided it into two subsections. In the first one we provide our coupling theorem between the exploration of the in-component and out-component of a randomly chosen vertex in $G_n(\kappa(1+\varphi_n))$ and the double tree $(\mathcal{T}_\mu^-(\kappa_m), \mathcal{T}_\mu^+(\kappa_m))$. The second subsection gives the proof of Theorem 3.3.10, which establishes the phase transition for the size of the largest strongly connected component.

3.4.3.1 Coupling with a double multi-type branching process.

Starting with a randomly chosen vertex in $G_n(\kappa(1+\varphi_n))$, say vertex i, we will perform a double exploration process that we will couple with a double multi-type branching process $\{\hat{\mathbf{Z}}_t^{(n)} : t \geq 0\}$ having "types" $\{1, \dots, n\}$. Note that these "types" are actually the *identities* of the vertices in $[n]$, so to avoid confusion with the actual *types* of each of the vertices, i.e., $\{\mathbf{X}_1, \dots, \mathbf{X}_n\}$, we will say that a vertex in the double tree has an *identity*, not a "type". The double tree is started at $\hat{\mathbf{Z}}_0^{(n)} = (\hat{Z}_{1,0}, \hat{Z}_{2,0}, \dots, \hat{Z}_{n,0})$, and is such that for $t \geq 1$, $\hat{\mathbf{Z}}_t^{(n)} = (\hat{Z}_{1,t}^-, \hat{Z}_{2,t}^-, \dots, \hat{Z}_{n,t}^-, \hat{Z}_{1,t}^+, \hat{Z}_{2,t}^+, \dots, \hat{Z}_{n,t}^+) \in \mathbb{N}^{2n}$, where $\hat{Z}_{j,t}^-$ denotes the number of individuals of *identity* j in the tth inbound generation of the double tree and $\hat{Z}_{j,t}^+$ denotes the number of individuals of *identity* j in the tth outbound generation of the double tree. Moreover, the number of offspring that each node in the double tree has is independent of all other nodes in the double tree, conditionally on the identity of the node. The initial vector $\hat{\mathbf{Z}}_0^{(n)}$ is set to equal $\mathbf{e}_i$, where $\mathbf{e}_i$ is the unit vector that has a one in position i and zeros elsewhere; note also that it does not have a $+/-$ superscript since it is at the center of the double tree.

In order to define the offspring distribution of nodes in the double tree, we fix a regular finitary kernel κ_m on $\mathcal{S} \times \mathcal{S}$ (see Definition 3.3.9) satisfying

$$0 \leq \kappa_m(\mathbf{x}, \mathbf{y}) \leq \kappa(\mathbf{x}, \mathbf{y}) \qquad \text{for all } \mathbf{x}, \mathbf{y} \in \mathcal{S},$$

and such that

$$\kappa_m(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{M_m} \sum_{j=1}^{M_m} c_{ij}^{(m)} 1(\mathbf{x} \in \mathcal{J}_i^{(m)}, \mathbf{y} \in \mathcal{J}_j^{(m)}),$$

for some partition $\{\mathcal{J}_i^{(m)} : 1 \leq i \leq M_m\}$ of $\mathcal{S}$ and some nonnegative constants $\{c_{ij}^{(m)} : 1 \leq i, j \leq M_m\}$, $M_m < \infty$. Now let the number of offspring of *identity* j that a node of *identity* i in the inbound tree, respectively outbound tree, has, be Poisson distributed with mean $r_{ji}^{(m,n)}$, resp. $\tilde{r}_{ij}^{(m,n)}$, where:

$$r_{ji}^{(m,n)} = \frac{\kappa_m(\mathbf{X}_j, \mathbf{X}_i)\mu(\mathcal{J}_{\theta(j)}^{(m)})}{n\mu_n(\mathcal{J}_{\theta(j)}^{(m)})} \qquad \text{and} \qquad \tilde{r}_{ij}^{(m,n)} = \frac{\kappa_m(\mathbf{X}_i, \mathbf{X}_j)\mu(\mathcal{J}_{\theta(j)}^{(m)})}{n\mu_n(\mathcal{J}_{\theta(j)}^{(m)})},$$

and $\theta(i) = j$ if and only if $\mathbf{X}_i \in \mathcal{J}_j^{(m)}$. We denote $\mathcal{T}_\mu^-(\kappa_m; \mathbf{X}_i)$ and $\mathcal{T}_\mu^+(\kappa_m; \mathbf{X}_i)$ the inbound and outbound trees, respectively, whose root is vertex i. Note that the trees $\mathcal{T}_\mu^-(\kappa_m; \mathbf{X}_i)$ and $\mathcal{T}_\mu^+(\kappa_m; \mathbf{X}_i)$ are conditionally independent (given $\mathscr{F}$) by construction.

Note: We point out that in the double tree *identities* can appear multiple times, unlike in the graph where they appear only once. In either case, *identities* take values in the set $[n] = \{1, 2, \ldots, n\}$.

Remark 3.4.3 *An important observation that will be used later is that the double tree* $\hat{\mathbf{Z}}_t^{(n)} = (\hat{Z}_{1,t}^-, \ldots, \hat{Z}_{n,t}^-, \hat{Z}_{1,t}^+, \ldots, \hat{Z}_{n,t}^+) \in \mathbb{N}^{2n}$ *defined above, conditional on* $\hat{Z}_0 = i \in [n]$, *has the same law as the double tree* $\tilde{\mathbf{Z}}_t^{(m)} = (\tilde{Z}_{1,t}^-, \ldots, \tilde{Z}_{M_m,t}^-, \tilde{Z}_{1,t}^+, \ldots, \tilde{Z}_{M_m,t}^+) \in \mathbb{N}^{2M_m}$, *whose offspring distributions are Poisson with means*

$$m_{ij}^- := c_{ji}^{(m)} \mu(\mathcal{J}_j^{(m)}) \qquad \textit{and} \qquad m_{ij}^+ := c_{ij}^{(m)} \mu(\mathcal{J}_j^{(m)}), \qquad 1 \leq i, j \leq M_m.$$

Moreover, the latter is the same as $(\mathcal{T}_\mu^-(\kappa_m; \mathbf{x}), \mathcal{T}_\mu^+(\kappa_m; \mathbf{x}))$ *for any* $\mathbf{x} \in \mathcal{J}_i^{(m)}$.

Recall that $Y_{ij} = 1(\text{arc } (i,j) \text{ is present in } G_n(\kappa(1 + \varphi_n)))$ is a Bernoulli random variable with success probability

$$p_{ij}^{(n)} = \frac{\kappa(\mathbf{X}_i, \mathbf{X}_j)(1 + \varphi_n(\mathbf{X}_i, \mathbf{X}_j))}{n} \wedge 1, \quad 1 \leq i \neq j \leq n, \qquad p_{ii}^{(n)} = 0.$$

We will couple Y_{ij} with a Poisson random variable Z_{ij} having mean $r_{ij}^{(m,n)}$ on the inbound side, and with a Poisson random variable $\tilde{Z}_{ij}$ having mean $\tilde{r}_{ij}^{(m,n)}$ on the outbound side, using a sequence $\{U_{ij} : 1 \leq i, j \leq n\}$ of i.i.d. Uniform$(0,1)$ random variables.

The exploration of the graph and the construction of the double tree are done by choosing a vertex uniformly at random among those which have not been explored. Starting with vertex i, we fix the number of vertices to explore in the in-component of i, say k_{in}, and the number of vertices to explore in the out-component of i, say k_{out}. A *step* in the exploration of the in-component (out-component) corresponds to identifying the inbound (outbound) neighbors of the vertex being explored. The exploration of the in-component continues until we have explored k_{in} vertices or until there are no more vertices to reveal, after which we proceed to explore the out-component for k_{out} steps or until there are no more vertices to reveal. Moreover, we allow k_{in} and k_{out} to be stopping times with respect to the history of the exploration process.

Vertices in the graph can have one of two labels: {inactive, active}, or they may be unlabelled. Active vertices are those that have been identified to be in the in-component, respectively out-component, of vertex i but whose inbound, respectively outbound, neighbors have not been revealed. Inactive vertices are all other vertices that have been revealed through the exploration process but that are not active; again, there is an inbound inactive set and an outbound inactive set. Inactive vertices on the inbound side have revealed all its inbound neighbors, but not necessarily all their outbound ones; symmetrically, inactive nodes on the outbound side have revealed all their outbound neighbors but not necessarily all their inbound ones.

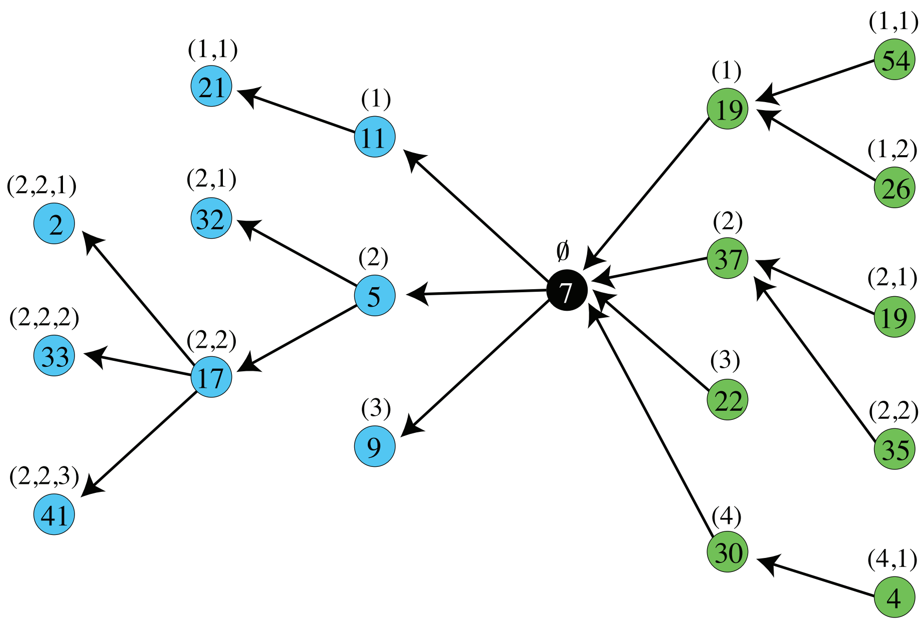

Figure 3.2: Exploration of the graph and coupled tree. We explore vertex 7 in the graph, which means the root of the double tree has *identity* $T_{\emptyset} = 7$. Node *identities* in the double tree are depicted inside the circles, whereas tree *labels* are right on top, e.g., node $(2, 2, 3)$ on the outbound tree has *identity* $T_{(2,2,3)} = 41$, whereas node (3) on the inbound tree has *identity* $T_{(3)} = 22$.

In the double tree we will say that a node is "active" if we have not yet sampled its offspring, and "inactive" if we have.

Notation: For $r = 0, 1, 2, \ldots$, and assuming the chosen vertex is i, let

A_r^- (A_r^+) = set of inbound (outbound) "active" vertices after having explored the first r vertices in the in-component (out-component) of vertex i.

I_r^- (I_r^+) = set of inbound (outbound) "inactive" vertices after having explored the first r vertices in the in-component (out-component) of vertex i.

T_r^- (T_r^+) = *identity* of the vertex being explored in step r, $r \geq 1$, of the exploration of the in-component (out-component) of vertex i.

$\hat{A}_r^-$ ($\hat{A}_r^+$) = set of "active" nodes in $\mathcal{T}_\mu^-(\kappa_m; \mathbf{X}_i)$ ($\mathcal{T}_n^+(\kappa_m; \mathbf{X}_i)$) after having sampled the offspring of the first r nodes in $\mathcal{T}_\mu^-(\kappa_m; \mathbf{X}_i)$ ($\mathcal{T}_\mu^+(\kappa_m; \mathbf{X}_i)$).

$\hat{I}_r^-$ ($\hat{I}_r^+$) = set of *identities* belonging to "inactive" nodes in $\mathcal{T}_\mu^-(\kappa_m; \mathbf{X}_i)$ ($\mathcal{T}_\mu^+(\kappa_m; \mathbf{X}_i)$) after having sampled the offspring of the first r nodes in $\mathcal{T}_\mu^-(\kappa_m; \mathbf{X}_i)$ ($\mathcal{T}_\mu^+(\kappa_m; \mathbf{X}_i)$).

$\hat{T}_r^-$ ($\hat{T}_r^+$) = *identity* of the node in $\mathcal{T}_\mu^-(\kappa_m; \mathbf{X}_i)$ ($\mathcal{T}_\mu^+(\kappa_m; \mathbf{X}_i)$) whose offspring are being sampled in step r; $r \geq 1$.

Exploration of the components of vertex i in the graph:
Fix k_{in} and k_{out}.

1) For the exploration of the in-component:

   Step 0: Label vertex i as "active" on the inbound side and set $A_0^- = \{i\}$, $I_0^- = \varnothing$.

   Step r, $1 \leq r \leq k_{in}$:

   Choose, uniformly at random, a vertex in A_{r-1}^-; let $T_r^- = i$ denote its *identity*.

   a) For $j = 1, 2, \ldots, n$, $j \neq i$:
      i. Realize $Y_{ji} = 1(U_{ji} > 1 - p_{ji}^{(n)})$. If $Y_{ji} = 0$ go to 1(a).
      ii. If $Y_{ji} = 1$ and vertex $j \in I_{r-1}^- \cup A_{r-1}^-$, do nothing. Go to 1(a).
      iii. If $Y_{ji} = 1$ and vertex j had no label, label it "active" on the inbound side. Go to 1(a).

   b) Once all the new inbound neighbors of vertex i have been identified and labeled "active", label vertex i as "inactive" on the inbound side.

   c) Define the sets $A_r^- = A_{r-1}^- \cup \{\text{new "active" vertices created in 1(a)(iii)}\} \setminus \{i\}$ and $I_r^- = I_{r-1}^- \cup \{i\}$. This completes Step r on the inbound side.

2) For the exploration of the out-component:

   Step 0: Label vertex i as "active" on the outbound side and set $A_0^+ = \{i\}$, $I_0^+ = \varnothing$.

   Step r, $1 \leq r \leq k_{out}$:

   Choose, uniformly at random, a vertex in A_{r-1}^+; let $T_r^+ = i$ denote its *identity*.

a) For $j = 1, 2, \ldots, n$, $j \neq i$, $j \notin I^-_{k_{in}} \cup A^-_{k_{in}}$:

   i. Realize $Y_{ij} = 1(U_{ij} > 1 - p^{(n)}_{ij})$. If $Y_{ij} = 0$ go to 2(a).

   ii. If $Y_{ij} = 1$ and vertex $j \in I^+_{r-1} \cup A^+_{r-1}$, do nothing. Go to 2(a).

   iii. If $Y_{ij} = 1$ and vertex j had no label, label it "active" on the outbound side. Go to 2(a).

b) Once all the new outbound neighbors of vertex i have been identified and labeled "active", label vertex i as "inactive" on the outbound side.

c) Define the sets $A^+_r = A^+_{r-1} \cup \{\text{new "active" vertices created in 2(a)(iii)}\} \setminus \{i\}$ and $I^+_r = I^+_{r-1} \cup \{i\}$. This completes Step r on the outbound side.

Note that by setting $k_{in} = \inf\{r \geq 1 : A^-_r = \varnothing\}$ and $k_{out} = \inf\{r \geq 1 : A^+_r = \varnothing\}$ we can fully explore the in-component and out-component of vertex i. We now explain how the coupled double tree is constructed.

Coupled construction of the double multi-type branching process:

Let $g^{-1}(u)$ denote the pseudo inverse of function g, i.e., $g^{-1}(u) = \inf\{x : u \leq g(x)\}$. Let G_{ji} and $\tilde{G}_{ij}$ be the distribution functions of Poisson random variables having means $r^{(m,n)}_{ji}$ and $\tilde{r}^{(m,n)}_{ij}$, respectively. On the double tree we use the index notation $\mathbf{i} = (i_1, \ldots, i_r)$ to identify nodes in the rth generation (inbound/outbound) of the double tree. Let $T_{\mathbf{i}}$ denote the *identity* of node $\mathbf{i}$; see Figure 3.2.

1) Construction of the inbound tree:

   Step 0: Set $\hat{\mathbf{Z}}^{(n)}_0 = \mathbf{e}_i$. Let $\hat{A}^-_0 = \{\emptyset\}$, $T_\emptyset = i$, $\hat{I}^-_0 = \varnothing$.

   Step r, $1 \leq r \leq k_{out}$:

   Choose a node in $\mathbf{i} \in \hat{A}^-_{r-1}$, uniformly at random; set $\hat{T}^-_r = T_{\mathbf{i}}$.

   I. If this is the first time *identity* $T_{\mathbf{i}}$ appears in the inbound tree, do as follows:

   a) For $j = 1, 2, \ldots, n$, $j \notin \{T_{\mathbf{i}}\}$:

      i. Realize $Z_{j,T_{\mathbf{i}}} = G^{-1}_{j,T_{\mathbf{i}}}(U_{j,T_{\mathbf{i}}})$. If $Z_{j,T_{\mathbf{i}}} = 0$ go to 1(I)(a).

      ii. If $Z_{j,T_{\mathbf{i}}} \geq 1$ label each of the newly created nodes as "active" on the inbound side. Go to 1(I)(a).

   b) For $j = T_{\mathbf{i}}$:

      i. Sample $Z^*_{j,T_{\mathbf{i}}}$ to be a Poisson random variable with mean $r^{(m,n)}_{j,T_{\mathbf{i}}}$, independently of everything else. If $Z^*_{j,T_{\mathbf{i}}} = 0$ go to 1(I)(c).

      ii. If $Z^*_{j,T_{\mathbf{i}}} \geq 1$ label each of the newly created nodes as "active" on the inbound side. Go to 1(I)(c).

   c) Once all the inbound offspring of node $\mathbf{i}$ have been identified, label *identity* $T_{\mathbf{i}}$ as "inactive" on the inbound side.

d) Define the sets $\hat{A}_r^- = \hat{A}_{r-1}^- \cup \{$new "active" nodes created in 1(I)(a)(ii) and 1(I)(b)(ii)$\} \setminus \{\mathbf{i}\}$ and $\hat{I}_r^- = \hat{I}_{r-1}^- \cup \{T_{\mathbf{i}}\}$. This completes Step r on the inbound side.

II. Else:

a) For $j = 1, 2, \ldots, n$:

i. Sample $Z_{j,T_{\mathbf{i}}}^*$ to be a Poisson random variable with mean $r_{j,T_{\mathbf{i}}}^{(m,n)}$, independently of everything else. If $Z_{j,T_{\mathbf{i}}}^* = 0$ go to 1(II)(a).

ii. If $Z_{j,T_{\mathbf{i}}}^* \geq 1$ label each of the newly created nodes as "active" on the inbound side. Go to 1(II)(a).

b) Once all the inbound offspring of node $\mathbf{i}$ have been identified, label *identity* $T_{\mathbf{i}}$ as "inactive" on the inbound side.

c) Define the sets $\hat{A}_r^- = \hat{A}_{r-1}^- \cup \{$new "active" nodes created in 1(II)(a)(ii)$\} \setminus \{\mathbf{i}\}$ and $\hat{I}_r^- = \hat{I}_{r-1}^- \cup \{T_{\mathbf{i}}\}$. This completes Step r on the inbound side.

2) Construction of the outbound tree:

Step 0: Set $\hat{A}_0^+ = \{\emptyset\}$, $T_\emptyset = i, \hat{I}_0^+ = \varnothing$.

Choose a node $\mathbf{i} \in \hat{A}_{r-1}^+$, uniformly at random; set $\hat{T}_r^+ = T_{\mathbf{i}}$.

I. If this is the first time *identity* $T_{\mathbf{i}}$ appears in the outbound tree, do as follows:

a) For $j = 1, 2, \ldots, n$, $j \notin \{T_{\mathbf{i}}\} \cup \{T_{\mathbf{j}} : T_{\mathbf{j}} \in \hat{I}_{k_{in}}^-$ or $\mathbf{j} \in \hat{A}_{k_{in}}^-\}$:

i. Realize $\tilde{Z}_{T_{\mathbf{i}},j} = \tilde{G}_{T_{\mathbf{i}},j}^{-1}(U_{T_{\mathbf{i}},j})$. If $\tilde{Z}_{T_{\mathbf{i}},j} = 0$ go to 2(I)(a).

ii. If $\tilde{Z}_{T_{\mathbf{i}},j} \geq 1$ label each of the newly created nodes as "active" on the outbound side. Go to 2(I)(a).

b) For $j \in \{T_{\mathbf{i}}\} \cup \{T_{\mathbf{j}} : T_{\mathbf{j}} \in \hat{I}_{k_{in}}^-$ or $\mathbf{j} \in \hat{A}_{k_{in}}^-\}$:

i. Sample $\tilde{Z}_{T_{\mathbf{i}},j}^*$ to be a Poisson random variable with mean $\tilde{r}_{T_{\mathbf{i}},j}^{(m,n)}$, independently of everything else. If $\tilde{Z}_{T_{\mathbf{i}},j}^* = 0$ go to 2(I)(b).

ii. If $\tilde{Z}_{T_{\mathbf{i}},j}^* \geq 1$ label each of the newly created nodes as "active" on the outbound side. Go to 2(I)(b).

c) Once all the outbound offspring of node $\mathbf{i}$ have been identified, label *identity* $T_{\mathbf{i}}$ as "inactive" on the outbound side.

d) Define the sets $\hat{A}_r^+ = \hat{A}_{r-1}^+ \cup \{$new "active" nodes created in 2(I)(a)(ii) and 2(I)(b)(ii)$\} \setminus \{\mathbf{i}\}$ and $\hat{I}_r^- = \hat{I}_{r-1}^- \cup \{T_{\mathbf{i}}\}$. This completes Step r on the outbound side.

II. Else:

a) For $j = 1, 2, \ldots, n$:

i. Sample $\tilde{Z}_{T_{\mathbf{i}},j}^*$ to be a Poisson random variable with mean $\tilde{r}_{T_{\mathbf{i}},j}^{(m,n)}$, independently of everything else. If $\tilde{Z}_{T_{\mathbf{i}},j}^* = 0$ go to 2(II)(a).

ii. If $\tilde{Z}^*_{j,T_{\mathbf{i}}} \geq 1$ label each of the newly created nodes as "active" on the outbound side. Go to 2(II)(a).

b) Once all the outbound offspring of node $\mathbf{i}$ have been identified, label *identity* $T_{\mathbf{i}}$ as "inactive" on the outbound side.

c) Define the sets $\hat{A}^+_r = \hat{A}^+_{r-1} \cup \{$new "active" nodes created in 2(II)(a)(ii)$\} \setminus \{\mathbf{i}\}$ and $\hat{I}^-_r = \hat{I}^-_{r-1} \cup \{T_{\mathbf{i}}\}$. This completes Step r on the outbound side.

Note: As long as the active sets in the graph and the double tree are the same, the chosen nodes in steps (1)(I) and (2)(I) are the same as the vertices chosen in steps (1) and (2) of the graph exploration process.

Definition 3.4.4 *We say that the coupling of the graph and the double multi-type branching process holds up to Step r on the inbound side if*

$$A^-_t = \{T_{\mathbf{j}} : \mathbf{j} \in \hat{A}^-_t\} \qquad \text{and} \qquad |A^-_t| = |\hat{A}^-_t| \quad \text{for all } 0 \leq t \leq r,$$

and up to Step r on the outbound side if

$$A^+_t = \{T_{\mathbf{j}} : \mathbf{j} \in \hat{A}^+_t\} \qquad \text{and} \qquad |A^+_t| = |\hat{A}^+_t| \quad \text{for all } 0 \leq t \leq r.$$

Let $T_\emptyset$ denote the identity *of the vertex whose in and out-components we want to explore. Define the stopping time τ^- to be the step in the graph exploration process of vertex $T_\emptyset$ during which the coupling breaks on the inbound side and τ^+ to be the step during which it breaks on the outbound side.*

Remark 3.4.5 *Note that $\tau^- = r$ if and only if either:*

a. *For any $j = 1, 2, \ldots, n$, $j \notin \{T^-_r\} \cup A^-_{r-1} \cup I^-_{r-1}$, we have $Z_{j,T^-_r} \neq Y_{j,T^-_r}$ in step (1)(I)(a)(i),*

b. *For any $j \in A^-_{r-1} \cup I^-_{r-1}$ we have $Z_{j,T^-_r} \geq 1$ in step (1)(I)(a)(i),*

c. *$Z^*_{T^-_r,T^-_r} \geq 1$ in step (1)(I)(b)(i),*

and $\tau^+ = r$ if and only if either:

d. *For any $j = 1, 2, \ldots, n$, $j \notin \{T^+_r\} \cup I^-_{k_{in}} \cup A^-_{k_{in}} \cup A^+_{r-1} \cup I^+_{r-1}$, we have $\tilde{Z}_{T^+_r,j} \neq Y_{T^+_r,j}$ in step (2)(I)(a)(i),*

e. *For any $j \in A^+_{r-1} \cup I^+_{r-1}$ we have $\tilde{Z}_{T^+_r,j} \geq 1$ in step (2)(I)(a)(i),*

f. *For any $j \in \{T^+_r\} \cup I^-_{k_{in}} \cup A^-_{k_{in}}$, we have $\tilde{Z}^*_{T^+_r,j} \geq 1$ in step (2)(I)(b)(i).*

We are now ready to state our main coupling result, which provides an explicit upper bound for the probability that the coupling breaks before we can determine whether both the in-component and the out-component of the vertex being explored have at least k vertices each or are fully explored.

Throughout the remainder of the paper, we use the notation $\mathbb{P}_i(\cdot) = \mathbb{E}[1(\cdot)|A_0 = \{i\}]$ and $\mathbb{E}_i[\cdot] = \mathbb{E}[\cdot|A_0 = \{i\}]$; also, $\|\mathbf{x}\|_1 = \sum_i |x_i|$ for any $\mathbf{x} \in \mathbb{R}^n$. Similarly to the definition of $\lambda_-(\mathbf{x})$ and $\lambda_+(\mathbf{x})$, define

$$\lambda_-^{(m)}(\mathbf{x}) = \int_{\mathcal{S}} \kappa_m(\mathbf{y}, \mathbf{x})\mu(d\mathbf{y}) \qquad \text{and} \qquad \lambda_+^{(m)}(\mathbf{x}) = \int_{\mathcal{S}} \kappa_m(\mathbf{x}, \mathbf{y})\mu(d\mathbf{y}),$$

$$\lambda_{m,n}^-(\mathbf{x}) = \int_{\mathcal{S}} \kappa_m(\mathbf{y}, \mathbf{x})\mu_n(d\mathbf{y}) \qquad \text{and} \qquad \lambda_{m,n}^+(\mathbf{x}) = \int_{\mathcal{S}} \kappa_m(\mathbf{x}, \mathbf{y})\mu_n(d\mathbf{y}),$$

and

$$\lambda_n^-(\mathbf{x}) = \int_{\mathcal{S}} \kappa(\mathbf{y}, \mathbf{x})\mu_n(d\mathbf{y}) \qquad \text{and} \qquad \lambda_n^+(\mathbf{x}) = \int_{\mathcal{S}} \kappa(\mathbf{x}, \mathbf{y})\mu_n(d\mathbf{y}).$$

Theorem 3.4.6 *Consider the exploration process described above along with its coupled double tree construction. Define for any fixed $k \in \mathbb{N}_-$ the stopping times $\sigma_k^- = \inf\{t \geq 1 : |A_t^-| + |I_t^-| \geq k$ or $A_t^- = \varnothing\}$ and $\sigma_k^+ = \{t \geq 1 : |A_t^+| + |I_t^+| \geq k$ or $A_t^+ = \varnothing\}$. For any $0 < \epsilon < 1/2$ and any $n, m \in \mathbb{N}_-$,*

$$\frac{1}{n}\sum_{i=1}^{n} \mathbb{P}_i\left(\{\tau^- \leq \sigma_k^-\} \cup \{\tau^+ \leq \sigma_k^+\}\right) \leq H(n, m, k, \epsilon),$$

where

$$H(n, m, k, \epsilon) = 1(\Omega_{m,n}^c) + 4\epsilon k^2 + 2\epsilon k^2 \left(1 + \sup_{\mathbf{x} \in \mathcal{S}} \lambda_-^{(m)}(\mathbf{x})\right) + 1(\Omega_{m,n}) \sum_{r=1}^{k} \sum_{s=0}^{r-1} \binom{r-1}{s} 2^{r-1-s}$$
$$\cdot \left\{ \int_{\mathcal{S}} (\Gamma_-^{(m,n)})^s g_{m,n,\epsilon}^-(\mathbf{x})\mu_n(d\mathbf{x}) + \int_{\mathcal{S}} (\Gamma_+^{(m,n)})^s g_{m,n,\epsilon}^+(\mathbf{x})\mu_n(d\mathbf{x}) \right\},$$

the linear integral operators $\Gamma_-^{(m,n)}$ and $\Gamma_+^{(m,n)}$ are defined in Lemma 3.4.9, the functions $g_{m,n,\epsilon}^-$ and $g_{m,n,\epsilon}^+$ are defined according to

$$g_{m,n,\epsilon}^-(\mathbf{X}_i) = \min\left\{1,\ (1 + 5\epsilon)\lambda_n^-(\mathbf{X}_i) - \lambda_{m,n}^-(\mathbf{X}_i) + (1 + \epsilon)\sum_{j=1}^{n} (p_{ji}^{(n)} + q_{ji}^{(n)})1(B_{ji}^c)\right\},$$

$$g_{m,n,\epsilon}^+(\mathbf{X}_i) = \min\left\{1,\ (1 + 5\epsilon)\lambda_n^+(\mathbf{X}_i) - \lambda_{m,n}^+(\mathbf{X}_i) + (1 + \epsilon)\sum_{j=1}^{n} (p_{ij}^{(n)} + q_{ij}^{(n)})1(B_{ij}^c)\right\},$$

$$\Omega_{m,n} = \bigcap_{t=1}^{M_m} \left\{ \left| \frac{\mu(\mathcal{J}_t^{(m)})}{\mu_n(\mathcal{J}_t^{(m)})} - 1 \right| 1(\mu_n(\mathcal{J}_t^{(m)}) > 0) < \epsilon \right\},$$

$$B_{ij} = \left\{ (1 - \epsilon)q_{ij}^{(n)} \leq p_{ij}^{(n)} \leq (1 + \epsilon)q_{ij}^{(n)},\ q_{ij}^{(n)} \leq \epsilon \right\}.$$

Moreover, $H(n,m,k,\epsilon) \xrightarrow{P} \hat{H}(m,k,\epsilon)$ *(defined in Lemma 3.4.10) as* $n \to \infty$ *with*

$$\lim_{m \nearrow \infty} \lim_{\epsilon \downarrow 0} \hat{H}(m,k,\epsilon) = 0$$

for any fixed $k \geq 1$*.*

Before proving the theorem, we will state and prove several preliminary results. The first one below gives an upper bound for the number of offspring sampled in each side of the double-tree $(\mathcal{T}_\mu^-(\kappa_m), \mathcal{T}_\mu^+(\kappa_m))$ up to step $\hat{\sigma}_k^-$ and step $\hat{\sigma}_k^+$, respectively.

Lemma 3.4.7 *Let* $\hat{\sigma}_k^- = \inf\{t \geq 1 : |\hat{A}_t^-| + |\hat{I}_t^-| \geq k$ *or* $\hat{A}_t^- = \varnothing\}$ *and* $\hat{\sigma}_k^+ = \inf\{t \geq 1 : |\hat{A}_t^+| + |\hat{I}_t^+| \geq k$ *or* $\hat{A}_t^+ = \varnothing\}$*. Then,*

$$\frac{1}{n}\sum_{i=1}^n \mathbb{E}_i\left[\left|\hat{I}_{\hat{\sigma}_k^-}^-\right| + \left|\hat{A}_{\hat{\sigma}_k^-}^-\right|\right] \leq k + k\sup_{\mathbf{x}\in\mathcal{S}} \lambda_-^{(m)}(\mathbf{x}) \quad \text{and} \quad \frac{1}{n}\sum_{i=1}^n \mathbb{E}_i\left[\left|\hat{I}_{\hat{\sigma}_k^+}^+\right| + \left|\hat{A}_{\hat{\sigma}_k^+}^+\right|\right] \leq k + k\sup_{\mathbf{x}\in\mathcal{S}} \lambda_+^{(m)}(\mathbf{x}).$$

Proof. Define $\mathcal{G}_r^-$ to be the sigma-algebra containing all the information of the exploration process of the in-component of vertex i up to the end of Step r and including the *identity* of T_{r+1}^-. Note that

$$\begin{aligned}
&\mathbb{E}_i\left[\left|\hat{I}_{\hat{\sigma}_k^-}^-\right| + \left|\hat{A}_{\hat{\sigma}_k^-}^-\right|\right] \\
&= \mathbb{E}_i\left[\left|\hat{I}_{\hat{\sigma}_k^- - 1}^-\right| + \left|\hat{A}_{\hat{\sigma}_k^- - 1}^-\right| + \sum_{j=1}^n Z_{j,\hat{T}_{\hat{\sigma}_k^-}^-}\right] \\
&\leq k - 1 + \sum_{r=1}^k \mathbb{E}_i\left[1(\hat{\sigma}_k^- = r)\sum_{j=1}^n Z_{j,\hat{T}_r^-}\right] \\
&= k - 1 + \sum_{r=1}^k \mathbb{E}_i\left[1(\hat{\sigma}_k^- > r-1)\mathbb{E}\left[1\left(\sum_{j=1}^n Z_{j,\hat{T}_r^-} \geq k - \left|\hat{A}_{r-1}^-\right| - \left|\hat{I}_{r-1}^-\right|\right)\sum_{j=1}^n Z_{j,\hat{T}_r^-}\middle|\mathcal{G}_{r-1}^-\right]\right].
\end{aligned}$$

Note that in the last equality the term that would correspond to $\{\hat{A}_r^- = \varnothing\}$ in the description of the event $\{\hat{\sigma}_k^- = r\}$ vanishes since $\sum_{j=1}^n Z_{j,\hat{T}_r^-} = 0$ in that case. Now use the observation that $\sum_{j=1}^n Z_{ji}$ is a Poisson random variable with mean $\sum_{j=1}^n r_{ji}^{(m,n)} = \lambda_-^{(m)}(\mathbf{X}_i)$, and the identity $E[X1(X \geq j)] \leq E[X] = \lambda$ when X is Poisson(λ), to obtain that

$$\begin{aligned}
&\sum_{r=1}^k \mathbb{E}_i\left[1(\hat{\sigma}_k^- > r-1)\mathbb{E}\left[1\left(\sum_{j=1}^n Z_{j,\hat{T}_r^-} \geq k - \left|\hat{A}_{r-1}^-\right| - \left|\hat{I}_{r-1}^-\right|\right)\sum_{j=1}^n Z_{j,\hat{T}_r^-}\middle|\mathcal{G}_{r-1}^-\right]\right] \\
&\leq \sum_{r=1}^k \mathbb{E}_i\left[1(\hat{\sigma}_k^- > r-1)\lambda_-^{(m)}(\mathbf{X}_{\hat{T}_r^-})\right] \\
&\leq k\sup_{\mathbf{x}\in\mathcal{S}} \lambda_-^{(m)}(\mathbf{x}).
\end{aligned}$$

The proof for the outbound tree is essentially the same and is therefore omitted. ■

The next result is a technical lemma giving an explicit upper bound for the ratio of independent Poisson random variables.

Lemma 3.4.8 *Let X, Y be independent Poisson random variables with means λ and μ, respectively. Let $a, b \geq 0$. Then,*

$$E\left[\frac{a+X}{b+X+Y}\cdot 1(b+X+Y\geq 1)\right] \leq \frac{2a}{b+1} + \frac{\lambda}{\lambda+\mu}(1-e^{-\lambda-\mu}).$$

Proof. Recall that X given $X+Y=n$ is a Binomial$(n, \lambda/(\lambda+\mu))$. Hence,

$$\begin{aligned}
&E\left[\frac{a+X}{b+X+Y}\cdot 1(b+X+Y\geq 1)\right] \\
&= E\left[\frac{a}{b}\cdot 1(X+Y=0, b\geq 1)\right] + E\left[\frac{a+X}{b+X+Y}\cdot 1(X+Y\geq 1)\right] \\
&= \frac{a}{b}1(b\geq 1)P(X+Y=0) + \sum_{n=1}^{\infty}\frac{E[a+X|X+Y=n]}{b+n}P(X+Y=n).
\end{aligned}$$

Now use the observation that X given $X+Y=n$ is a binomial with parameters $(n, \lambda/(\mu+\lambda))$ to obtain that

$$\begin{aligned}
&\sum_{n=1}^{\infty}\frac{E[a+X|X+Y=n]}{b+n}P(X+Y=n) \\
&= \sum_{n=1}^{\infty}\frac{a+n\lambda/(\mu+\lambda)}{b+n}P(X+Y=n) \\
&= a\sum_{n=1}^{\infty}\frac{1}{b+n}P(X+Y=n) + \frac{\lambda}{\mu+\lambda}\sum_{n=1}^{\infty}\frac{n}{b+n}P(X+Y=n) \\
&\leq \frac{a}{b+1}P(X+Y\geq 1) + \frac{\lambda}{\mu+\lambda}P(X+Y\geq 1) \\
&= \left(\frac{a}{b+1}+\frac{\lambda}{\lambda+\mu}\right)P(X+Y\geq 1).
\end{aligned}$$

Using the observation that $(a/b)1(b\geq 1)\leq 2a/(b+1)$ gives that

$$E\left[\frac{a+X}{b+X+Y}\cdot 1(b+X+Y\geq 1)\right] \leq \frac{2a}{b+1} + \frac{\lambda}{\lambda+\mu}P(X+Y\geq 1),$$

which completes the proof. ■

The following result constitutes a key step of the proof of Theorem 3.4.6 by providing an upper estimate for the distribution of the *identities* of the active nodes $\hat{T}_r^-$ and $\hat{T}_r^+$.

Lemma 3.4.9 *Let h be a nonnegative function in $\mathcal{S}$, then*

$$\mathbb{E}_i\left[1(\hat{A}^-_{r-1}\neq\varnothing)h(\mathbf{X}_{\hat{T}^-_r})\right]\leq\sum_{s=0}^{r-1}\binom{r-1}{s}2^{r-1-s}(\Gamma_-^{(m,n)})^s h(\mathbf{X}_i)$$

and

$$\mathbb{E}_i\left[1(\hat{A}^+_{r-1}\neq\varnothing)h(\mathbf{X}_{\hat{T}^+_r})\right]\leq\sum_{s=0}^{r-1}\binom{r-1}{s}2^{r-1-s}(\Gamma_+^{(m,n)})^s h(\mathbf{X}_i),$$

where $\Gamma_-^{(m,n)}$ and $\Gamma_+^{(m,n)}$ are the following linear integral operators:

$$\Gamma_-^{(m,n)}h(\mathbf{x})=\int_{\mathcal{S}}\frac{\mu(\mathcal{J}^{(m)}_{\vartheta(\mathbf{y})})}{\mu_n(\mathcal{J}^{(m)}_{\vartheta(\mathbf{y})})}\cdot\frac{(1-e^{-\lambda_-^{(m)}(\mathbf{x})})}{\lambda_-^{(m)}(\mathbf{x})}\cdot\kappa_m(\mathbf{y},\mathbf{x})h(\mathbf{y})\,\mu_n(d\mathbf{y})$$

and

$$\Gamma_+^{(m,n)}h(\mathbf{x})=\int_{\mathcal{S}}\frac{\mu(\mathcal{J}^{(m)}_{\vartheta(\mathbf{y})})}{\mu_n(\mathcal{J}^{(m)}_{\vartheta(\mathbf{y})})}\cdot\frac{(1-e^{-\lambda_+^{(m)}(\mathbf{x})})}{\lambda_+^{(m)}(\mathbf{x})}\cdot\kappa_m(\mathbf{x},\mathbf{y})h(\mathbf{y})\,\mu_n(d\mathbf{y}),$$

with $\vartheta(\mathbf{x})=t$ if and only if $\mathbf{x}\in\mathcal{J}^{(m)}_t$.

Proof. Let $\mathbf{W}^-_t=(W^-_{t,1},\dots,W^-_{t,n})$ denote the process that keeps track of the *identities* of the vertices in the active set $\hat{A}^-_t$ for $t\geq 0$; that is, $W^-_{t,j}$ denotes the number of tree nodes with *identity* j in $\hat{A}^-_t$. Then,

$$\begin{aligned}\mathbb{P}_i(\hat{A}^-_{r-1}\neq\varnothing,\hat{T}^-_r=l)&=\mathbb{E}_i\left[1\left(\|\mathbf{W}^-_{r-1}\|_1\geq 1\right)\mathbb{P}(\hat{T}^-_r=l|\mathbf{W}^-_{r-1})\right]\\&=\mathbb{E}_i\left[\frac{W^-_{r-1,l}}{\|\mathbf{W}^-_{r-1}\|_1}\cdot 1\left(\|\mathbf{W}^-_{r-1}\|_1\geq 1\right)\right]\\&=\mathbb{E}_i\left[\mathbb{E}\left[\frac{W^-_{r-1,l}}{\|\mathbf{W}^-_{r-1}\|_1}\cdot 1\left(\|\mathbf{W}^-_{r-1}\|_1\geq 1\right)\middle|\mathbf{W}^-_{r-2}\right]1(\|\mathbf{W}^-_{r-2}\|_1\geq 1)\right],\end{aligned}$$

where $\|(x_1,\dots,x_n)\|_1=|x_1|+\cdots+|x_n|$. Now let $(h_1,\dots,h_n)=(h(\mathbf{X}_1),\dots,h(\mathbf{X}_n))$ and note that

$$\begin{aligned}\mathbb{E}_i\left[1(\hat{A}^-_{r-1}\neq\varnothing)h(\mathbf{X}_{\hat{T}^-_r})\right]&=\sum_{l=1}^n h_l\mathbb{P}_i(\hat{A}^-_{r-1}\neq\varnothing,\hat{T}^-_r=l)\\&=\mathbb{E}_i\left[\sum_{l=1}^n h_l\mathbb{E}\left[\frac{W^-_{r-1,l}}{\|\mathbf{W}^-_{r-1}\|_1}\cdot 1\left(\|\mathbf{W}^-_{r-1}\|_1\geq 1\right)\middle|\mathbf{W}^-_{r-2}\right]1(\|\mathbf{W}^-_{r-2}\|_1\geq 1)\right].\end{aligned}$$

Moreover, provided $\|\mathbf{W}^-_{r-2}\|_1 \geq 1$, we have

$$\begin{aligned}
&\mathbb{E}\left[\frac{W^-_{r-1,l}}{\|\mathbf{W}^-_{r-1}\|_1}\cdot 1\left(\|\mathbf{W}^-_{r-1}\|_1 \geq 1\right)\middle|\, \mathbf{W}^-_{r-2}\right] \\
&= \sum_{s=1}^{n} \mathbb{P}(\hat{T}^-_{r-1} = s|\mathbf{W}^-_{r-2})\,\mathbb{E}\left[\frac{W^-_{r-1,l}}{\|\mathbf{W}^-_{r-1}\|_1}\cdot 1\left(\|\mathbf{W}^-_{r-1}\|_1 \geq 1\right)\middle|\, \mathbf{W}^-_{r-2}, \{\hat{T}^-_{r-1} = s\}\right] \\
&= \sum_{1\leq s\leq n} \frac{W^-_{r-2,s}}{\|\mathbf{W}^-_{r-2}\|_1}\mathbb{E}\left[\frac{W^-_{r-2,l} + Z_{ls} - 1(s=l)}{\sum_{j=1}^n (W^-_{r-2,j} + Z_{js}) - 1}\cdot 1\left(\sum_{j=1}^n (W^-_{r-2,j} + Z_{js}) \geq 2\right)\middle|\, \mathbf{W}^-_{r-2}\right] \\
&\leq \sum_{s=1}^{n} \frac{W^-_{r-2,s}}{\|\mathbf{W}^-_{r-2}\|_1}\mathbb{E}\left[\frac{W^-_{r-2,l} + Z_{ls}}{\sum_{j=1}^n (W^-_{r-2,j} + Z_{js}) - 1}\cdot 1\left(\sum_{j=1}^n (W^-_{r-2,j} + Z_{js}) \geq 2\right)\middle|\, \mathbf{W}^-_{r-2}\right].
\end{aligned}$$

Now use Lemma 3.4.8 with $a = W^-_{r-2,l}$, $b = \sum_{j=1}^n W^-_{r-2,j} - 1$, $X = Z_{ls}$ and $Y = \sum_{j\neq l} Z_{js}$ to obtain that

$$\begin{aligned}
&\sum_{s=1}^{n} \frac{W^-_{r-2,s}}{\|\mathbf{W}^-_{r-2}\|_1}\mathbb{E}\left[\frac{W^-_{r-2,l} + Z_{ls}}{\sum_{j=1}^n (W^-_{r-2,j} + Z_{js}) - 1}\cdot 1\left(\sum_{j=1}^n (W^-_{r-2,j} + Z_{js}) \geq 2\right)\middle|\, \mathbf{W}^-_{r-2}\right] \\
&\leq \sum_{s=1}^{n} \frac{W^-_{r-2,s}}{\|\mathbf{W}^-_{r-2}\|_1}\left(\frac{2W^-_{r-2,l}}{\|\mathbf{W}^-_{r-2}\|_1} + \frac{r^{(m,n)}_{ls}}{\sum_{j=1}^n r^{(m,n)}_{js}}\left(1 - e^{-\sum_{j=1}^n r^{(m,n)}_{js}}\right)\right) \\
&=: \frac{2W^-_{r-2,l}}{\|\mathbf{W}^-_{r-2}\|_1} + \sum_{s=1}^{n} \frac{W^-_{r-2,s}}{\|\mathbf{W}^-_{r-2}\|_1}\cdot \gamma^{(m,n)}_{ls},
\end{aligned}$$

where

$$\gamma^{(m,n)}_{ls} = \frac{r^{(m,n)}_{ls}}{\sum_{j=1}^n r^{(m,n)}_{js}}\left(1 - e^{-\sum_{j=1}^n r^{(m,n)}_{js}}\right) = r^{(m,n)}_{ls}\frac{(1 - e^{-\lambda^{(m)}_-(\mathbf{X}_s)})}{\lambda^{(m)}_-(\mathbf{X}_s))},$$

and we use the convention that $(1 - e^{-0})/0 \equiv 1$. It follows that

$$\begin{aligned}
&\mathbb{E}_i\left[1(\hat{A}^-_{r-1} \neq \varnothing)h(\mathbf{X}_{\hat{T}^-_r})\right] \\
&\leq \mathbb{E}_i\left[\sum_{l=1}^{n} h_l\left\{\frac{2W^-_{r-2,l}}{\|\mathbf{W}^-_{r-2}\|_1} + \sum_{s=1}^{n}\frac{W^-_{r-2,s}}{\|\mathbf{W}^-_{r-2}\|_1}\cdot\gamma^{(m,n)}_{ls}\right\}1(\|\mathbf{W}^-_{r-2}\|_1 \geq 1)\right] \\
&= 2\mathbb{E}_i\left[1(\hat{A}^-_{r-2} \neq \varnothing)h_{\hat{T}^-_{r-1}}\right] + \mathbb{E}_i\left[\sum_{s=1}^{n}\frac{W^-_{r-2,s}}{\|\mathbf{W}^-_{r-2}\|_1}1(\|\mathbf{W}^-_{r-2}\|_1 \geq 1)\sum_{l=1}^{n} h_l\cdot\gamma^{(m,n)}_{ls}\right] \\
&= 2\mathbb{E}_i\left[1(\hat{A}^-_{r-2} \neq \varnothing)h(\mathbf{X}_{\hat{T}^-_{r-1}})\right] + \mathbb{E}_i\left[1(\hat{A}_{r-2} \neq \varnothing)\Gamma^{(m,n)}_- h(\mathbf{X}_{\hat{T}^-_{r-1}})\right].
\end{aligned}$$

Letting $a_{r,s} = \mathbb{E}_i\left[1(\hat{A}^-_{r-1} \neq \varnothing)(\Gamma_-^{(m,n)})^s h(\mathbf{X}_{\hat{T}_r^-})\right]$, and iterating $r-2$ times we obtain that

$$a_{r,0} \leq 2a_{r-1,0} + a_{r-1,1} \leq \sum_{s=0}^{r-1} \binom{r-1}{s} 2^{r-1-s} a_{1,s},$$

which yields

$$\mathbb{E}_i\left[1(\hat{A}^-_{r-1} \neq \varnothing) h(\mathbf{X}_{\hat{T}_r^-})\right] \leq \sum_{s=0}^{r-1} \binom{r-1}{s} 2^{r-1-s} (\Gamma_-^{(m,n)})^s h(\mathbf{X}_i).$$

The proof for $\mathbb{E}_i\left[1(\hat{A}^+_{r-1} \neq \varnothing) h(\mathbf{X}_{\hat{T}_r^+})\right]$ is essentially the same and is therefore omitted. ■

Lemma 3.4.10 *Let $H(n,m,k,\epsilon)$ be defined as in Theorem 3.4.6, then*

$$H(n,m,k,\epsilon) \xrightarrow{P} \hat{H}(m,k,\epsilon) \qquad n \to \infty,$$

where

$$\hat{H}(m,k,\epsilon) = 4\epsilon k^2 + 2\epsilon k^2 \left(1 + \sup_{\mathbf{x}\in\mathcal{S}} \lambda_-^{(m)}(\mathbf{x})\right) + \sum_{r=1}^{k}\sum_{s=0}^{r-1} \binom{r-1}{s} 2^{r-1-s}$$
$$\cdot \left\{ \int_{\mathcal{S}} (\Gamma_-^{(m)})^s g^-_{m,\epsilon}(\mathbf{x})\mu(d\mathbf{x}) + \int_{\mathcal{S}} (\Gamma_+^{(m)})^s g^+_{m,\epsilon}(\mathbf{x})\mu(d\mathbf{x}) \right\},$$

where

$$g^-_{m,\epsilon}(\mathbf{x}) = \min\left\{1,\ (1+5\epsilon)\lambda_-(\mathbf{x}) - \lambda_-^{(m)}(\mathbf{x})\right\},$$
$$g^+_{m,\epsilon}(\mathbf{x}) = \min\left\{1,\ (1+5\epsilon)\lambda_+(\mathbf{x}) - \lambda_+^{(m)}(\mathbf{x})\right\},$$

and the linear integral operators $\Gamma_-^{(m)}$ and $\Gamma_+^{(m)}$ are given by

$$\Gamma_-^{(m)} h(\mathbf{x}) = \int_{\mathcal{S}} \frac{(1 - e^{-\lambda_-^{(m)}(\mathbf{x})})}{\lambda_-^{(m)}(\mathbf{x})} \cdot \kappa_m(\mathbf{y},\mathbf{x}) h(\mathbf{y})\mu(d\mathbf{y}),$$
$$\Gamma_+^{(m)} h(\mathbf{x}) = \int_{\mathcal{S}} \frac{(1 - e^{-\lambda_+^{(m)}(\mathbf{x})})}{\lambda_+^{(m)}(\mathbf{x})} \cdot \kappa_m(\mathbf{x},\mathbf{y}) h(\mathbf{y})\mu(d\mathbf{y}).$$

Furthermore, for any fixed $k \geq 1$,

$$\lim_{m\nearrow\infty} \lim_{\epsilon\downarrow 0} \hat{H}(m,k,\epsilon) = 0.$$

Proof. Start by noting that Assumption 3.3.1(a) implies that $1(\Omega_{m,n}) \xrightarrow{P} 1$ as $n \to \infty$, so the convergence of $H(n, m, k, \epsilon)$ will follow once we show that

$$\int_{\mathcal{S}} (\Gamma_-^{(m,n)})^s g_{m,n,\epsilon}^-(\mathbf{x})\mu_n(d\mathbf{x}) \xrightarrow{P} \int_{\mathcal{S}} (\Gamma_-^{(m)})^s g_{m,\epsilon}^-(\mathbf{x})\mu(d\mathbf{x}) \tag{3.4.11}$$

and

$$\int_{\mathcal{S}} (\Gamma_+^{(m,n)})^s g_{m,n,\epsilon}^+(\mathbf{x})\mu_n(d\mathbf{x}) \xrightarrow{P} \int_{\mathcal{S}} (\Gamma_+^{(m)})^s g_{m,\epsilon}^+(\mathbf{x})\mu(d\mathbf{x}) \tag{3.4.12}$$

as $n \to \infty$ for any fixed $s \in \mathbb{N}$. Let

$$w_m^-(\mathbf{y},\mathbf{x}) = \frac{(1 - e^{-\lambda_-^{(m)}(\mathbf{x})})}{\lambda_-^{(m)}(\mathbf{x})} \cdot \kappa_m(\mathbf{y},\mathbf{x}) \qquad \text{and} \qquad r_n(\mathbf{y}) = \frac{\mu(\mathcal{J}_{\vartheta(\mathbf{y})}^{(m)})}{\mu_n(\mathcal{J}_{\vartheta(\mathbf{y})}^{(m)})},$$

and note that for any function h and $\mathbf{x} \in \mathcal{J}_j^{(m)}$, we have

$$\Gamma_-^{(m,n)} h(\mathbf{x}) = \int_{\mathcal{S}} r_n(\mathbf{y}) w_m^-(\mathbf{y},\mathbf{x}) h(\mathbf{y})\mu_n(d\mathbf{y}) =: \sum_{i=1}^{M_m} \mathcal{I}_i^{(m,n)}(h) d_{i,j}^{(m,n)},$$

where

$$\begin{aligned} d_{i,j}^{(m,n)} &= r_n(\mathbf{y}) w_m^-(\mathbf{y},\mathbf{x}) \quad \text{for all } \mathbf{y} \in \mathcal{J}_i^{(m)}, \mathbf{x} \in \mathcal{J}_j^{(m)}, \\ \mathcal{I}_i^{(m,n)}(h) &= \int_{\mathcal{J}_i^{(m)}} h(\mathbf{y})\mu_n(d\mathbf{y}). \end{aligned}$$

In general, for $s \geq 1$, we have that

$$(\Gamma_-^{(m,n)})^s h(\mathbf{x}) = \sum_{i=1}^{M_m} \mathcal{I}_i^{(m,n)}(h)((\mathbf{D}^{(m,n)})^s)_{i,j} \qquad \text{for } \mathbf{x} \in \mathcal{J}_j^{(m)}$$

and

$$\int_{\mathcal{S}} (\Gamma_-^{(m,n)})^s h(\mathbf{x})\mu_n(d\mathbf{x}) = \sum_{j=1}^{M_m} \sum_{i=1}^{M_m} \mathcal{I}_i^{(m,n)}(h)((\mathbf{D}^{(m,n)})^s)_{i,j}\mu_n(\mathcal{J}_j^{(m)}),$$

where $\mathbf{D}^{(m,n)}$ is the $M_m \times M_m$ matrix whose (i,j)th component is $d_{i,j}^{(m,n)}$. Define $\mathbf{D}^{(m)}$ to be the matrix whose (i,j)th component is $d_{i,j}^{(m)} = w_m^-(\mathbf{y},\mathbf{x})$ for all $\mathbf{y} \in \mathcal{J}_i^{(m)}, \mathbf{x} \in \mathcal{J}_j^{(m)}$. Since by Assumption 3.3.1(a) we have that $\mu_n(\mathcal{J}_j^{(m)}) \xrightarrow{P} \mu(\mathcal{J}_j^{(m)})$ and $d_{i,j}^{(m,n)} \xrightarrow{P} d_{i,j}^{(m)}$ as $n \to \infty$ for all $1 \leq i,j \leq M_m$, it follows that

$$\lim_{n\to\infty} \int_{\mathcal{S}} (\Gamma_-^{(m,n)})^s g_{m,n,\epsilon}^-(\mathbf{x})\mu_n(d\mathbf{x}) = \sum_{j=1}^{M_m} \sum_{i=1}^{M_m} \left(\lim_{n\to\infty} \mathcal{I}_i^{(m,n)}(g_{m,n,\epsilon}^-)\right) ((\mathbf{D}^{(m)})^s)_{i,j}\mu(\mathcal{J}_j^{(m)}),$$

assuming the last limit exists for each $1 \leq i \leq M_m$. To see that it does let $\hat{g}^-_{m,n,\epsilon}(\mathbf{x}) = \min\left\{1,\, (1+5\epsilon)\lambda^-_n(\mathbf{x}) - \lambda^-_{m,n}(\mathbf{x})\right\}$ and note that by Lemma 3.4.2,

$$\left|\mathcal{I}_i^{(m,n)}(g^-_{m,n,\epsilon}) - \mathcal{I}_i^{(m,n)}(\hat{g}^-_{m,n,\epsilon})\right| \leq (1+\epsilon)\frac{1}{n}\sum_{l\in\mathcal{J}_i^{(m)}}\sum_{j=1}^{n}(p_{jl}^{(n)} + q_{jl}^{(n)})1(B_{jl}^c) \xrightarrow{P} 0$$

as $n \to \infty$. Now let $\mathbf{X}^{(n)}$ and $\mathbf{Y}^{(n)}$ be conditionally i.i.d. random variables (given $\mathscr{F}$) having distribution μ_n (as constructed in Lemma 3.4.1). Assumption 3.3.1 implies that $(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) \Rightarrow (\mathbf{X}, \mathbf{Y})$ as $n \to \infty$, where $\mathbf{X}$ and $\mathbf{Y}$ are i.i.d. with distribution μ, and Lemma 3.4.2 gives $\mathbb{E}\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)})\right] \xrightarrow{P} E[\kappa(\mathbf{X}, \mathbf{Y})]$. Therefore, by bounded convergence,

$$\begin{aligned}
\mathcal{I}_i^{(m,n)}(\hat{g}^-_{m,n,\epsilon}) &= \mathbb{E}\left[1(\mathbf{X}^{(n)} \in \mathcal{J}_i^{(m)}) \min\left\{1,\, (1+5\epsilon)\mathbb{E}\left[\kappa(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) - \kappa_m(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) \middle|\, \mathbf{X}^{(n)}\right]\right\}\right] \\
&\xrightarrow{P} E\left[1(\mathbf{X} \in \mathcal{J}_i^{(m)}) \min\left\{1,\, (1+5\epsilon)E\left[\kappa(\mathbf{X}, \mathbf{Y}) - \kappa_m(\mathbf{X}, \mathbf{Y}) \middle|\, \mathbf{X}\right]\right\}\right] \\
&=: \mathcal{I}_i^{(m)}(g^-_{m,\epsilon}),
\end{aligned}$$

as $n \to \infty$. We conclude that

$$\mathcal{I}_i^{(m,n)}(g^-_{m,n,\epsilon}) \xrightarrow{P} \mathcal{I}_i^{(m)}(g^-_{m,\epsilon})$$

as $n \to \infty$, and noting that

$$\sum_{j=1}^{M_m}\sum_{i=1}^{M_m} \mathcal{I}_i^{(m)}(g^-_{m,\epsilon})((\mathbf{D}^{(m)})^s)_{i,j}\mu(\mathcal{J}_j^{(m)}) = \int_{\mathcal{S}} (\Gamma_-^{(m)})^s g^-_{m,\epsilon}(\mathbf{x})\mu(d\mathbf{x})$$

completes the proof of (3.4.11). The proof for (3.4.12) is essentially the same and is therefore omitted. This concludes the proof that $H(n, m, k, \epsilon) \xrightarrow{P} \hat{H}(m, k, \epsilon)$ as $n \to \infty$. To compute the limit of $\hat{H}(m, k, \epsilon)$ note that by monotone convergence,

$$\lim_{\epsilon\downarrow 0} \hat{H}(m, k, \epsilon) = \sum_{k=1}^{k}\sum_{s=0}^{r-1}\binom{r-1}{s}2^{r-1-s}\left\{\int_{\mathcal{S}}(\Gamma_-^{(m)})^s g_m^-(\mathbf{x})\mu(d\mathbf{x}) + \int_{\mathcal{S}}(\Gamma_+^{(m)})^s g_m^+(\mathbf{x})\mu(d\mathbf{x})\right\},$$

where $g_m^{\pm}(\mathbf{x}) = \min\left\{1, \lambda_{\pm}(\mathbf{x}) - \lambda_{\pm}^{(m)}(\mathbf{x})\right\}$. Now let Γ_- and Γ_+ be the linear integral operators defined by

$$\Gamma_- h(\mathbf{x}) = \int_{\mathcal{S}} \frac{1 - e^{-\lambda_-(\mathbf{x})}}{\lambda_-(\mathbf{x})} \cdot \kappa(\mathbf{y}, \mathbf{x})h(\mathbf{y})\mu(d\mathbf{y}) \quad \text{and} \quad \Gamma_+ h(\mathbf{x}) = \int_{\mathcal{S}} \frac{1 - e^{-\lambda_+(\mathbf{x})}}{\lambda_+(\mathbf{x})} \cdot \kappa(\mathbf{x}, \mathbf{y})h(\mathbf{y})\mu(d\mathbf{y})$$

and note that by monotone convergence,

$$\lim_{m\nearrow\infty} \Gamma_{\pm}^{(m)} h(\mathbf{x}) = \Gamma_{\pm} h(\mathbf{x})$$

for any nonnegative function h. Moreover, for any $h : \mathcal{S} \to [0,1]$, we have that $\Gamma_\pm^{(m)} h(\mathbf{x}), \Gamma_\pm h(\mathbf{x}) \in [0,1]$, and therefore, the bounded convergence theorem gives

$$\lim_{m \nearrow \infty} \int_{\mathcal{S}} (\Gamma_\pm^{(m)})^s g_m^\pm(\mathbf{x}) \mu(d\mathbf{x}) = \int_{\mathcal{S}} (\Gamma_\pm)^s \left(\lim_{m \nearrow \infty} g_m^\pm \right) (\mathbf{x}) \mu(d\mathbf{x}) = 0.$$

This completes the proof. ■

We are now ready to give the proof of Theorem 3.4.6.

Proof of Theorem 3.4.6. To start, note that

$$\begin{aligned} &\mathbb{P}_i \left(\{\tau^- \leq \sigma_k^-\} \cup \{\tau^+ \leq \sigma_k^+\} \right) \\ &\leq \left\{ \mathbb{P}_i \left(\tau^- \leq \sigma_k^- \right) + \mathbb{P}_i \left(\{\tau^- > \sigma_k^-\} \cap \{\tau^+ \leq \sigma_k^+\} \right) \right\} 1(\Omega_{m,n}) + 1(\Omega_{m,n}^c), \end{aligned}$$

where the event $\Omega_{m,n}$ is defined in the statement of the theorem. To analyze the two probabilities, define $\mathcal{G}_m^-$ to be the sigma-algebra containing all the information of the exploration process of the in-component of vertex i up to the end of Step m and including the *identity* of T_{m+1}^-, and let $\mathcal{G}_m^+$ be the sigma-algebra containing all the information of the exploration process of the in-component of vertex i up to Step σ_k^-, and of its out-component up to the end of Step m, including the *identity* of T_{m+1}^+; note that $\mathcal{G}_m^- \subseteq \mathcal{G}_r^+$ for all $0 \leq m \leq \sigma_k^-$ and any $r \geq 0$. Next, for any $r \geq 1$ define the events

$$\begin{aligned} E_r^- &= \{|I_r^-| + |A_r^-| < k\}, \\ E_r^+ &= \{|I_r^+| + |A_r^+| < k\}, \end{aligned}$$

$$C_i^-(r) = \left\{ \max_{j \in [n], j \notin \{i\} \cup A_{r-1}^- \cup I_{r-1}^-} |Z_{ji} - Y_{ji}| + \max_{j \in A_{r-1}^- \cup I_{r-1}^-} Z_{ji} + Z_{ii}^* = 0 \right\},$$

$$C_i^+(r) = \left\{ \max_{j \in [n], j \notin \{i\} \cup I_{\sigma_k^-}^- \cup A_{\sigma_k^-}^- \cup A_{r-1}^+ \cup I_{r-1}^+} |\tilde{Z}_{ij} - Y_{ij}| + \max_{j \in A_{r-1}^+ \cup I_{r-1}^+} \tilde{Z}_{ij} + \max_{j \in \{i\} \cup I_{\sigma_k^-}^- \cup A_{\sigma_k^-}^-} \tilde{Z}_{ji}^* = 0 \right\}.$$

Now, use Remark 3.4.5 to obtain that on the event $\Omega_{m,n}$,

$$
\begin{aligned}
&\mathbb{P}_i\left(\tau^- \le \sigma_k^-\right) + \mathbb{P}_i\left(\{\tau^- > \sigma_k^-\} \cap \{\tau^+ \le \sigma_k^+\}\right) \\
&= \sum_{r=1}^k \left\{\mathbb{P}_i\left(r = \tau^- \le \sigma_k^-\right) + \mathbb{P}_i\left(\{\tau^- > \sigma_k^-\} \cap \{r = \tau^+ \le \sigma_k^+\}\right)\right\} \\
&\le \sum_{r=1}^k \mathbb{P}_i\left(\tau^- > r-1, A_{r-1}^- \ne \varnothing, E_{r-1}^-, (C_{T_r^-}^-(r))^c\right) \\
&\quad + \sum_{r=1}^k \mathbb{P}_i\left(\tau^- > \sigma_k^-,\, \tau^+ > r-1, A_{r-1}^+ \ne \varnothing, E_{r-1}^+, (C_{T_r^+}^+(r))^c\right) \\
&= \sum_{r=1}^k \mathbb{E}_i\left[1(\tau^- > r-1, A_{r-1}^- \ne \varnothing, E_{r-1})\mathbb{P}\left((C_{T_r^-}^-(r))^c \middle| \mathcal{G}_{r-1}^-\right)\right] \\
&\quad + \sum_{r=1}^k \mathbb{E}_i\left[1(\tau^- > \sigma_k^-,\, \tau^+ > r-1, A_{r-1}^+ \ne \varnothing, E_{r-1})\mathbb{P}\left((C_{T_r^+}^+(r))^c \middle| \mathcal{G}_{r-1}^+\right)\right].
\end{aligned}
$$

To analyze the two conditional probabilities in the last expressions, note that the union bound and the independence of the $\{U_{ij} : 1 \le i, j \le n\}$ from everything else give

$$
\begin{aligned}
\mathbb{P}\left((C_{T_r^-}^-(r))^c \middle| \mathcal{G}_{r-1}^-\right) \le &\sum_{j\in[n], j\notin\{T_r^-\}\cup A_{r-1}^-\cup I_{r-1}^-} \mathbb{P}(|Z_{j,T_r^-} - Y_{j,T_r^-}| > 0|T_r^-) &(3.4.13)\\
&+ \sum_{j\in\{T_r^-\}\cup A_{r-1}^-\cup I_{r-1}^-} \mathbb{P}(Z_{j,T_r^-} > 0|T_r^-), &(3.4.14)
\end{aligned}
$$

and

$$
\begin{aligned}
\mathbb{P}\left((C_{T_r^+}^+(r))^c \middle| \mathcal{G}_{r-1}^+\right) \le &\sum_{j\in[n], j\notin\{T_r^+\}\cup I_{\sigma_k^-}^-\cup A_{\sigma_k^-}^-\cup A_{r-1}^+\cup I_{r-1}^+} \mathbb{P}(|\tilde{Z}_{T_r^+,j} - Y_{T_r^+,j}| > 0|T_r^+) &(3.4.15)\\
&+ \sum_{j\in\{T_r^+\}\cup I_{\sigma_k^-}^-\cup A_{\sigma_k^-}^-\cup A_{r-1}^+\cup I_{r-1}^+} \mathbb{P}(\tilde{Z}_{T_r^+,j} > 0|T_r^+). &(3.4.16)
\end{aligned}
$$

To analyze (3.4.13) note that on the event B_{ji} we have that $(1-\epsilon)q_{ji}^{(n)} \le p_{ji}^{(n)} < (1+$

$\epsilon)q_{ji}^{(n)} \leq (1+\epsilon)\epsilon < 1$, which implies that on the event B_{ji} we have

$$\begin{aligned}
\mathbb{P}(|Y_{ji} - Z_{ji}| > 0) &= (p_{ji}^{(n)} - r_{ji}^{(m,n)})1(p_{ji}^{(n)} > r_{ji}^{(m,n)}) \\
&\quad + (e^{-r_{ji}^{(m,n)}} - 1 + p_{ji}^{(n)})1(1 - e^{-r_{ji}^{(m,n)}} < p_{ji}^{(n)} \leq r_{ji}^{(m,n)}) \\
&\quad + (1 - p_{ji}^{(n)} - e^{-r_{ji}^{(m,n)}})1(p_{ji}^{(n)} < 1 - e^{-r_{ji}^{(m,n)}}) + 1 - e^{-r_{ji}^{(m,n)}} - e^{-r_{ji}^{(m,n)}} r_{ji}^{(m,n)} \\
&\leq (p_{ji}^{(n)} - r_{ji}^{(m,n)})1(p_{ji}^{(n)} > r_{ji}^{(m,n)}) + (p_{ji}^{(n)} - r_{ji}^{(m,n)})1(1 - e^{-r_{ji}^{(m,n)}} < p_{ji}^{(n)} \leq r_{ji}^{(m,n)}) \\
&\quad + (r_{ji}^{(m,n)} - p_{ji}^{(n)})1(p_{ji}^{(n)} < 1 - e^{-r_{ji}^{(m,n)}}) + (r_{ji}^{(m,n)})^2 \\
&= |p_{ji}^{(n)} - r_{ji}^{(m,n)}| + (r_{ji}^{(m,n)})^2,
\end{aligned}$$

where we have used the inequalities $e^{-x} - 1 \leq -x + x^2/2$, $1 - e^{-x} \leq x$, and $1 - e^{-x} - e^{-x}x \leq x^2/2$ for $x \geq 0$. It follows that if we let $q_{ji}^{(m,n)} = \kappa_m(\mathbf{X}_j, \mathbf{X}_i)/n$, then, on the event $\Omega_{m,n}$, where we have $(1-\epsilon)q_{ji}^{(m,n)} \leq r_{ji}^{(m,n)} \leq (1+\epsilon)q_{ji}^{(m,n)}$, we have

$$\begin{aligned}
&\mathbb{P}(|Y_{ji} - Z_{ji}| > 0)1(B_{ji}) \\
&\leq \left(|p_{ji}^{(n)} - r_{ji}^{(m,n)}| + (r_{ji}^{(m,n)})^2\right) 1(B_{ji}) \\
&\leq \left(|p_{ji}^{(n)} - q_{ji}^{(n)}| + q_{ji}^{(n)} - q_{ji}^{(m,n)} + |q_{ji}^{(m,n)} - r_{ji}^{(m,n)}| + (1+\epsilon)^2(q_{ji}^{(m,n)})^2\right) 1(B_{ji}) \\
&\leq \epsilon q_{ji}^{(n)} + q_{ji}^{(n)} - q_{ji}^{(m,n)} + \epsilon q_{ji}^{(m,n)} + (1+\epsilon)^2 \epsilon q_{ji}^{(m,n)} \\
&\leq (1+5\epsilon)q_{ji}^{(n)} - q_{ji}^{(m,n)}.
\end{aligned}$$

On the other hand, note that on the event $\Omega_{m,n}$ we have

$$\begin{aligned}
\mathbb{P}(|Y_{ji} - Z_{ji}| > 0)1(B_{ji}^c) &\leq \mathbb{P}(Y_{ji} + Z_{ji} > 0)1(B_{ji}^c) \\
&\leq \min\left\{1, p_{ji}^{(n)} + r_{ji}^{(m,n)}\right\} 1(B_{ji}^c) \\
&\leq (1+\epsilon)(p_{ji}^{(n)} + q_{ji}^{(n)})1(B_{ji}^c) =: \mathcal{B}_n(\mathbf{X}_j, \mathbf{X}_i).
\end{aligned}$$

Hence, on the event $\Omega_{m,n}$, (3.4.13) is bounded from above by

$$\begin{aligned}
&\sum_{j\in[n]} \left\{(1+5\epsilon)q_{j,T_r^-}^{(n)} - q_{j,T_r^-}^{(m,n)}\right\} + \sum_{j\in[n], j\notin\{T_r^-\}\cup A_{r-1}^-\cup I_{r-1}^-} \mathcal{B}_n(\mathbf{X}_j, \mathbf{X}_{T_r^-}) \\
&\leq (1+5\epsilon)\lambda_n^-(\mathbf{X}_{T_r^-}) - \lambda_{m,n}^-(\mathbf{X}_{T_r^-}) + \sum_{j\in[n], j\notin\{T_r^-\}\cup A_{r-1}^-\cup I_{r-1}^-} \mathcal{B}_n(\mathbf{X}_j, \mathbf{X}_{T_r^-}),
\end{aligned}$$

where we have used the observation that $\sum_{j=1}^n q_{ji}^{(n)} = \lambda_n^-(\mathbf{X}_i)$ and $\sum_{j=1}^n q_{ji}^{(m,n)} = \int_{\mathcal{S}} \kappa_m(\mathbf{y}, \mathbf{X}_i)\mu_n(d\mathbf{y}) = \lambda_{m,n}^-(\mathbf{X}_i)$.

To analyze (3.4.14), note that on the event $\Omega_{m,n}$,

$$\mathbb{P}(Z_{ji} \geq 1) = 1 - e^{-r_{ji}^{(m,n)}} \leq r_{ji}^{(m,n)} \leq (1+\epsilon)q_{ji}^{(m,n)}1(B_{ji}) + (1+\epsilon)q_{ji}^{(m,n)}1(B_{ji}^c) \leq 2\epsilon + \mathcal{B}_n(\mathbf{X}_j, \mathbf{X}_i).$$

We have thus obtained that, on the event $\Omega_{m,n}$,

$$\begin{aligned}\mathbb{P}\left((C^-_{T^-_r})^c\middle|\,\mathcal{G}^-_{r-1}\right) &\le (1+5\epsilon)\lambda^-_n(\mathbf{X}_{T^-_r}) - \lambda^-_{m,n}(\mathbf{X}_{T^-_r}) + \sum_{j\in[n],j\notin\{T^-_r\}\cup A^-_{r-1}\cup I^-_{r-1}} \mathcal{B}_n(\mathbf{X}_j,\mathbf{X}_{T^-_r})\\ &\quad + \sum_{j\in\{T^-_r\}\cup A^-_{r-1}\cup I^-_{r-1}} \left(2\epsilon + \mathcal{B}_n(\mathbf{X}_j,\mathbf{X}_{T^-_r})\right)\\ &\le (1+5\epsilon)\lambda^-_n(\mathbf{X}_{T^-_r}) - \lambda^-_{m,n}(\mathbf{X}_{T^-_r}) + 2\epsilon\left|\{T^-_r\}\cup A^-_{r-1}\cup I^-_{r-1}\right| + \sum_{j\in[n]}\mathcal{B}_n(\mathbf{X}_j,\mathbf{X}_{T^-_r}).\end{aligned}$$

The same arguments yield that, on the event $\Omega_{m,n}$, (3.4.15) is bounded by

$$(1+5\epsilon)\lambda^+_n(\mathbf{X}_{T^+_r}) - \lambda^+_{m,n}(\mathbf{X}_{T^+_r}) + \sum_{j\in[n],j\notin\{T^+_r\}\cup I^-_{\sigma^-_k}\cup A^-_{\sigma^-_k}\cup A^+_{r-1}\cup I^+_{r-1}} \mathcal{B}_n(\mathbf{X}_{T^+_r},\mathbf{X}_j),$$

and (3.4.16) is bounded by

$$\sum_{j\in\{T^+_r\}\cup I^-_{\sigma^-_k}\cup A^-_{\sigma^-_k}\cup A^+_{r-1}\cup I^+_{r-1}} \left(2\epsilon + \mathcal{B}_n(\mathbf{X}_{T^+_r},\mathbf{X}_j)\right).$$

Hence, on the event $\Omega_{m,n}$,

$$\begin{aligned}\mathbb{P}\left((C^+_{T^+_r})^c\middle|\,\mathcal{G}^+_{r-1}\right) &\le (1+5\epsilon)\lambda^+_n(\mathbf{X}_{T^+_r}) - \lambda^+_{m,n}(\mathbf{X}_{T^+_r}) + 2\epsilon\left|\{T^+_r\}\cup I^-_{\sigma^-_k}\cup A^-_{\sigma^-_k}\cup A^+_{r-1}\cup I^+_{r-1}\right|\\ &\quad + \sum_{j\in[n]}\mathcal{B}_n(\mathbf{X}_{T^+_r},\mathbf{X}_j).\end{aligned}$$

To simplify the notation, define the functions:

$$\begin{aligned}g^-_{m,n,\epsilon}(\mathbf{X}_l) &= \min\left\{1,\ (1+5\epsilon)\lambda^-_n(\mathbf{X}_l) - \lambda^-_{m,n}(\mathbf{X}_l) + \sum_{j\in[n]}\mathcal{B}_n(\mathbf{X}_j,\mathbf{X}_l)\right\} \quad \text{and}\\ g^-_{m,n,\epsilon}(\mathbf{X}_l) &= \min\left\{1,\ (1+5\epsilon)\lambda^+_n(\mathbf{X}_l) - \lambda^+_{m,n}(\mathbf{X}_l) + \sum_{j\in[n]}\mathcal{B}_n(\mathbf{X}_l,\mathbf{X}_j)\right\},\end{aligned}$$

and note that by using the inequality $\min\{1,x+y\}\le x+\min\{1,y\}$, we obtain

$$\begin{aligned}\mathbb{P}\left((C^-_{T^-_r})^c\middle|\,\mathcal{G}^-_{r-1}\right) &\le g^-_{m,n,\epsilon}(\mathbf{X}_{T^-_r}) + 2\epsilon\left|\{T^-_r\}\cup A^-_{r-1}\cup I^-_{r-1}\right| \qquad \text{and}\\ \mathbb{P}\left((C^+_{T^+_r})^c\middle|\,\mathcal{G}^+_{r-1}\right) &\le g^+_{m,n,\epsilon}(\mathbf{X}_{T^+_r}) + 2\epsilon\left|\{T^+_r\}\cup I^-_{\sigma^-_k}\cup A^-_{\sigma^-_k}\cup A^+_{r-1}\cup I^+_{r-1}\right|.\end{aligned}$$

It follows that on the event $\Omega_{m,n}$ we have

$$\begin{aligned}
&\mathbb{P}_i\left(\tau^- \le \sigma_k^-\right) + \mathbb{P}_i\left(\{\tau^- > \sigma_k^-\} \cap \{\tau^+ \le \sigma_k^+\}\right) \\
&\le \sum_{r=1}^k \mathbb{E}_i\left[1(\tau^- > r-1, A_{r-1}^- \neq \varnothing, E_{r-1})\left(g_{m,n,\epsilon}^-(\mathbf{X}_{T_r^-}) + 2\epsilon\left|\{T_r^-\} \cup A_{r-1}^- \cup I_{r-1}^-\right|\right)\right] \\
&\quad + \sum_{r=1}^k \mathbb{E}_i\left[1(\tau^- > \sigma_k^-,\, \tau^+ > r-1, A_{r-1}^+ \neq \varnothing, E_{r-1})\right. \\
&\qquad\qquad \left.\cdot\left(g_{m,n,\epsilon}^+(\mathbf{X}_{T_r^+}) + 2\epsilon\left|\{T_r^+\} \cup I_{\sigma_k^-}^- \cup A_{\sigma_k^-}^- \cup A_{r-1}^+ \cup I_{r-1}^+\right|\right)\right] \\
&\le \sum_{r=1}^k \mathbb{E}_i\left[1(\hat{A}_{r-1}^- \neq \varnothing)\left(g_{m,n,\epsilon}^-(\mathbf{X}_{\hat{T}_r^-}) + 2\epsilon k\right)\right] \\
&\quad + \sum_{r=1}^k \mathbb{E}_i\left[1(\hat{A}_{r-1}^+ \neq \varnothing)\left(g_{m,n,\epsilon}^+(\mathbf{X}_{\hat{T}_r^+}) + 2\epsilon\left(k + \left|\hat{I}_{\hat{\sigma}_k^-}^- \cup \left\{T_{\mathbf{i}} : \mathbf{i} \in \hat{A}_{\hat{\sigma}_k^-}^-\right\}\right|\right)\right)\right] \\
&\le 4\epsilon k^2 + 2\epsilon k \mathbb{E}_i\left[\left|\hat{I}_{\hat{\sigma}_k^-}^-\right| + \left|\hat{A}_{\hat{\sigma}_k^-}^-\right|\right] \\
&\quad + \sum_{r=1}^k \mathbb{E}_i\left[1(\hat{A}_{r-1}^- \neq \varnothing) g_{m,n,\epsilon}^-(\mathbf{X}_{\hat{T}_r^-})\right] + \sum_{r=1}^k \mathbb{E}_i\left[1(\hat{A}_{r-1}^+ \neq \varnothing) g_{m,n,\epsilon}^+(\mathbf{X}_{\hat{T}_r^+})\right],
\end{aligned}$$

where $\hat{T}_r^-$ and $\hat{T}_r^+$ are the *identities* of the rth "active" nodes to be explored in the inbound and outbound multi-type branching processes, respectively, and $\hat{\sigma}_k^\pm = \inf\{t \ge 1 : |\hat{A}_t^\pm| + |\hat{I}_t^\pm| \ge k \text{ or } \hat{A}_t^\pm = \varnothing\}$.

Next, use Lemma 3.4.9 to obtain that for $r \ge 1$,

$$\mathbb{E}_i\left[1(\hat{A}_{r-1}^- \neq \varnothing) g_{m,n,\epsilon}^-(\mathbf{X}_{\hat{T}_r^-})\right] \le \sum_{s=0}^{r-1} \binom{r-1}{s} 2^{r-1-s} (\Gamma_-^{(m,n)})^s g_{m,n,\epsilon}^-(\mathbf{X}_i)$$

and

$$\mathbb{E}_i\left[1(\hat{A}_{r-1}^- \neq \varnothing) g_{m,n,\epsilon}^+(\mathbf{X}_{\hat{T}_r^+})\right] \le \sum_{s=0}^{r-1} \binom{r-1}{s} 2^{r-1-s} (\Gamma_+^{(m,n)})^s g_{m,n,\epsilon}^+(\mathbf{X}_i),$$

where $\Gamma_-^{(m,n)}$ and $\Gamma_+^{(m,n)}$ are the linear integral operators defined in Lemma 3.4.9.

Averaging over all $1 \le i \le n$ and using Lemma 3.4.7 to bound $n^{-1}\sum_{i=1}^n \mathbb{E}_i\left[\left|\hat{I}_{\hat{\sigma}_k^-}^-\right| + \left|\hat{A}_{\hat{\sigma}_k^-}^-\right|\right]$, we obtain

$$\frac{1}{n}\sum_{i=1}^{n}\left(\mathbb{P}_i\left(\tau^- \le \sigma_k^-\right) + \mathbb{P}_i\left(\{\tau^- > \sigma_k^-\} \cap \{\tau^+ \le \sigma_k^+\}\right)\right) 1(\Omega_{m,n})$$
$$\le 4\epsilon k^2 + 2\epsilon k^2 \left(1 + \sup_{\mathbf{x}\in\mathcal{S}} \lambda_-^{(m)}(\mathbf{x})\right)$$
$$+ 1(\Omega_{m,n}) \sum_{r=1}^{k}\sum_{s=0}^{r-1}\binom{r-1}{s} 2^{r-1-s}\left\{\int_{\mathcal{S}} (\Gamma_-^{(m,n)})^s g_{m,n,\epsilon}^-(\mathbf{x})\mu_n(d\mathbf{x}) + \int_{\mathcal{S}} (\Gamma_+^{(m,n)})^s g_{m,n,\epsilon}^+(\mathbf{x})\mu_n(d\mathbf{x})\right\}$$
$$=: H(n,m,k,\epsilon) - 1(\Omega_{m,n}^c).$$

The upper bound for the limit of $H(n,m,k,\epsilon)$ as $n\to\infty$ is given in Lemma 3.4.10. This completes the proof. ■

As a last proof in this section, we use Theorem 3.4.6 to prove Theorem 3.3.4, the result establishing the limiting distribution of the degrees in $G_n(\kappa(1+\varphi_n))$. The latter can also be proven directly using similar arguments as those used in the proof of Theorem 3.4.6, but we choose to do it this way to avoid repetition.

Proof of Theorem 3.3.4. Let

$$D_{n,i}^- = \sum_{j\neq i} Y_{ji} \qquad \text{and} \qquad D_{n,i}^+ = \sum_{j\neq i} Y_{ij}$$

and define

$$Z_{n,i}^- = \sum_{j=1}^{n} Z_{ji} \qquad \text{and} \qquad Z_{n,i}^+ = \sum_{j=1}^{n} \tilde{Z}_{ij},$$

where Z_{ji} is Poisson with mean $r_{ji}^{(m,n)}$ and $\tilde{Z}_{ij}$ is Poisson with mean $\tilde{r}_{ij}^{(m,n)}$. Then,

$$\left(D_{n,\xi}^-, D_{n,\xi}^+\right) = \left(D_{n,\xi}^- - Z_{n,\xi}^-,\ D_{n,\xi}^+ - Z_{n,\xi}^+\right) + \left(Z_{n,\xi}^-,\ Z_{n,\xi}^+\right),$$

where since $\sum_{j=1}^n r_{ji}^{(m,n)} = \lambda_-^{(m)}(\mathbf{X}_i)$ and $\sum_{j=1}^n \tilde{r}_{ij}^{(m,n)} = \lambda_+^{(m)}(\mathbf{X}_i)$, we obtain that

$$\mathbb{P}\left(Z_{n,\xi}^- = k, Z_{n,\xi}^+ = l\right) = \frac{1}{n}\sum_{i=1}^{n} \frac{e^{-\lambda_-^{(m)}(\mathbf{X}_i)}(\lambda_-^{(m)}(\mathbf{X}_i))^k}{k!}\cdot\frac{e^{-\lambda_+^{(m)}(\mathbf{X}_i)}(\lambda_+^{(m)}(\mathbf{X}_i))^l}{l!}$$
$$\xrightarrow{P} \int_{\mathcal{S}} \frac{e^{-\lambda_-^{(m)}(\mathbf{x})}(\lambda_-^{(m)}(\mathbf{x}))^k}{k!}\cdot\frac{e^{-\lambda_+^{(m)}(\mathbf{x})}(\lambda_+^{(m)}(\mathbf{x}))^l}{l!}\,\mu(d\mathbf{x})$$

for any $k,l\ge 0$, as $n\to\infty$ (by the bounded convergence theorem). Moreover, by Theorem 3.4.6,

$$\mathbb{P}\left(|D_{n,\xi}^- - Z_{n,\xi}^-| + |D_{n,\xi}^+ - Z_{n,\xi}^+| > 0\right) = \frac{1}{n}\sum_{i=1}^{n}\mathbb{P}_i\left(\{\tau^-\le\sigma_1^-\}\cup\{\tau^+\le\sigma_1^+\}\right) \le H(n,m,1,\epsilon)$$

for any $0 < \epsilon < 1/2$. Therefore, for $(Z^-_{(m)}, Z^+_{(m)})$ constructed on the same probability space as $\left(Z^-_{n,\xi}, Z^+_{n,\xi}\right)$, with $Z^-_{(m)}$ and $Z^+_{(m)}$ conditionally independent (given $\mathbf{X}$) Poisson random variables with parameters $\lambda^{(m)}_-(\mathbf{X})$ and $\lambda^{(m)}_+(\mathbf{X})$, and $\mathbf{X}$ distributed according to μ, we obtain that

$$\limsup_{n\to\infty} P\left(|D^-_{n,\xi} - Z^-_{(m)}| + |D^+_{n,\xi} - Z^+_{(m)}| > 0\right) \leq \limsup_{n\to\infty} E\left[H(n,m,1,\epsilon) \wedge 1\right] = \hat{H}(m,1,\epsilon),$$

where $\lim_{m\nearrow\infty} \lim_{\epsilon\downarrow 0} \hat{H}(m,1,\epsilon) = 0$ by Lemma 3.4.10. Taking the limit as $\epsilon \downarrow 0$ followed by $m \nearrow \infty$ and noting that $(Z^-_{(m)}, Z^+_{(m)}) \to (Z^-, Z^+)$ a.s., where (Z^-, Z^+) are conditionally independent (given $\mathbf{X}$) Poisson random variables with parameters $\lambda_-(\mathbf{X})$ and $\lambda_+(\mathbf{X})$, gives the weak convergence statement of the theorem.

To obtain the convergence of the expectations note that

$$E[D^-_{n,\xi}] = E[D^+_{n,\xi}] = \frac{1}{n} E\left[\sum_{i=1}^n \sum_{j=1}^n p^{(n)}_{ji}\right] \to \iint_{\mathcal{S}^2} \kappa(\mathbf{x},\mathbf{y})\mu(d\mathbf{x})\mu(d\mathbf{y})$$

as $n \to \infty$ by Assumption 3.3.1(d). Now note that

$$\iint_{\mathcal{S}^2} \kappa(\mathbf{x},\mathbf{y})\mu(d\mathbf{x})\mu(d\mathbf{y}) = E[\lambda_-(\mathbf{X})] = E[\lambda_+(\mathbf{X})] = E[Z^-] = E[Z^+].$$

This completes the proof. ■

3.4.3.2 Size of the largest strongly connected component

This last section of the paper contains the proof of Theorem 3.3.10, the phase transition for the existence of a giant strongly connected component. As mentioned earlier, the idea is to use Theorem 3.4.6 to couple the exploration of the graph $G_n(\kappa(1+\varphi_n))$ starting from a given vertex with a double tree $(\mathcal{T}^-_\mu(\kappa_m), \mathcal{T}^+_\mu(\kappa_m))$ for a kernel κ_m that takes at most a finite number of different values.

Recall from Section 3.3.3 that $(\mathcal{T}^-_\mu(\kappa;\mathbf{x}), \mathcal{T}^+_\mu(\kappa;\mathbf{x}))$ denotes the double multi-type Galton-Watson process having root of type $\mathbf{x} \in \mathcal{S}$, and whose offspring distributions are given by (3.3.1). Let $\rho^{\geq k}_-(\kappa;\mathbf{x})$ (respectively, $\rho^{\geq k}_+(\kappa;\mathbf{x})$) be the probability that the total population of $\mathcal{T}^-_\mu(\kappa;\mathbf{x})$ (respectively, $\mathcal{T}^+_\mu(\kappa;\mathbf{x})$) is at least k. Define also $\rho_-(\kappa;\mathbf{x})$ (respectively, $\rho_+(\kappa;\mathbf{x})$) to be its survival probability, i.e., the probability that its total population is infinite. The averaged joint survival probability is defined as

$$\rho(\kappa) = \int_{\mathcal{S}} \rho_+(\kappa;\mathbf{x})\rho_-(\kappa;\mathbf{x})\mu(d\mathbf{x}).$$

Similarly, for any $k \in \mathbb{N}_-$, we define $\rho^{\geq k}(\kappa) = \int_{\mathcal{S}} \rho^{\geq k}_+(\kappa;\mathbf{x})\rho^{\geq k}_-(\kappa;\mathbf{x})\mu(d\mathbf{x})$.

In addition, we will require from here on that the kernel κ_m be regular finitary (see Definition 3.3.9) and quasi-irreducible (see Definition 3.3.8). The following lemma is taken from [14] and it provides the existence of a sequence of partitions $\{\mathscr{J}_m\}_{m\geq 1}$ of $\mathcal{S}$ over which we can define a sequence of regular finitary kernels.

Lemma 3.4.11 (Lemma 7.1 in [14]) *there exists a sequence of partitions* $\{\mathscr{J}_m : m \geq 1\}$ *of* $\mathcal{S}$*, with* $\mathscr{J}_m = \{\mathcal{J}_1^{(m)}, \ldots, \mathcal{J}_{M_m}^{(m)}\}$*, such that*

i) each $\mathcal{J}_i^{(m)}$ *is measurable and* $\mu(\partial \mathcal{J}_i^{(m)}) = 0$*,*

ii) for each m*,* $\mathscr{J}_{m+1}$ *refines* $\mathscr{J}_m$*, i.e., each* $\mathcal{J}_i^{(m)} = \bigcup_{j \in I_i^{(m)}} \mathcal{J}_j^{(m+1)}$ *for some index set* $I_i^{(m)}$*,*

iii) for a.e. $\mathbf{x} \in \mathcal{S}$*, diam*$(\mathcal{J}_{\vartheta(\mathbf{x})}^{(m)}) \to 0$ *as* $m \to \infty$*, where* $\vartheta(\mathbf{x}) = j$ *if and only if* $\mathbf{x} \in \mathcal{J}_j^{(m)}$*.*

Before we construct the sequence of quasi-irreducible regular finitary kernels that we need, we define for notational convenience the following relation.

Definition 3.4.12 *Let* $\tilde{\kappa}$ *be a kernel on* $\mathcal{S} \times \mathcal{S}$ *and let* $\mathscr{J} = \{\mathcal{J}_1, \ldots, \mathcal{J}_M\}$ *be a finite partition of* $\mathcal{S}$*. Then, we say that set* $A \subseteq \mathcal{S}$ *is* inbound-accessible *(respectively,* outbound-accessible*) from* $\mathbf{x} \in \mathcal{S}$ *with respect to* $(\tilde{\kappa}, \mathscr{J})$*, denoted* $\mathbf{x} \to A$ *(respectively,* $\mathbf{x} \leftarrow A$*), if there exists* $\{u_1, \ldots, u_k\} \subseteq \{1, \ldots, M\}$ *such that:*

i) $\tilde{\kappa}(\mathbf{x}, \mathbf{y}) > 0$ *for all* $\mathbf{y} \in \mathcal{J}_{u_1}$*,*

ii) $\tilde{\kappa} > 0$ *on* $\mathcal{J}_{u_i} \times \mathcal{J}_{u_{i+1}}$ *(respectively,* $\tilde{\kappa} > 0$ *on* $\mathcal{J}_{u_{i+1}} \times \mathcal{J}_{u_i}$*) for all* $1 \leq i < k$*,*

iii) $\mu(\mathcal{J}_{u_i}) > 0$ *for all* $1 \leq i \leq k$*, and*

iv) $\mathcal{J}_{u_k} \subseteq A$*.*

Remark 3.4.13 *Note that if we take* $\mathscr{J}_m = \{\mathcal{J}_1^{(m)}, \ldots, \mathcal{J}_{M_m}^{(m)}\}$ *as constructed in Lemma 3.4.11, and we let* $\tilde{\kappa}_m$ *satisfy* $\tilde{\kappa}_m \leq \tilde{\kappa}_{m+1}$ *a.e., then if* $\mathbf{x} \to A$ *(*$\mathbf{x} \leftarrow A$*) with respect to* $(\tilde{\kappa}_{m_0}, \mathscr{J}_{m_0})$ *for some* $m_0 \geq 1$*, then* $\mathbf{x} \to A$ *(*$\mathbf{x} \leftarrow A$*) with respect to* $(\tilde{\kappa}_m, \mathscr{J}_m)$ *for any* $m \geq m_0$*, since each* $\mathcal{J}_{u_i}^{(m)}$ *in part (iii) of Definition 3.4.12 must contain at least one subset* $\mathcal{J}_t^{(m+1)} \subseteq \mathcal{J}_{u_i}^{(m)}$ *with* $\mu(\mathcal{J}_t^{(m+1)}) > 0$*.*

We now give a result that states that we can always find a sequence of quasi-irreducible regular finitary kernels which converges monotonically to κ and can be used to approximate from below $\kappa(1 + \varphi_n)$. Its proof follows that of Lemma 7.3 in [14], with some variations due to the directed nature of our kernels.

Lemma 3.4.14 *For any continuous kernel* κ *and any* φ_n *satisfying Assumption 3.3.1, there exists a sequence* $\{\tilde{\kappa}_m\}_{m \geq 1}$ *of regular finitary kernels on* $\mathcal{S} \times \mathcal{S}$*, measurable with respect to* $\mathscr{F}$*, with the following properties.*

1. $\tilde{\kappa}_m(\mathbf{x}, \mathbf{y}) \nearrow \kappa(\mathbf{x}, \mathbf{y})$ *in probability as* $m \to \infty$ *for a.e.* $(\mathbf{x}, \mathbf{y}) \in \mathcal{S} \times \mathcal{S}$
2. $\tilde{\kappa}_m(\mathbf{x}, \mathbf{y}) \leq \inf_{n \geq m} \kappa(\mathbf{x}, \mathbf{y})(1 + \varphi_n(\mathbf{x}, \mathbf{y}))$ *for every* $(\mathbf{x}, \mathbf{y}) \in \mathcal{S} \times \mathcal{S}$*.*

3. If κ is quasi-irreducible, then so is κ_m for all large m.

Proof. We may assume that $\kappa > 0$ on a set of positive measure, as otherwise we may take $\kappa_m \equiv 0$ for every m and there is nothing to prove. We will construct the sequence $\{\kappa_m : m \geq 1\}$ in two stages. First, we construct a sequence $\{\tilde{\kappa}_m : m \geq 1\}$ where each $\tilde{\kappa}_m$ is regular finitary and satisfies conditions (a) and (b); then we use this sequence to obtain $\{\kappa_m : m \geq 1\}$ satisfying (c).

To this end, construct the sequence of partitions $\{\mathscr{J}_m\}_{m\geq 1}$ according to Lemma 3.4.11 and define

$$\tilde{\kappa}_m(\mathbf{x},\mathbf{y}) := \inf\left\{\kappa(\mathbf{x}',\mathbf{y}') \wedge \inf_{n\geq m} \kappa(\mathbf{x}',\mathbf{y}')(1+\varphi_n(\mathbf{x}',\mathbf{y}')) : \mathbf{x}' \in \mathcal{J}^{(m)}_{\vartheta(\mathbf{x})},\ \mathbf{y}' \in \mathcal{J}^{(m)}_{\vartheta(\mathbf{y})}\right\}.$$

Note that the properties of $\{\mathscr{J}_m : m \geq 1\}$, and the assumption on φ_n imply that

$$\tilde{\kappa}_m(\mathbf{x},\mathbf{y}) \nearrow \kappa(\mathbf{x},\mathbf{y}) \quad \text{in probability} \quad \text{as } m \to \infty, \quad \text{for a.e. } (\mathbf{x},\mathbf{y}) \in \mathcal{S}\times\mathcal{S}.$$

Moreover, for $n \geq m$ we have that

$$\tilde{\kappa}_m(\mathbf{x},\mathbf{y}) \leq \kappa(\mathbf{x},\mathbf{y})(1+\varphi_n(\mathbf{x},\mathbf{y})) \qquad \text{for all } (\mathbf{x},\mathbf{y}) \in \mathcal{S}\times\mathcal{S}.$$

Hence, $\kappa_m = \tilde{\kappa}_m$ satisfies conditions (a) and (b) in the statement of the lemma.

To prove (c) assume from now on that κ is quasi-irreducible. In fact, without loss of generality we may assume that κ is irreducible, since it suffices to construct κ_m to be quasi-irreducible on the restriction $\mathcal{S}'\times\mathcal{S}'$ where κ is irreducible and then set it to be zero outside of $\mathcal{S}'\times\mathcal{S}'$.

The first step of the proof ensures the existence of a directed cycle $\mathcal{C} \subseteq \mathcal{S}$ for some $m_1 \geq 1$. The second step uses $\mathcal{C}$ to construct a set on which $\tilde{\kappa}_m$ is irreducible. To establish the existence of $\mathcal{C}$, note that if $\tilde{\kappa}_m = 0$ a.e. for all $m \geq 1$, it would imply that $\kappa = 0$ a.e., which would contradict the irreducibility of κ. Therefore, there must exist some $m_0 \geq 1$ and indexes $1 \leq r,s,t \leq M_{m_0}$ such that $\tilde{\kappa}_{m_0} > 0$ on $(\mathcal{J}^{(m_0)}_t \times \mathcal{J}^{(m_0)}_r)$ and on $(\mathcal{J}^{(m_0)}_r \times \mathcal{J}^{(m_0)}_s)$, with $\mu(\mathcal{J}^{(m_0)}_t)\mu(\mathcal{J}^{(m_0)}_r)\mu(\mathcal{J}^{(m_0)}_s) > 0$.

Claim: for any set $A \subseteq \mathcal{S}$ for which there exists a set $D \subseteq \mathcal{S}$ such that $\mu(D) > 0$ and $\tilde{\kappa}_m > 0$ on $D \times A$ (respectively, $A \times D$), the sequence of sets $\{B_m(A)\}_{m\geq 1}$ (respectively, $\{\tilde{B}_m(A)\}_{m\geq 1}$) defined according to $B_m(A) = \{\mathbf{x} \in \mathcal{S} : \mathbf{x} \to A \text{ w.r.t. } (\tilde{\kappa}_m, \mathscr{J}_m)\}$ (respectively, $\tilde{B}_m(A) = \{\mathbf{x} \in \mathcal{S} : \mathbf{x} \leftarrow A \text{ w.r.t. } (\tilde{\kappa}_m, \mathscr{J}_m)\}$) satisfy: 1) $B_m(A) \subseteq B_{m+1}(A)$ (respectively, $\tilde{B}_m(A) \subseteq \tilde{B}_{m+1}(A)$), and 2) $\mu\left(\bigcup_{m=1}^\infty B_m(A)\right) = 1$ $\Big($respectively, $\mu\Big(\bigcup_{m=1}^\infty \tilde{B}_m(A)\Big) = 1\Big)$.

To prove the claim note that Remark 3.4.13 implies (1). To see that (2) holds, let $B(A) = \bigcup_{m=1}^\infty B_m(A)$ and note that from the definition of $B(A)$ we have $\kappa = 0$ a.e. on $B(A)^c \times B(A)$, and the irreducibility of κ implies that either $\mu(B(A)^c) = 0$ or $\mu(B(A)) = 0$; since $\mu(B(A)) \geq \mu(D) > 0$, it must be that $\mu(B(A)^c) = 0$, which implies that $\mu(B(A)) = 1$. The symmetric arguments yield the claim for $\{\tilde{B}_m(A)\}$.

Now apply the inbound part of the claim to $A = \mathcal{J}_r^{(m_0)}$ and $D = \mathcal{J}_t^{(m_0)}$ to obtain that there exists $m_1 \geq m_0$ such that $\mu(B_{m_1}(\mathcal{J}_r^{(m_0)}) \cap \mathcal{J}_s^{(m_0)}) > 0$, which in turn implies there exists a set $\mathcal{J}_{s'}^{(m_1)} \subseteq \mathcal{J}_s^{(m_0)}$ such that $\mu(\mathcal{J}_{s'}^{(m_1)}) > 0$ and $\mathbf{x} \to \mathcal{J}_r^{(m_0)}$ for all $\mathbf{x} \in \mathcal{J}_{s'}^{(m_1)}$. In other words, there exist sets $\{\mathcal{J}_{u_0}^{(m_1)}, \ldots, \mathcal{J}_{u_k}^{(m_1)}\}$ satisfying $\mu(\mathcal{J}_{u_i}^{(m_1)}) > 0$ for all $0 \leq i \leq k$, $\mathcal{J}_{u_0}^{(m_1)} = \mathcal{J}_{s'}^{(m_1)}$, $\mathcal{J}_{u_k}^{(m_1)} \subseteq \mathcal{J}_r^{(m_0)}$, and $\tilde{\kappa}_{m_1} > 0$ on $\mathcal{J}_{u_i}^{(m_1)} \times \mathcal{J}_{u_{i+1}}^{(m_1)}$ for all $0 \leq i < k$. Since $0 < \tilde{\kappa}_{m_0} \leq \tilde{\kappa}_{m_1}$ on $\mathcal{J}_{u_k}^{(m_1)} \times \mathcal{J}_{u_0}^{(m_1)}$ by construction, we have that the set $\mathcal{C} = \bigcup_{i=0}^k \mathcal{J}_{u_i}^{(m_1)}$ defines a directed cycle.

Next, construct the sequences $\{B_m(\mathcal{C})\}_{m\geq 1}$ and $\{\tilde{B}_m(\mathcal{C})\}_{m\geq 1}$ according to the claim, and define

$$\kappa_m(\mathbf{x}, \mathbf{y}) = \tilde{\kappa}_m(\mathbf{x}, \mathbf{y}) 1(\mathbf{x} \in (B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C})), \, \mathbf{y} \in (B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C}))).$$

Note that $\kappa_m \nearrow \kappa$ in probability as $m \to \infty$ since $\tilde{\kappa}_m \nearrow \kappa$ in probability and

$$\mu\left(\bigcup_{m=1}^{\infty} (B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C})) \right) \geq 1 - \mu\left(\bigcap_{m=1}^{\infty} B_m(\mathcal{C})^c \right) - \mu\left(\bigcap_{m=1}^{\infty} \tilde{B}_m(\mathcal{C})^c \right) = 1.$$

It remains to show that κ_m restricted to $(B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C})) \times (B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C}))$ is irreducible. To see this, let $A \subseteq (B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C}))$ and suppose $\kappa_m = 0$ on $A \times (A^c \cap B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C}))$. Note that since $\tilde{\kappa}_{m_1} > 0$ on each $\mathcal{J}_{u_i}^{(m_1)} \times \mathcal{J}_{u_{i+1}}^{(m_1)}$, then it must be that either $\mathcal{C} \subseteq A$ or $\mathcal{C} \subseteq A^c$. Suppose that it is the former, and note that for any $\mathbf{x} \in A^c \cap B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C})$ there exist indexes $\{v_1, \ldots, v_l\}$ and $\{w_1, \ldots, w_j\}$ such that

$$\tilde{\kappa}_{m_1} > 0 \text{ on } \mathcal{J}_{v_i}^{(m_1)} \times \mathcal{J}_{v_{i+1}}^{(m_1)}, \, 0 \leq i \leq l, \, \mu(\mathcal{J}_{v_i}^{(m_1)}) > 0, \, 1 \leq i \leq l, \, \mathcal{J}_{v_l}^{(m_1)} \subseteq \mathcal{C},$$

and

$$\tilde{\kappa}_{m_1} > 0 \text{ on } \mathcal{J}_{w_{i+1}}^{(m_1)} \times \mathcal{J}_{w_i}^{(m_1)}, \, 0 \leq i \leq j, \, \mu(\mathcal{J}_{w_i}^{(m_1)}) > 0, \, 1 \leq i \leq j, \, \mathcal{J}_{w_j}^{(m_1)} \subseteq \mathcal{C},$$

where $\mathcal{J}_{v_0}^{(m_1)} = \mathcal{J}_{w_0}^{(m_1)} = \mathcal{J}_{\vartheta(\mathbf{x})}^{(m_1)}$. Moreover, $\mu(\mathcal{J}_{\vartheta(\mathbf{x})}^{(m_1)}) > 0$ would imply that $\mathcal{J}_{v_i}^{(m_1)} \subseteq B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C})$ for all $1 \leq i \leq l$ and $\mathcal{J}_{w_h}^{(m_1)} \subseteq B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C})$ for all $1 \leq h \leq j$, since they would all lie on a directed cycle of positive measure, but this contradicts our assumption that $\tilde{\kappa}_{m_1} = 0$ on $A \times A^c \cap B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C})$. Hence, it must be that $\mu(\mathcal{J}_{\vartheta(\mathbf{x})}^{(m_1)}) = 0$ for all $\mathbf{x} \in A^c \cap B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C})$, and therefore, $\mu(A^c \cap B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C})) = 0$. The same argument gives that if $\mathcal{C} \subseteq A^c \cap B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C})$ then $\mu(A) = 0$. We conclude that κ_m restricted to $(B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C})) \times (B_m(\mathcal{C}) \cap \tilde{B}_m(\mathcal{C}))$ is irreducible. This completes the proof. ■

The following lemma establishes the relationships between $\rho(\kappa_m)$, $\rho^{\geq k}(\kappa_m)$, $\rho^{\geq k}(\kappa)$, and $\rho(\kappa)$.

Lemma 3.4.15 *Let $\{\kappa_m\}_{m\geq 1}$ be a sequence of kernels on $(\mathcal{S}, \mu)$ increasing a.e. to κ. Then, the following limits hold:*

1. *$\rho^{\geq k}(\kappa; \mathbf{x}) \searrow \rho(\kappa; \mathbf{x})$ a.e. $\mathbf{x}$ and $\rho^{\geq k}(\kappa) \searrow \rho(\kappa)$ as $k \to \infty$.*
2. *For every $k \geq 1$, $\rho^{\geq k}(\kappa_m; \mathbf{x}) \nearrow \rho^{\geq k}(\kappa; \mathbf{x})$ for a.e. $\mathbf{x}$ and $\rho^{\geq k}(\kappa_m) \nearrow \rho^{\geq k}(\kappa)$ as $m \to \infty$.*

3. $\rho(\kappa_m; \mathbf{x}) \nearrow \rho(\kappa; \mathbf{x})$ *for a.e.* $\mathbf{x}$ *and* $\rho(\kappa_m) \nearrow \rho(\kappa)$ *as* $m \to \infty$.

Proof. By Lemma 9.5 in [14], we have that $\rho_+^{\geq k}(\kappa; \mathbf{x}) \searrow \rho_+(\kappa; \mathbf{x})$ and $\rho_-^{\geq k}(\kappa; \mathbf{x}) \searrow \rho_-(\kappa; \mathbf{x})$ as $k \to \infty$ for a.e. $\mathbf{x}$. Then, by the monotone convergence theorem, we have

$$\begin{aligned}\lim_{k\to\infty} \rho^{\geq k}(\kappa) &= \lim_{k\to\infty} \int_{\mathcal{S}} \rho_+^{\geq k}(\kappa; \mathbf{s})\rho_-^{\geq k}(\kappa; \mathbf{s})\mu(d\mathbf{s}) \\ &= \int_{\mathcal{S}} \lim_{k\to\infty} \rho_+^{\geq k}(\kappa; \mathbf{s})\rho_-^{\geq k}(\kappa; \mathbf{s})\mu(d\mathbf{s}) \\ &= \int_{\mathcal{S}} \rho_+(\kappa; \mathbf{s})\rho_-(\kappa; \mathbf{s})\mu(d\mathbf{s}) = \rho(\kappa),\end{aligned}$$

which establishes (a).

By Theorem 6.5(i) in [14] we have that for any fixed $k \geq 1$, $\rho_+^{\geq k}(\kappa_m; \mathbf{x}) \nearrow \rho_+^{\geq k}(\kappa; \mathbf{x})$ and $\rho_-^{\geq k}(\kappa_m; \mathbf{x}) \nearrow \rho_-^{\geq k}(\kappa; \mathbf{x})$ as $m \to \infty$ for a.e. $\mathbf{x}$, which together with monotone convergence as above implies (b).

Part (c) follows from part (a) applied to the kernel κ_m, followed by part (b), to obtain that

$$\lim_{m\to\infty} \rho(\kappa_m; \mathbf{x}) = \lim_{m\to\infty} \lim_{k\to\infty} \rho^{\geq k}(\kappa_m; \mathbf{x}) = \lim_{k\to\infty} \lim_{m\to\infty} \rho^{\geq k}(\kappa_m; \mathbf{x}) = \lim_{k\to\infty} \rho^{\geq k}(\kappa; \mathbf{x}) = \rho(\kappa; \mathbf{x})$$

for a.e. $\mathbf{x}$. Then use monotone convergence as above. ■

Recall the definition of the operators T_κ^- and T_κ^+ given in Section 3.3.3, as well as of their spectral radii $r(T_\kappa^-)$ and $r(T_\kappa^+)$. The strict positivity of $\rho(\kappa)$, which ensures the existence of a giant strongly connected component, is characterized below. As a preliminary result, we establish the phase transition for regular finitary, quasi-irreducible kernels first.

Proposition 3.4.16 *Suppose that* $\tilde{\kappa}$ *is a regular finitary, quasi-irreducible, kernel on the type-space* $\mathcal{S}$ *with respect to measure* μ. *Then,* $r(T_{\tilde{\kappa}}^-) = r(T_{\tilde{\kappa}}^+)$ *and we have that* $\rho(\tilde{\kappa}) > 0$ *if and only if* $r(T_{\tilde{\kappa}}^-) > 1$. *Moreover, there exist nonnegative, non-zero eigenfunctions* f_- *and* f_+, *such that* $T_{\tilde{\kappa}}^- f^- = r(T_{\tilde{\kappa}}^-)f_-$ *and* $T_{\tilde{\kappa}}^+ f^+ = r(T_{\tilde{\kappa}}^+)f_+$, *and they are the only (up to multiplicative constants and sets of measure zero) nonnegative, non-zero eigenfunctions of* $T_{\tilde{\kappa}}^-$ *and* $T_{\tilde{\kappa}}^+$, *respectively.*

Proof. Since $\tilde{\kappa}$ is quasi-irreducible, there exists $\mathcal{S}^* \subseteq \mathcal{S}$ such that $\tilde{\kappa}$ restricted to $\mathcal{S}^*$ is irreducible and $\mu(\mathcal{S}^*) > 0$. Also, since $\tilde{\kappa}$ is regular finitary, there exists a finite partition $\{\mathcal{J}_i : 1 \leq i \leq M\}$ such that $\tilde{\kappa}$ is constant on $\mathcal{J}_i \times \mathcal{J}_j$. Next, define

$$\mathcal{S}' = \bigcup_{i=1}^{M} \{\mathcal{J}_i \cap \mathcal{S}^* : \mu(\mathcal{J}_i \cap \mathcal{S}^*) > 0\},$$

and define the kernel $\kappa'(\mathbf{x}, \mathbf{y}) = \mu(\mathcal{S}')\tilde{\kappa}(\mathbf{x}, \mathbf{y})$ for $\mathbf{x}, \mathbf{y} \in \mathcal{S}'$. Note that κ' is regular finitary and irreducible on $\mathcal{S}'$ and $\mu(\mathcal{S}') = \mu(\mathcal{S}^*)$. Moreover, if we let $\mu'(A) = \mu(A)/\mu(\mathcal{S}')$ for $A \subseteq \mathcal{S}'$, and let $\{\mathcal{J}'_i : 1 \leq i \leq M'\}$ denote the partition of $\mathcal{S}'$ such that κ' is constant on $\mathcal{J}'_i \times \mathcal{J}'_j$, then $\mu'(\mathcal{J}'_i) > 0$ for all $1 \leq i \leq M'$.

Next, consider the double tree $(\mathcal{T}^-_{\mu'}(\kappa'), \mathcal{T}^+_{\mu'}(\kappa'))$ on the type-space $\mathcal{S}'$ with respect to measure μ'. Note that each of these trees can be thought of as a multi-type branching process with M' types (one associated to each of the $\mathcal{J}'_i$) each having positive probability. We will show that:

1. the survival probability $\rho(\tilde{\kappa}) = \mu(\mathcal{S}')\rho'(\kappa')$, where

$$\rho'(\kappa') = \int_{\mathcal{S}'} \rho'_-(\kappa';\mathbf{x})\rho'_+(\kappa';\mathbf{x})\mu'(d\mathbf{x}),$$

   and $\rho'_-(\kappa';\mathbf{x})$, $\rho'_+(\kappa';\mathbf{x})$ are the survival probabilities of the trees $\mathcal{T}^-_{\mu'}(\kappa';\mathbf{x})$ and $\mathcal{T}^+_{\mu'}(\kappa';\mathbf{x})$, respectively; and

2. the spectral radii of the operators $T^{\pm}_{\tilde{\kappa}}$ on $\mathcal{S}$ and $T^{\pm}_{\kappa'}$ on $\mathcal{S}'$ are the same.

To prove (a), note that since types $\mathbf{x} \in (\mathcal{S}^*)^c$ are isolated (since $\tilde{\kappa}(\mathbf{x},\mathbf{y}) = 0$ for $\mathbf{x} \in (\mathcal{S}')^c$ or $\mathbf{y} \in (\mathcal{S}')^c$) and $\mathcal{S}^* \cap (\mathcal{S}')^c$ has measure zero, then they do not contribute to the survival probabilities of $\mathcal{T}^-_{\mu}(\tilde{\kappa})$ and $\mathcal{T}^+_{\mu}(\tilde{\kappa})$, which implies that

$$\rho(\tilde{\kappa}) = \int_{\mathcal{S}} \rho_+(\tilde{\kappa};\mathbf{x})\rho_-(\tilde{\kappa};\mathbf{x})\mu(d\mathbf{x}) = \mu(\mathcal{S}') \int_{\mathcal{S}'} \rho_+(\tilde{\kappa};\mathbf{x})\rho_-(\tilde{\kappa};\mathbf{x})\mu'(d\mathbf{x}).$$

Now note that the trees $\mathcal{T}^{\pm}_{\mu}(\tilde{\kappa})$ and $\mathcal{T}^{\pm}_{\mu'}(\kappa')$ have the same law when their roots belong to $\mathcal{S}'$ since the number of offspring of type $\mathbf{y} \in \mathcal{S}'$ that an individual of type $\mathbf{x} \in \mathcal{S}'$ on the tree $\mathcal{T}^-_{\mu'}(\kappa')$ has, is Poisson distributed with mean

$$\int_{\mathcal{S}'} \kappa'(\mathbf{y},\mathbf{x})\mu'(d\mathbf{x}) = \int_{\mathcal{S}'} \mu(\mathcal{S}')\tilde{\kappa}(\mathbf{y},\mathbf{x})\mu(d\mathbf{x})/\mu(\mathcal{S}') = \int_{\mathcal{S}} \tilde{\kappa}(\mathbf{y},\mathbf{x})\mu(d\mathbf{x}),$$

which is equal to the corresponding distribution in $\mathcal{T}_{\mu}(\tilde{\kappa})$. The same argument yields the result for $\mathcal{T}^+_{\mu}(\tilde{\kappa})$ and $\mathcal{T}^+_{\mu'}(\kappa')$. Hence, we have that $\rho_{\pm}(\tilde{\kappa};\mathbf{x}) = \rho_{\pm}(\kappa';\mathbf{x})$ for $\mathbf{x} \in \mathcal{S}'$, and therefore,

$$\rho(\tilde{\kappa}) = \mu(\mathcal{S}')\rho'(\kappa').$$

To establish (b), note that if $f'_{\pm}$ is the nonnegative eigenfunction associated to $r(T^{\pm}_{\kappa'})$ on $\mathcal{S}'$, then $f_{\pm}(\mathbf{x}) = f'_{\pm}(\mathbf{x})1(\mathbf{x} \in \mathcal{S}')$ satisfies

$$(T^-_{\tilde{\kappa}} f_-)(\mathbf{x}) = \int_{\mathcal{S}} \tilde{\kappa}(\mathbf{y},\mathbf{x})f_-(\mathbf{y})\mu(d\mathbf{y}) = \int_{\mathcal{S}'} \kappa'(\mathbf{y},\mathbf{x})f'_-(\mathbf{y})\mu'(d\mathbf{y}) = r(T^-_{\kappa'})f'_-(\mathbf{x}) = r(T^-_{\kappa'})f_-(\mathbf{x})$$

for $\mathbf{x} \in \mathcal{S}'$, while for $\mathbf{x} \in (\mathcal{S}')^c$ we have $(T^-_{\tilde{\kappa}} f_-)(\mathbf{x}) = 0$ since $\tilde{\kappa}(\mathbf{y},\mathbf{x}) = 0$ for all $\mathbf{y} \in \mathcal{S}$. Therefore, $r(T^-_{\kappa'})$ is an eigenvalue of $T^-_{\tilde{\kappa}}$, which implies that $r(T^-_{\kappa'}) \leq r(T^-_{\tilde{\kappa}})$; similarly, $r(T^+_{\kappa'})$ is an eigenvalue of $T^+_{\tilde{\kappa}}$ and $r(T^+_{\kappa'}) \leq r(T^+_{\tilde{\kappa}})$. For the opposite inequality, suppose $f_{\pm}$ is a nonnegative eigenvector associated to $r(T^{\pm}_{\tilde{\kappa}})$ and set $f'_{\pm}$ to be its restriction to $\mathcal{S}'$. Then note that for $\mathbf{x} \in \mathcal{S}'$,

$$(T^{-}_{\kappa'} f'_{-})(\mathbf{x}) = \int_{\mathcal{S}'} \kappa'(\mathbf{y},\mathbf{x}) f'_{-}(\mathbf{y}) \mu'(d\mathbf{y}) = \int_{\mathcal{S}} \tilde{\kappa}(\mathbf{y},\mathbf{x}) f_{-}(\mathbf{y}) \mu(d\mathbf{y}) = r(T^{-}_{\tilde{\kappa}}) f_{-}(\mathbf{x}) = r(T^{-}_{\tilde{\kappa}}) f'_{-}(\mathbf{x}),$$

and therefore, $r(T^{-}_{\tilde{\kappa}})$ is an eigenvalue of $T^{-}_{\kappa'}$ and therefore $r(T^{-}_{\tilde{\kappa}}) \leq r(T^{-}_{\kappa'})$. Similarly, $r(T^{+}_{\tilde{\kappa}}) \leq r(T^{+}_{\kappa'})$. We conclude that

$$r(T^{\pm}_{\tilde{\kappa}}) = r(T^{\pm}_{\kappa'}).$$

To see that $r(T^{-}_{\kappa'}) = r(T^{+}_{\kappa'})$ we first point out that $\mathcal{T}^{-}_{\mu'}(\kappa')$ and $\mathcal{T}^{+}_{\mu'}(\kappa')$ can be thought of as irreducible multi-type Galton-Watson processes with a finite number of types and mean progeny matrices $\mathbf{M}^{-} = (m^{-}_{ij})$ and $\mathbf{M}^{+} = (m^{+}_{ij})$, respectively, where $m^{-}_{ij} = c_{ji}\mu'(\mathcal{J}'_j)$, $m^{+}_{ij} = c_{ij}\mu'(\mathcal{J}'_j)$, and $\kappa'(\mathbf{x},\mathbf{y}) = \sum_{i=1}^{M'} \sum_{j=1}^{M'} c_{ij} 1(\mathbf{x} \in \mathcal{J}'_i, \mathbf{y} \in \mathcal{J}'_j)$. Moreover, the operators $T^{-}_{\kappa'}$ and $T^{+}_{\kappa'}$ satisfy

$$T^{\pm}_{\kappa'} f = \mathbf{M}^{\pm}\mathbf{v} \qquad \text{for } \mathbf{v} = (v_1, \dots, v_{M'})^T \in \mathbb{R}^{M'} \text{ and } f(\mathbf{x}) = v_i 1(\mathbf{x} \in \mathcal{J}'_i),\ \mathbf{x} \in \mathcal{S}'.$$

That $\mathbf{M}^{-}$ and $\mathbf{M}^{+}$ have the same spectral radius follows from noting that $\mathbf{M}^{-} = \mathbf{CD}$ and $\mathbf{M}^{+} = \mathbf{C}^T\mathbf{D} = (\mathbf{DC})^T$ for $\mathbf{D} = \mathrm{diag}(\mu'(\mathcal{J}'_1), \dots, \mu'(\mathcal{J}'_{M'}))$ and $\mathbf{C} = (c_{ij})$, which implies that the eigenvalues of $\mathbf{M}^{+}$ are the complex conjugates of those of $\mathbf{DC}$, which in turn are the same as those of $\mathbf{CD}$.

The if and only if statement for the survival probabilities now follows from Theorem 8 in [8] (see also Theorems 2.1 and 2.2 in Chapter 2 of [79]), which states that

$$\rho'_{\pm}(\kappa'; \mathbf{x}) > 0 \text{ for all } \mathbf{x} \in \mathcal{S}' \qquad \text{if and only if} \qquad r(\mathbf{M}^{\pm}) > 1,$$

where $r(\mathbf{M}^{\pm}) = r(T^{\pm}_{\kappa'})$ is the spectral radius of $\mathbf{M}^{\pm}$.

The existence of the eigenfunctions f_{-} and f_{+} on $\mathcal{S}$ follows from the Perron-Frobenius theorem (see Theorem 1.5 in [104]), which guarantees the existence of strictly positive eigenfunctions f'_{-} and f'_{+} on $\mathcal{S}'$ such that $T^{\pm}_{\kappa'} f'_{\pm} = r(T^{\pm}_{\kappa'}) f'_{\pm}$, by setting $f_{\pm}(\mathbf{x}) = f'_{\pm}(\mathbf{x}) 1(\mathbf{x} \in \mathcal{S}')$. Moreover, f'_{-} and f'_{+} are the only (up to multiplicative constants) nonnegative, non-zero eigenfunctions of the operators $T^{-}_{\kappa'}$ and $T^{+}_{\kappa'}$, respectively. To see that the nonnegative eigenfunctions f_{-} and f_{+} are also unique (up to multiplicative constants and sets of measure zero) note that any other nonnegative eigenfunction g_{-} of $T^{-}_{\tilde{\kappa}}$ associated to a positive eigenvalue λ would have to satisfy

$$(T^{-}_{\tilde{\kappa}} g_{-})(\mathbf{x}) = \int_{\mathcal{S}} \tilde{\kappa}(\mathbf{y},\mathbf{x}) g_{-}(\mathbf{y}) \mu(d\mathbf{y}) = 0 \qquad \text{for } \mathbf{x} \in (\mathcal{S}^*)^c,$$

since $\tilde{\kappa}(\mathbf{x},\mathbf{y}) = 0$ for $\mathbf{x} \in (\mathcal{S}^*)^c$, and

$$(T^{-}_{\tilde{\kappa}} g_{-})(\mathbf{x}) = \int_{\mathcal{S}'} \kappa'(\mathbf{y},\mathbf{x}) g_{-}(\mathbf{y}) \mu'(d\mathbf{y}) = \lambda g_{-}(\mathbf{x}) \qquad \text{for } \mathbf{x} \in \mathcal{S}'$$

which would imply λ is a positive eigenvalue of $T^-_{\kappa'}$ with a nonnegative, non-zero, eigenfunction. The uniqueness of f'_- then gives that $g_-(\mathbf{x}) = \alpha f'_-(\mathbf{x})$ for $\mathbf{x} \in \mathcal{S}'$ for some constant $\alpha > 0$. Finally, since $\mu(\mathcal{S}^* \cap (\mathcal{S}')^c) = 0$, we conclude that $g_-(\mathbf{x}) = \alpha f_-(\mathbf{x})$ a.e. The same arguments give that any other nonnegative eigenfunction g_+ of $T^+_{\tilde{\kappa}}$ would have to satisfy $g_+(\mathbf{x}) = \beta f_+(\mathbf{x})$ a.e. This completes the proof. ■

We now use the regular finitary and quasi-irreducible case to establish the result for general irreducible kernels. As pointed out in Remark 3.3.12, the result does not provide a full if and only if condition for the strict positivity of $\rho(\kappa)$, since when the operators T^-_κ and T^+_κ are unbounded we cannot guarantee the continuity of the spectral radii of the sequence of operators $T^-_{\kappa_m}$ and $T^+_{\kappa_m}$.

Lemma 3.4.17 *Suppose that κ is irreducible on the type-space $\mathcal{S}$ with respect to measure μ. Then, if $\rho(\kappa) > 0$ we have $r(T^-_\kappa) > 1$ and $r(T^+_\kappa) > 1$. Moreover, if there exists a regular finitary quasi-irreducible kernel $\tilde{\kappa}$ such that $\tilde{\kappa} \leq \kappa$ a.e. and $r(T^-_{\tilde{\kappa}}) > 1$ (equivalently, $r(T^+_{\tilde{\kappa}}) > 1$), then $\rho(\kappa) > 0$.*

Proof. Suppose first that $\rho(\kappa) > 0$. Now use Lemma 3.4.14 and Lemma 3.4.15 to obtain that $\rho(\kappa_m) > 0$ for some quasi-irreducible, regular finitary, kernel κ_m such that $\kappa_m(\mathbf{x}, \mathbf{y}) \leq \kappa(\mathbf{x}, \mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in \mathcal{S}$. By Proposition 3.4.16 we have that the spectral radii of the operators $T^-_{\kappa_m}$ and $T^+_{\kappa_m}$ satisfy $r(T^-_{\kappa_m}) = r(T^+_{\kappa_m}) > 1$. By monotonicity of the spectral radius, we conclude that $r(T^-_\kappa) \geq r(T^-_{\kappa_m}) > 1$ and $r(T^+_\kappa) \geq r(T^+_{\kappa_m}) > 1$.

For the converse, note that if $\tilde{\kappa} \leq \kappa$ a.e. and $r(T^-_{\tilde{\kappa}}) > 1$, then by Proposition 3.4.16 we have that $\rho(\tilde{\kappa}) > 0$. Since $\rho(\tilde{\kappa}) \leq \rho(\kappa)$, the result follows. ■

The last preliminary result before proving Theorem 3.3.10 provides the key estimates obtained through Theorem 3.4.6, since it relates the indicator random variables for each vertex i to have in-component and out-component of size at least k with the corresponding probabilities in the double-tree $(\mathcal{T}^-_\mu(\kappa_m; \mathbf{X}_i), \mathcal{T}^+_\mu(\kappa_m; \mathbf{X}_i))$.

Proposition 3.4.18 *For any $k \geq 1$ and $i \in [n]$, define $\chi^{\geq k}_{n,i}$ to be the indicator function of the event that vertex i has in-component and out-component both of size at least k in the graph $G_n(\kappa(1 + \varphi_n))$. Then, for any $0 < \epsilon < 1/2$, we have*

$$\left| \frac{1}{n} \sum_{i=1}^n \mathbb{E}\left[\chi^{\geq k}_{n,i}\right] - \frac{1}{n} \sum_{i=1}^n \rho^{\geq k}_-(\kappa_m; \mathbf{X}_i) \rho^{\geq k}_+(\kappa_m; \mathbf{X}_i) \right| \leq H(n, m, k, \epsilon),$$

$$\frac{1}{n^2} \sum_{i=1}^n \sum_{j \neq i} \mathbb{E}\left[\left(\chi^{\geq k}_{n,i} - \mathbb{E}\left[\chi^{\geq k}_{n,i}\right]\right)\left(\chi^{\geq k}_{n,j} - \mathbb{E}\left[\chi^{\geq k}_{n,j}\right]\right)\right] \leq K(n, m, k) + 3H(n, m, k, \epsilon),$$

where

$$K(n, m, k) := \frac{4(k+1)\log n}{n} \sup_{\mathbf{x}, \mathbf{y} \in \mathcal{S}} \kappa_m(\mathbf{x}, \mathbf{y}) + \frac{k}{\log n}\left(2 + \sup_{\mathbf{x} \in \mathcal{S}} \lambda^{(m)}_-(\mathbf{x}) + \sup_{\mathbf{x} \in \mathcal{S}} \lambda^{(m)}_+(\mathbf{x})\right),$$

and $H(n, m, k, \epsilon)$ is defined in Theorem 3.4.6.

Proof. To derive the first bound construct a coupling between the graph exploration processes of the in-component and out-component of vertex i and the double tree $(\mathcal{T}_\mu^-(\kappa_m;\mathbf{X}_i), \mathcal{T}_\mu^+(\kappa_m;\mathbf{X}_i))$, as described in Section 3.4.3.1. Define τ^- and τ^+ to be the steps in the construction when the coupling breaks on the inbound, respectively outbound, sides, and let $\sigma_k^- = \inf\{t \geq 1 : |A_t^-| + |I_t^-| \geq k \text{ or } A_t^- = \varnothing\}$ and $\sigma_k^+ = \inf\{t \geq 1 : |A_t^+| + |I_t^+| \geq k \text{ or } A_t^+ = \varnothing\}$. Note that at time $\sigma_k^- \vee \sigma_k^+$ it is possible to determine whether both the in-component and out-component of vertex i have at least k vertices or not. To simplify the notation, let $\rho^{\geq k}(\kappa_m;\mathbf{x}) = \rho_-^{\geq k}(\kappa_m;\mathbf{x})\rho_+^{\geq k}(\kappa_m;\mathbf{x})$.

$$
\begin{aligned}
\frac{1}{n}\sum_{i=1}^n \mathbb{E}\left[\chi_{n,i}^{\geq k}\right] &= \frac{1}{n}\sum_{i=1}^n \mathbb{P}\left(\chi_{n,i}^{\geq k} = 1\right) \\
&\leq \frac{1}{n}\sum_{i=1}^n \mathbb{P}_i\left(\chi_{n,i}^{\geq k} = 1,\, \tau^- \geq \sigma_k^-, \tau^+ \geq \sigma_k^+\right) + \frac{1}{n}\sum_{i=1}^n \mathbb{P}_i\left(\{\tau^- < \sigma_k^-\} \cup \{\tau^+ < \sigma_k^+\}\right) \\
&\leq \frac{1}{n}\sum_{i=1}^n \mathbb{P}\left(\text{both } \mathcal{T}_\mu^-(\kappa_m;\mathbf{X}_i) \text{ and } \mathcal{T}_\mu^+(\kappa_m;\mathbf{X}_i) \text{ have at least } k \text{ nodes}\right) + H(n,m,k,\epsilon) \\
&= \frac{1}{n}\sum_{i=1}^n \rho^{\geq k}(\kappa_m;\mathbf{X}_i) + H(n,m,k,\epsilon),
\end{aligned}
$$

where we used Theorem 3.4.6 to obtain that $n^{-1}\sum_{i=1}^n \mathbb{P}_i\left(\{\tau^- < \sigma_k^-\} \cup \{\tau^+ < \sigma_k^+\}\right) \leq H(n,m,k,\epsilon)$.

The other direction follows because

$$
\begin{aligned}
\frac{1}{n}\sum_{i=1}^n \mathbb{E}\left[\chi_{n,i}^{\geq k}\right] &\geq \frac{1}{n}\sum_{i=1}^n \mathbb{P}_i\left(\chi_{n,i}^{\geq k} = 1,\, \tau^- \geq \sigma_k^-, \tau^+ \geq \sigma_k^+\right) \\
&\geq \frac{1}{n}\sum_{i=1}^n \mathbb{P}\left(\text{both } \mathcal{T}_\mu^-(\kappa_m;\mathbf{X}_i) \text{ and } \mathcal{T}_\mu^+(\kappa_m;\mathbf{X}_i) \text{ have at least } k \text{ nodes}\right) \\
&\quad - \frac{1}{n}\sum_{i=1}^n \mathbb{P}_i\left(\{\tau^- < \sigma_k^-\} \cup \{\tau^+ < \sigma_k^+\}\right) \\
&\geq \frac{1}{n}\sum_{i=1}^n \rho^{\geq k}(\kappa_m;\mathbf{X}_i) - H(n,m,k,\epsilon).
\end{aligned}
$$

For the second inequality, first note that

$$
\begin{aligned}
&\frac{1}{n^2}\sum_{i=1}^n \sum_{j\neq i} \mathbb{E}\left[\left(\chi_{n,i}^{\geq k} - \mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\right)\left(\chi_{n,j}^{\geq k} - \mathbb{E}\left[\chi_{n,j}^{\geq k}\right]\right)\right] \\
&= \frac{1}{n^2}\sum_{i=1}^n \sum_{j\neq i} \mathbb{E}\left[\chi_{n,i}^{\geq k}\chi_{n,j}^{\geq k}\right] - \frac{1}{n^2}\sum_{i=1}^n \sum_{j\neq i} \mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\mathbb{E}\left[\chi_{n,j}^{\geq k}\right].
\end{aligned}
$$

To estimate $\mathbb{E}\left[\chi_{n,i}^{\geq k}\chi_{n,j}^{\geq k}\right]$ we will assume that we first explore the inbound and outbound neighborhood of vertex i up to the time both its in-component and out-component have at least k vertices or there are no more vertices to explore, i.e., we will explore the in-component of vertex i up to time $\sigma_{k,i}^-$ and its out-component up to time $\sigma_{k,i}^+$. Note that we have added the subscript i, relative to the notation introduced in Section 3.4.3.1, to emphasize that the exploration starts at vertex i. Next, define $\mathcal{F}_{k,i}$ to be the sigma-algebra generated by the exploration of the in-component and out-component of vertex i, as described in Section 3.4.3.1, up to Step $\sigma_{k,i}^-$ on the inbound side and up to Step $\sigma_{k,i}^+$ on the outbound side. Define $\mathcal{N}_i^{(k)} = I^+_{\sigma^+_{k,i}} \cup I^-_{\sigma^-_{k,i}} \cup A^-_{\sigma^-_{k,i}} \cup A^+_{\sigma^+_{k,i}}$ to be the set of vertices discovered during that exploration. Now explore the in-component and out-component of vertex j, as described in Section 3.4.3.1, up to Step $\sigma_{k,j}^-$ on the inbound side and up to Step $\sigma_{k,j}^+$ on the outbound side; let $\mathcal{N}_j^{(k)}$ be the corresponding set of vertices discovered during the exploration of vertex j.

Define $C_{ij} = \left\{\mathcal{N}_i^{(k)} \cap \mathcal{N}_j^{(k)} = \varnothing\right\}$ and note that,

$$\mathbb{E}\left[\chi_{n,i}^{\geq k}\chi_{n,j}^{\geq k}\right] \leq \mathbb{E}\left[\chi_{n,i}^{\geq k}\chi_{n,j}^{\geq k}1(C_{ij})\right] + \mathbb{E}\left[1(C_{ij}^c)\right] = \mathbb{E}\left[\chi_{n,i}^{\geq k}\mathbb{E}\left[\chi_{n,j}^{\geq k}1(C_{ij})\middle|\mathcal{F}_{k,i}\right]\right] + \mathbb{P}(C_{ij}^c).$$

To analyze the conditional expectation, observe that

$$\mathbb{E}\left[\chi_{n,j}^{\geq k}1(C_{ij})\middle|\mathcal{F}_{k,i}\right] = \mathbb{E}\left[\chi_{n,j}^{\geq k}\middle|\mathcal{F}_{k,i}, C_{ij}\right]\mathbb{P}(C_{ij}|\mathcal{F}_{k,i}),$$

where, due to the independence among the arcs, we have that conditionally on $\mathcal{F}_{k,i}$ and C_{ij}, the random variable $\chi_{n,j}^{\geq k}$ has the same distribution as the indicator function of the event that vertex j has in-component and out-component both of size at least k on the graph $G_n(\kappa_{n,i})$, with

$$\kappa_{n,i}(\mathbf{X}_s, \mathbf{X}_t) = \kappa(\mathbf{X}_s, \mathbf{X}_t)(1 + \varphi_n(\mathbf{X}_s, \mathbf{X}_t))1(s \notin \mathcal{N}_i^{(k)}, t \notin \mathcal{N}_i^{(k)}).$$

Now note that since $\kappa_{n,i} \leq \kappa(1 + \varphi_n)$ for any realization of $\mathcal{N}_i^{(k)} \subseteq [n]$, we have

$$\mathbb{E}\left[\chi_{n,j}^{\geq k}\middle|\mathcal{F}_{k,i}, C_{ij}\right] \leq \mathbb{E}\left[\chi_{n,j}^{\geq k}\right],$$

from where it follows that

$$\frac{1}{n^2}\sum_{i=1}^n\sum_{j\neq i}\mathbb{E}\left[\chi_{n,i}^{\geq k}\chi_{n,j}^{\geq k}\right] \leq \frac{1}{n^2}\sum_{i=1}^n\sum_{j\neq i}\left(\mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\mathbb{E}\left[\chi_{n,j}^{\geq k}\right] + \mathbb{P}(C_{ij}^c)\right),$$

which in turn implies that

$$\frac{1}{n^2}\sum_{i=1}^n\sum_{j\neq i}\mathbb{E}\left[\left(\chi_{n,i}^{\geq k} - \mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\right)\left(\chi_{n,j}^{\geq k} - \mathbb{E}\left[\chi_{n,j}^{\geq k}\right]\right)\right] \leq \frac{1}{n^2}\sum_{i=1}^n\sum_{j\neq i}\mathbb{P}(C_{ij}^c).$$

Similarly to what was done on the graph, define $\hat{\mathcal{N}}_i^{(k)} = \hat{I}^+_{\hat{\sigma}^+_{k,i}} \cup \hat{I}^-_{\hat{\sigma}^-_{k,i}} \cup \left\{ T_{\mathbf{i}} : \mathbf{i} \in \hat{A}^-_{\hat{\sigma}^-_{k,i}} \right\} \cup \left\{ T_{\mathbf{i}} : \mathbf{i} \in \hat{A}^+_{\hat{\sigma}^+_{k,i}} \right\}$ to be the set of *identities* that appear during the construction of the double tree
$(\mathcal{T}^-_\mu(\kappa_m; \mathbf{X}_i), \mathcal{T}^+_\mu(\kappa_m; \mathbf{X}_i))$ up to Step $\hat{\sigma}^-_{k,i}$ on the inbound side, and up to Step $\hat{\sigma}^+_{k,i}$ on the outbound side. Let $\hat{C}_{ij} = \left\{ \hat{\mathcal{N}}_i^{(k)} \cap \hat{\mathcal{N}}_j^{(k)} = \varnothing \right\}$. We then have

$$\begin{aligned}
\mathbb{P}(C^c_{ij}) &\leq \mathbb{P}(C^c_{ij}, \tau^-_i > \sigma^-_{k,i}, \tau^+_i > \sigma^+_{k,i}, \tau^-_j > \sigma^-_{k,j}, \tau^+_j > \sigma^+_{k,j}) 1(\Omega_{m,n}) + 1(\Omega^c_{m,n}) \\
&\quad + \mathbb{P}(\{\tau^-_i \leq \sigma^-_{k,i}\} \cup \{\tau^+_i \leq \sigma^+_{k,i}\}) + \mathbb{P}(\{\tau^-_j \leq \sigma^-_{k,j}\} \cup \{\tau^+_j \leq \sigma^+_{k,j}\}) \\
&\leq 1(\Omega^c_{m,n}) + \mathbb{P}(\hat{C}^c_{ij}, |\hat{\mathcal{N}}_i^{(k)}| \leq \log n) 1(\Omega_{m,n}) + \mathbb{P}(|\hat{\mathcal{N}}_i^{(k)}| > \log n) \\
&\quad + \mathbb{P}_i(\{\tau^- \leq \sigma^-_k\} \cup \{\tau^+ \leq \sigma^+_k\}) + \mathbb{P}_j(\{\tau^- \leq \sigma^-_k\} \cup \{\tau^+ \leq \sigma^+_k\}),
\end{aligned}$$

where the event $\Omega_{m,n}$ is defined in Theorem 3.4.6.

To bound the first probability on the right-hand side, define $\hat{\mathcal{F}}_{k,i}$ to be the sigma-algebra generated by the construction of the double tree whose root has *identity* i, up to Step $\hat{\sigma}^-_{k,i}$ on the inbound side and up to Step $\hat{\sigma}^+_{k,i}$ on the outbound side. Now note that

$$\hat{C}_{ij} = \{j \notin \hat{\mathcal{N}}_i^{(k)}\} \cap \left(\bigcap_{r=1}^{\hat{\sigma}^+_{k,j}} \bigcap_{t \in \hat{\mathcal{N}}_i^{(k)}} \{\tilde{Z}_{\hat{T}^+_{r,j},t} = 0\} \right) \cap \left(\bigcap_{r=1}^{\hat{\sigma}^-_{k,j}} \bigcap_{t \in \hat{\mathcal{N}}_i^{(k)}} \{Z_{t,\hat{T}^-_{r,j}} = 0\} \right)$$

where $\hat{T}^-_{r,j}$ and $\hat{T}^+_{r,j}$ are the rth active *identities* to have their offspring sampled in the double tree whose root is j. Moreover, if we define $B_s = \bigcap_{t \in \hat{\mathcal{N}}_i^{(k)}} \{Z_{ts} = 0\}$ and $\tilde{B}_s = \bigcap_{t \in \hat{\mathcal{N}}_i^{(k)}} \{\tilde{Z}_{st} = 0\}$, then

$$\{j \notin \hat{\mathcal{N}}_i^{(k)}\} = B_j \cap \tilde{B}_j \qquad \text{and} \qquad \hat{C}_{ij} = B_j \cap \tilde{B}_j \cap \left(\bigcap_{r=1}^{\hat{\sigma}^+_{k,j}} \tilde{B}_{\hat{T}^+_{r,j}} \right) \cap \left(\bigcap_{r=1}^{\hat{\sigma}^-_{k,j}} B_{\hat{T}^-_{r,j}} \right),$$

and therefore, since $\hat{\sigma}^-_{k,j}, \hat{\sigma}^-_{k,j} \leq k$, the union bound gives

$$
\begin{aligned}
\mathbb{P}(\hat{C}^c_{ij}|\hat{\mathcal{F}}^{(k)}_i) &\leq \mathbb{P}\left(B^c_j \cup \left(\bigcup_{r=1}^{\hat{\sigma}^-_{k,j}} B^c_{\hat{T}^-_{r,j}} \right) \middle| \hat{\mathcal{F}}^{(k)}_i \right) + \mathbb{P}\left(\tilde{B}^c_j \cup \left(\bigcup_{r=1}^{\hat{\sigma}^+_{k,j}} \tilde{B}^c_{\hat{T}^+_{r,j}} \right) \middle| \hat{\mathcal{F}}^{(k)}_i \right) \\
&\leq \mathbb{E}\left[1(B^c_j) + \sum_{r=1}^{\hat{\sigma}^-_{k,j}} 1\left(B_j \cap \bigcap_{s=1}^{r-1} B_{\hat{T}^-_{s,j}} \cap B^c_{\hat{T}^-_{r,j}} \right) \middle| \hat{\mathcal{F}}^{(k)}_i \right] \\
&\quad + \mathbb{E}\left[1(\tilde{B}^c_j) + \sum_{r=1}^{\hat{\sigma}^+_{k,j}} 1\left(\tilde{B}_j \cap \bigcap_{s=1}^{r-1} \tilde{B}_{\hat{T}^+_{s,j}} \cap \tilde{B}^c_{\hat{T}^+_{r,j}} \right) \middle| \hat{\mathcal{F}}^{(k)}_i \right] \\
&\leq \mathbb{E}\left[1(B^c_j) + \sum_{r=1}^{k} 1\left(\hat{A}^-_{r-1,j} \neq \varnothing, \bigcap_{s=1}^{r} \{\hat{T}^-_{s,j} \notin \hat{\mathcal{N}}^{(k)}_i\}, B^c_{\hat{T}^-_{r,j}} \right) \middle| \hat{\mathcal{F}}^{(k)}_i \right] \qquad (3.4.17) \\
&\quad + \mathbb{E}\left[1(\tilde{B}^c_j) + \sum_{r=1}^{k} 1\left(\hat{A}^+_{r-1,j} \neq \varnothing, \bigcap_{s=1}^{r} \{\hat{T}^+_{s,j} \notin \hat{\mathcal{N}}^{(k)}_i\}, \tilde{B}^c_{\hat{T}^+_{r,j}} \right) \middle| \hat{\mathcal{F}}^{(k)}_i \right], \qquad (3.4.18)
\end{aligned}
$$

where $\hat{A}^-_{r,j}$ and $\hat{A}^+_{r,j}$ are the rth inbound and outbound active sets in the construction of the double tree started at j. Now note that the event $\bigcap_{s=1}^r \{\hat{T}^-_{s,j} \notin \hat{\mathcal{N}}^{(k)}_i\}$ implies that none of the $\{U_{s,\hat{T}^-_{r,j}} : 1 \leq s \leq n\}$ have been used in the construction of the double tree started at i, hence

$$
\mathbb{P}\left(\hat{A}^-_{r-1,j} \neq \varnothing, \bigcap_{s=1}^{r} \{\hat{T}^-_{s,j} \notin \hat{\mathcal{N}}^{(k)}_i\}, B^c_{\hat{T}^-_{r,j}} \middle| \hat{\mathcal{F}}^{(k)}_i \right) \leq \mathbb{E}\left[1(\hat{A}^-_{r-1,j} \neq \varnothing) Q(\hat{\mathcal{N}}^{(k)}_i, \hat{T}^-_{r,j}) \right],
$$

where for any set $V \subseteq [n]$ and any $s \in [n]$ we define

$$
\begin{aligned}
Q(V,s) &= \mathbb{P}\left(\bigcup_{t\in V} \{Z_{ts} \geq 1\} \right) \leq \sum_{t\in V} P(Z_{ts} \geq 1) = \sum_{t \in V} (1 - e^{-r^{(m,n)}_{ts}}) \\
&\leq \sum_{t\in V} r^{(m,n)}_{ts} \leq \frac{R_n}{n} \sum_{t\in V} \kappa_m(\mathbf{X}_t, \mathbf{X}_s) \leq \frac{R_n}{n} |V| \sup_{\mathbf{x},\mathbf{y}\in\mathcal{S}} \kappa_m(\mathbf{x},\mathbf{y}),
\end{aligned}
$$

and $R_n = \max_{1\leq t\leq M_m} 1(\mu_n(\mathcal{J}^{(m)}_t) > 0)\mu(\mathcal{J}^{(m)}_t)/\mu_n(\mathcal{J}^{(m)}_t)$. Since $\mathbb{P}(B^c_j|\hat{\mathcal{F}}^{(k)}_i) \leq Q(\hat{\mathcal{N}}^{(k)}_i, j)$, we obtain that (3.4.17) is bounded from above by

$$
Q(\hat{\mathcal{N}}^{(k)}_i, j) + \sum_{r=1}^{k} \mathbb{E}\left[1(\hat{A}^-_{r-1,j} \neq \varnothing) Q(\hat{\mathcal{N}}^{(k)}_i, \hat{T}^-_{r,j}) \middle| \hat{\mathcal{F}}^{(k)}_i \right] \leq \frac{R_n(k+1)}{n} |\hat{\mathcal{N}}^{(k)}_i| \sup_{\mathbf{x},\mathbf{y}\in\mathcal{S}} \kappa_m(\mathbf{x},\mathbf{y}).
$$

Similarly, (3.4.18) is bounded from above by

$$
\frac{R_n(k+1)}{n} |\hat{\mathcal{N}}^{(k)}_i| \sup_{\mathbf{x},\mathbf{y}\in\mathcal{S}} \kappa_m(\mathbf{x},\mathbf{y}).
$$

It follows that

$$\mathbb{P}(\hat{C}_{ij}^c|\hat{\mathcal{F}}_i^{(k)}) \leq \frac{2R_n(k+1)}{n}|\hat{\mathcal{N}}_i^{(k)}| \sup_{\mathbf{x},\mathbf{y}\in\mathcal{S}} \kappa_m(\mathbf{x},\mathbf{y}),$$

which in turn implies that for any $i, j \in [n]$,

$$\begin{aligned}\mathbb{P}(\hat{C}_{ij}^c, |\hat{\mathcal{N}}_i^{(k)}| < \log n) 1(\Omega_{m,n}) &= \mathbb{E}\left[\mathbb{P}(\hat{C}_{ij}^c|\hat{\mathcal{F}}_i^{(k)}) 1(|\hat{\mathcal{N}}_i^{(k)}| < \log n)\right] 1(\Omega_{m,n}) \\ &\leq \frac{4(k+1)\log n}{n} \sup_{\mathbf{x},\mathbf{y}\in\mathcal{S}} \kappa_m(\mathbf{x},\mathbf{y}),\end{aligned}$$

and we have used the observation that on $\Omega_{m,n}$ we have $R_n \leq 1+\epsilon \leq 2$.

Using this estimate we obtain that

$$\begin{aligned}\frac{1}{n^2}\sum_{i=1}^n \sum_{j\neq i} \mathbb{P}(C_{ij}^c) &\leq \frac{1}{n^2}\sum_{i=1}^n \sum_{j\neq i} \left\{ 1(\Omega_{m,n}^c) + \mathbb{P}(\hat{C}_{ij}^c,\, |\hat{\mathcal{N}}_i^{(k)}| \leq \log n) 1(\Omega_{m,n}) + \mathbb{P}(|\hat{\mathcal{N}}_i^{(k)}| > \log n) \right\} \\ &\quad + \frac{2}{n}\sum_{i=1}^n \mathbb{P}_i(\{\tau^- \leq \sigma_k^-\} \cup \{\tau^+ \leq \sigma_k^+\}) \\ &\leq 1(\Omega_{m,n}^c) + \frac{4(k+1)\log n}{n} \sup_{\mathbf{x},\mathbf{y}\in\mathcal{S}} \kappa_m(\mathbf{x},\mathbf{y}) + \frac{1}{n}\sum_{i=1}^n \mathbb{P}(|\hat{\mathcal{N}}_i^{(k)}| > \log n) \\ &\quad + \frac{2}{n}\sum_{i=1}^n \mathbb{P}_i(\{\tau^- \leq \sigma_k^-\} \cup \{\tau^+ \leq \sigma_k^+\}).\end{aligned}$$

To complete the proof, apply Theorem 3.4.6 to obtain

$$1(\Omega_{m,n}^c) + \frac{2}{n}\sum_{i=1}^n \mathbb{P}_i(\{\tau^- \leq \sigma_k^-\} \cup \{\tau^+ \leq \sigma_k^+\}) \leq 3H(n,m,k,\epsilon),$$

and Markov's inequality followed by Lemma 3.4.7 to get

$$\frac{1}{n}\sum_{i=1}^n \mathbb{P}(|\hat{\mathcal{N}}_i^{(k)}| > \log n) \leq \frac{1}{n\log n}\sum_{i=1}^n \mathbb{E}\left[|\hat{\mathcal{N}}_i^{(k)}|\right] \leq \frac{k}{\log n}\left(2 + \sup_{\mathbf{x}\in\mathcal{S}} \lambda_-^{(m)}(\mathbf{x}) + \sup_{\mathbf{x}\in\mathcal{S}} \lambda_+^{(m)}(\mathbf{x})\right).$$

■

We now use Proposition 3.4.18 to show that the number of vertices with in-component and out-component both of size at least k converges in probability.

Proposition 3.4.19 *Define the in-component and out-component of vertex $v \in [n]$, $R^-(v)$ and $R^+(v)$, respectively, as in Theorem 3.3.11. Let $N_n^{\geq k} = \{v \in [n] : |R^-(v)| \geq k$ and $|R^+(v)| \geq k\}$. Then,*

$$\frac{|N_n^{\geq k}|}{n} \xrightarrow{P} \rho^{\geq k}(\kappa), \qquad n \to \infty.$$

Proof. Define $\{\chi_{n,i}^{\geq k}\}_{i\in[n]}$ as in Proposition 3.4.18. We start by noting that for any $m \geq 1$ we have

$$\left|\frac{|N_n^{\geq k}|}{n} - \rho^{\geq k}(\kappa)\right| \leq \left|\frac{|N_n^{\geq k}|}{n} - \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\right|$$
$$+ \left|\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\left[\chi_{n,i}^{\geq k}\right] - \frac{1}{n}\sum_{i=1}^{n}\rho_-^{\geq k}(\kappa_m;\mathbf{X}_i)\rho_+^{\geq k}(\kappa_m;\mathbf{X}_i)\right| \quad (3.4.19)$$
$$+ \left|\frac{1}{n}\sum_{i=1}^{n}\rho_-^{\geq k}(\kappa_m;\mathbf{X}_i)\rho_+^{\geq k}(\kappa_m;\mathbf{X}_i) - \rho^{\geq k}(\kappa_m)\right| + \left|\rho^{\geq k}(\kappa_m) - \rho^{\geq k}(\kappa)\right|$$

Moreover, by Proposition 3.4.18 we have that for any $0 < \epsilon < 1/2$, (3.4.19) is bounded by $H(n,m,k,\epsilon)$, where $H(n,m,k,\epsilon)$ is defined in Theorem 3.4.6 and satisfies $H(n,m,k,\epsilon) \xrightarrow{P} \hat{H}(m,k,\epsilon)$ for some other function $\hat{H}(m,k,\epsilon)$ (defined in Lemma 3.4.10) as $n \to \infty$, where for any fixed $k \geq 1$ we have

$$\lim_{m\nearrow\infty}\lim_{\epsilon\downarrow 0}\hat{H}(m,k,\epsilon) = 0.$$

Also, by the bounded convergence theorem we have that for any $m, k \in \mathbb{N}_-$,

$$\frac{1}{n}\sum_{i=1}^{n}\rho_-^{\geq k}(\kappa_m;\mathbf{X}_i)\rho_+^{\geq k}(\kappa_m;\mathbf{X}_i) = \int_{\mathcal{S}}\rho_-^{\geq k}(\kappa_m;\mathbf{x})\rho_+^{\geq k}(\kappa_m;\mathbf{x})\mu_n(d\mathbf{x}) \xrightarrow{P} \rho^{\geq k}(\kappa_m) \qquad n \to \infty,$$

and by Lemma 3.4.15 we have that

$$\lim_{m\nearrow\infty}\rho^{\geq k}(\kappa_m) = \rho^{\geq k}(\kappa)$$

in probability, since $\kappa_m \nearrow \kappa$ in probability. Therefore, for any $m \geq 1$ and $0 < \epsilon < 1/2$ we have

$$\limsup_{n\to\infty}\left|\frac{|N_n^{\geq k}|}{n} - \rho^{\geq k}(\kappa)\right| \leq \limsup_{n\to\infty}\left|\frac{|N_n^{\geq k}|}{n} - \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\right| + \hat{H}(m,k,\epsilon) + \left|\rho^{\geq k}(\kappa_m) - \rho^{\geq k}(\kappa)\right|.$$

and by taking $\epsilon \downarrow 0$ followed by $m \nearrow \infty$ we obtain that the following limit holds in probability

$$\limsup_{n\to\infty}\left|\frac{|N_n^{\geq k}|}{n} - \rho^{\geq k}(\kappa)\right| \leq \limsup_{n\to\infty}\left|\frac{|N_n^{\geq k}|}{n} - \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\right|.$$

It remains to show that this last limit is zero. To do this, start by using Proposition 3.4.18

again to obtain that for any $m \geq 1$ and $0 < \epsilon < 1/2$, we have that on the event $\Omega_{m,n}$,

$$
\begin{aligned}
&\mathbb{E}\left[\left(\frac{N_n^{\geq k}}{n} - \frac{1}{n}\sum_{i=1}^n \mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\right)^2\right] \\
&= \frac{1}{n^2}\left\{\sum_{i=1}^n \mathbb{E}\left[\left(\chi_{n,i}^{\geq k} - \mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\right)^2\right] + \sum_{i=1}^n\sum_{j\neq i}\mathbb{E}\left[\left(\chi_{n,i}^{\geq k} - \mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\right)\left(\chi_{n,j}^{\geq k} - \mathbb{E}\left[\chi_{n,j}^{\geq k}\right]\right)\right]\right\} \\
&\leq \frac{1}{n^2}\sum_{i=1}^n \mathbb{E}\left[\left(\chi_{n,i}^{\geq k} - \mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\right)^2\right] + K(n,m,k) + 3H(n,m,k,\epsilon) \\
&\leq \frac{1}{n} + K(n,m,k) + 3H(n,m,k,\epsilon),
\end{aligned}
$$

where $K(n,m,k)$ is defined in Proposition 3.4.18 and satisfies $K(n,m,k) \xrightarrow{P} 0$ as $n \to \infty$ for any fixed m, ϵ. The bounded convergence theorem now gives

$$
\lim_{n\to\infty} E\left[\left(\frac{N_n^{\geq k}}{n} - \frac{1}{n}\sum_{i=1}^n \mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\right)^2\right] = E\left[\lim_{n\to\infty}\mathbb{E}\left[\left(\frac{N_n^{\geq k}}{n} - \frac{1}{n}\sum_{i=1}^n \mathbb{E}\left[\chi_{n,i}^{\geq k}\right]\right)^2\right]\right] = 3\hat{H}(m,k,\epsilon),
$$

and taking the limit as $\epsilon \downarrow 0$ followed by $m \nearrow \infty$ completes the proof. ■

We are now ready to prove Theorem 3.3.10, the phase transition for the existence of a giant strongly connected component in $G_n(\kappa(1+\varphi_n))$.

Proof of Theorem 3.3.10. By Lemma 3.4.14, there exists a sequence of kernels $\{\kappa_m : m \geq 1\}$ defined on $\mathcal{S} \times \mathcal{S}$, measurable with respect to $\mathscr{F}$, such that κ_m is quasi-irreducible, regular finitary, and such that for any $n \geq m$, we have

$$
\kappa_m(\mathbf{x},\mathbf{y}) \leq \kappa(\mathbf{x},\mathbf{y})(1+\varphi_n(\mathbf{x},\mathbf{y})) \qquad \text{for all } \mathbf{x},\mathbf{y} \in \mathcal{S}.
$$

Proof of the lower bound: We will start by proving a lower bound for the largest strongly connected component of $G_n(\kappa,\varphi_n)$. To this end, note that we can construct a coupling between $G_n(\kappa(1+\varphi_n))$ and $G_n(\kappa_m)$ such that every arc in $G_n(\kappa_m)$ is also in $G_n(\kappa(1+\varphi_n))$ P-a.s. It follows that

$$
\mathcal{C}_1(G_n(\kappa(1+\varphi_n))) \supseteq \mathcal{C}_1(G_n(\kappa_m)) \quad P\text{-a.s.}
$$

The idea is now to apply Theorem 1 in [12] to $G_n(\kappa_m)$, however, that theorem requires that the kernel κ_m be irreducible, whereas κ_m is only quasi-irreducible. To address this issue, we construct a third graph as follows. Let $\mathcal{S}^*$ be the restriction of $\mathcal{S}$ where κ_m is irreducible and set

$$
\mathcal{S}' = \bigcup_{i=1}^{M_m}\left\{\mathcal{J}_i^{(m)} \cap \mathcal{S}^* : \mu(\mathcal{J}_i^{(m)}) > 0\right\}.
$$

To avoid trivial cases, assume from now on that $\mu(\mathcal{S}') > 0$.

Now let $V_{n'} = \{1 \leq i \leq n : \mathbf{X}_i \in \mathcal{S}'\}$ denote the set of vertices in $G_n(\kappa_m)$ that have types in $\mathcal{S}'$ and let n' denote its cardinality. Note that n' is random, but measurable with respect to $\mathscr{F}$. Next, fix $0 < \delta < 1$ and define the kernel $\kappa'(\mathbf{x}, \mathbf{y}) = (1-\delta)\mu(\mathcal{S}')\kappa_m(\mathbf{x}, \mathbf{y})$ and the graph $G_{n'}(\kappa')$ whose arc probabilities are given by

$$p_{ij}^{(n')} = \frac{(1-\delta)\mu(\mathcal{S}')\kappa_m(\mathbf{X}_i, \mathbf{X}_j)}{n'} \wedge 1, \qquad i, j \in [n'],\, i \neq j.$$

Note that $G_{n'}(\kappa')$ is a graph on the type space $\mathcal{S}'$ whose types are distributed according to measure $\mu_n'(A) := \mu_n(A)/\mu_n(\mathcal{S}')$ for any $A \subseteq \mathcal{S}'$. Moreover, κ' is irreducible on $\mathcal{S}'$ with each of its induced types, i.e., the sets $\mathcal{J}_i^{(m)} \cap \mathcal{S}'$, having strictly positive measure. Now note that since $n\mu_n(\mathcal{S}') = n'$ and $\mu_n(\mathcal{S}') \xrightarrow{P} \mu(\mathcal{S}')$ as $n \to \infty$, then

$$p_{ij}^{(n')} = \frac{(1-\delta)\mu(\mathcal{S}')\kappa_m(\mathbf{X}_i, \mathbf{X}_j)}{n\mu_n(\mathcal{S}')} \wedge 1 \leq \frac{\kappa_m(\mathbf{X}_i, \mathbf{X}_j)}{n} \wedge 1, \qquad i, j \in [n'],\, i \neq j,$$

for all sufficiently large n. Therefore, there exists a coupling such that every arc in $G_{n'}(\kappa')$ is also in $G_n(\kappa_m)$, and therefore, for all sufficiently large n,

$$\mathcal{C}_1(G_n(\kappa_m)) \supseteq \mathcal{C}_1(G_{n'}(\kappa')) \quad P\text{-a.s.}$$

Now use Theorem 1 in [12] to obtain that for every $\epsilon > 0$

$$P\left(\left| \frac{|\mathcal{C}_1(G_{n'}(\kappa'))|}{n'} - \rho'(\kappa') \right| > \epsilon \right) \to 0 \qquad n \to \infty,$$

where

$$\rho'(\kappa') = \int_{\mathcal{S}'} \rho'_-(\kappa'; \mathbf{x})\rho'_+(\kappa'; \mathbf{x})\mu'(d\mathbf{x}),$$

and $\rho'_-(\kappa'; \mathbf{x}), \rho'_+(\kappa'; \mathbf{x})$ are the survival probabilities of the trees $\mathcal{T}_{\mu'}^-(\kappa')$ and $\mathcal{T}_{\mu'}^+(\kappa')$, respectively, defined on the type space $\mathcal{S}'$ with respect to the measure $\mu'(A) = \mu(A)/\mu(\mathcal{S}')$ for $A \subseteq \mathcal{S}'$.

By the arguments in the proof of Proposition 3.4.16, we have that $\rho((1-\delta)\kappa_m) = \mu(\mathcal{S}')\rho'(\kappa')$, where

$$\rho((1-\delta)\kappa) = \int_{\mathcal{S}} \rho_-((1-\delta)\kappa_m; \mathbf{x})\rho_+((1-\delta)\kappa; \mathbf{x})\mu(d\mathbf{x}),$$

and $\rho_-((1-\delta)\kappa_m; \mathbf{x})$, $\rho_+((1-\delta)\kappa_m; \mathbf{x})$ are the survival probabilities of the trees $\mathcal{T}_\mu^-((1-\delta)\kappa_m)$ and $\mathcal{T}_\mu^+((1-\delta)\kappa_m)$, defined on the type space $\mathcal{S}$.

Hence,

$$\frac{|\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))|}{n} \geq \frac{|\mathcal{C}_1(G_n((1-\delta)\kappa_m))|}{n} \geq \frac{|\mathcal{C}_1(G_{n'}(\kappa'))|}{n'} \cdot \frac{n'}{n} \xrightarrow{P} \rho'(\kappa')\mu(\mathcal{S}') = \rho((1-\delta)\kappa_m),$$

as $n \to \infty$. Now use Lemma 3.4.15 to obtain that the following limits hold in probability

$$\lim_{m \nearrow \infty} \lim_{\delta \downarrow 0} \rho((1-\delta)\kappa_m) = \lim_{\delta \downarrow 0} \lim_{m \nearrow \infty} \rho((1-\delta)\kappa_m) = \rho(\kappa),$$

from where we conclude that for any $\epsilon > 0$,

$$P\left(\frac{|\mathcal{C}_1(G_n(\kappa, \varphi_n))|}{n} - \rho(\kappa) < -\epsilon \right) \to 0 \qquad n \to \infty.$$

Proof of the upper bound: For any $k, m \geq 1$ let $\rho_-^{\geq k}(\kappa_m; \mathbf{x})$ $(\rho_+^{\geq k}(\kappa_m; \mathbf{x}))$ denote the probability that the tree $\mathcal{T}_\mu^-(\kappa_m; \mathbf{x})$ $(\mathcal{T}_\mu^+(\kappa_m; \mathbf{x}))$ has a population of at least k nodes. Define for $k \geq 1$ the set $N_n^{\geq k}$ as in Proposition 3.4.19; $N_n^{\geq k}$ is the set of vertices in $G_n(\kappa(1+\varphi_n))$ with both large in-component and large out-component. Now note that provided $\liminf_{n\to\infty} |\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))| = \infty$, we have that for any fixed $k \geq 1$,

$$|\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))| \leq |N_n^{\geq k}| \qquad \text{for all sufficiently large } n.$$

It follows that

$$\frac{|\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))|}{n} - \rho(\kappa) \leq \frac{|N_n^{\geq k}|}{n} - \rho^{\geq k}(\kappa) + \rho^{\geq k}(\kappa) - \rho(\kappa).$$

Now use Proposition 3.4.19 to obtain that $|N_n^{\geq k}|/n \xrightarrow{P} \rho^{\geq k}(\kappa)$ as $n \to \infty$ for any fixed $k \geq 1$. Now use Lemma 3.4.15 to obtain that $\rho^{\geq k}(\kappa) \nearrow \rho(\kappa)$ as $k \nearrow \infty$, which completes the proof of the upper bound.

Proof of the phase transition: It follows from Lemma 3.4.17. ■

We now proof Theorem 3.3.11, which provides a more detailed description of the giant strongly connected component and of the bow-tie structure it determines.

Proof of Theorem 3.3.11. In view of Theorem 3.3.10, we know that $\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))$ contains asymptotically $n\rho(\kappa)$ vertices, and therefore, $\mathcal{C}_1(G_n(\kappa(1+\varphi_n))) \subseteq L_n^+ \cap L_n^-$ with high probability. To show the reverse subset relation fix $0 < \delta < 1$ and $m \geq 1$ and construct the kernel $\kappa'(\mathbf{x}, \mathbf{y}) = (1-\delta)\mu(\mathcal{S}')\kappa_m(\mathbf{x}, \mathbf{y})$ on the type space $\mathcal{S}' \subseteq \mathcal{S}$ just as in the proof of Theorem 3.3.10, so that κ' is regular finitary and irreducible on $\mathcal{S}'$. Now construct the graph $G_{n'}(\kappa')$ using a coupling ensuring that

$$\mathcal{C}_1(G_{n'}(\kappa')) \subseteq \mathcal{C}_1(G_n(\kappa(1+\varphi_n))) \qquad P\text{-a.s.},$$

as was done in the proof of Theorem 3.3.10. Now define the sets $L_{m,\delta,n}^- = \{v \in [n'] : |R^-(v)| \geq (\log n)/n\}$ and $L_{m,\delta,n}^+ = \{v \in [n'] : |R^+(v)| \geq (\log n)/n\}$ relative to graph $G_{n'}(\kappa')$. A close inspection of the proof of Theorem 1 in [12] shows that $L_{m,\delta,n}^+ \cap L_{m,\delta,n}^-$ (denoted

$B'(\omega_1)$ in [12]) is strongly connected with high probability (specifically, see the proof of (22) in [12]). It follows that

$$\lim_{n\to\infty} P\left(L^+_{m,\delta,n}\cap L^-_{m,\delta,n}\subseteq \mathcal{C}_1(G_{n'}(\kappa'))\subseteq \mathcal{C}_1(G_n(\kappa(1+\varphi_n)))\subseteq L^+_n\cap L^-_n\right)=1.$$

Now note that $L^+_{m,\delta,n}\nearrow L^+_n$ and $L^-_{m,\delta,n}\nearrow L^-_n$ as $m\to\infty$ and $\delta\to 0$, which implies that

$$\lim_{n\to\infty} P\left(\mathcal{C}_1(G_n(\kappa(1+\varphi_n)))=L^+_n\cap L^-_n\right)=1.$$

To establish the limits for $|L^+_n|$ and $|L^-_n|$, let $L^{\geq k}_{-,n}=\{v\in[n]:|R^-(v)|\geq k\}$ and $L^{\geq k}_{+,n}=\{v\in[n]:|R^+(v)|\geq k\}$ and note that $|L^\pm_n|\leq |L^{\geq k}_{\pm,n}|$ for any $1\leq k\leq (\log n)/n$ and $|L^{\geq k}_{\pm,n}|\leq |L^\pm_n|$ for any $k\geq (\log n)/n$. A straightforward adaptation of Proposition 3.4.19 can be used to obtain

$$\frac{|L^{\geq k}_{-,n}|}{n}\xrightarrow{P}\int_{\mathcal{S}}\rho^{\geq k}_-(\kappa;\mathbf{x})\mu(d\mathbf{x})\qquad\text{and}\qquad \frac{|L^{\geq k}_{+,n}|}{n}\xrightarrow{P}\int_{\mathcal{S}}\rho^{\geq k}_+(\kappa;\mathbf{x})\mu(d\mathbf{x})$$

as $n\to\infty$. Monotone convergence gives $\int_{\mathcal{S}}\rho^{\geq k}_\pm(\kappa;\mathbf{x})\mu(d\mathbf{x})\searrow\int_{\mathcal{S}}\rho_\pm(\kappa;\mathbf{x})\mu(d\mathbf{x})$ as $k\nearrow\infty$, which yields the result. ■

We end the paper with the proof of Proposition 3.3.13, which states the main results for the rank-1 kernel case.

Proof of Proposition 3.3.13. The first two statements follow immediately from noting that $E[\kappa_+(\mathbf{X})]E[\kappa_-(\mathbf{X})]=\iint_{\mathcal{S}^2}\kappa(\mathbf{x},\mathbf{y})\mu(d\mathbf{x})\mu(d\mathbf{y})$. The third one follows from noting that $\lambda_-(\mathbf{X})=\kappa_-(\mathbf{X})E[\kappa_+(\mathbf{X})]$ and $\lambda_+(\mathbf{X})=\kappa_+(\mathbf{X})E[\kappa_-(\mathbf{X})]$.

To establish (d) assume first that $\rho(\kappa)>0$. Now use Lemma 3.4.14 (applied to $\kappa(\mathbf{x},\mathbf{y})=\kappa_-(\mathbf{y})$ and $\kappa(\mathbf{x},\mathbf{y})=\kappa_+(\mathbf{x})$ separately) to obtain that there exists a sequence of kernels $\{\kappa^+_m(\mathbf{x}):m\geq 1\}$ and $\{\tilde{\kappa}^-_m(\mathbf{x}):m\geq 1\}$ such that: 1) $0\leq\kappa^\pm_m(\mathbf{x})\leq\kappa_\pm(\mathbf{x})$ for all $\mathbf{x}\in\mathcal{S}$, 2) each is piecewise constant taking only a finite number of values, and 3) $\kappa^\pm_m(\mathbf{x})\nearrow\kappa_\pm(\mathbf{x})$ in probability for a.e. $\mathbf{x}\in\mathcal{S}$ as $m\to\infty$. Now set $B_m=\{\mathbf{x}\in\mathcal{S}:\kappa^-_m(\mathbf{x})>0,\kappa^+_m(\mathbf{x})>0\}$ and define

$$\kappa_m(\mathbf{x},\mathbf{y})=\kappa^+_m(\mathbf{x})\kappa^-_m(\mathbf{y})1(\mathbf{x}\in B_m,\mathbf{y}\in B_m).$$

Note that κ_m is regular finitary and is strictly positive on $B_m\times B_m$. Hence, the only set $A\subseteq B_m$ satisfying $\kappa_m=0$ on $A\times(A^c\cap B_m)$ is $A=\varnothing$ or $A^c\cap B_m=\varnothing$, implying the irreducibility of κ_m on $B_m\times B_m$. Moreover, since $\kappa_->0$ and $\kappa_+>0$ a.e. in order for κ to be irreducible, we have that $\kappa_m\nearrow\kappa$ in probability as $m\to\infty$.

Next, use Lemma 3.4.15 to obtain that $\rho(\kappa)=\lim_{m\to\infty}\rho(\kappa_m)$, and therefore, $\rho(\kappa_m)>0$ for some m sufficiently large. By Proposition 3.4.16 this implies that the spectral radii of the operators $T^-_{\kappa_m}$ and $T^+_{\kappa_m}$ are strictly larger than one. Now note that the functions $f^-_m(\mathbf{x})=\kappa^-_m(\mathbf{x})$ and $f^+_m(\mathbf{x})=\kappa^+_m(\mathbf{x})$ are nonnegative and satisfy

$$\begin{aligned}T^-_{\kappa_m}f^-_m(\mathbf{x})&=\int_{\mathcal{S}}\kappa^+_m(\mathbf{y})\kappa^-_m(\mathbf{x})f^-_m(\mathbf{y})\mu(d\mathbf{x})=\kappa^-_m(\mathbf{x})\int_{\mathcal{S}}\kappa^+_m(\mathbf{y})f^-_m(\mathbf{y})\mu(d\mathbf{y})\\&=f^-_m(\mathbf{x})\int_{\mathcal{S}}\kappa^+_m(\mathbf{y})\kappa^-_m(\mathbf{y})\mu(d\mathbf{y}),\end{aligned}$$

and therefore, $r_m := \int_{\mathcal{S}} \kappa_m^+(\mathbf{y})\kappa_m^-(\mathbf{y})\mu(d\mathbf{y})$ is an eigenvalue of $T_{\kappa_m}^-$. Similarly, r_m is an eigenvalue of $T_{\kappa_m}^+$ associated to the nonnegative eigenfunction f_m^+. Since we may assume that $\kappa_m^-(\mathbf{x})$ and $\kappa_m^+(\mathbf{x})$ are different from zero for sufficiently large m, then Proposition 3.4.16 gives that $r_m = r(T_{\kappa_m}^{\pm}) > 1$. Taking the limit as $m \to \infty$ gives that

$$E[\kappa_+(\mathbf{X})\kappa_-(\mathbf{X})] = \lim_{m\to\infty} r_m > 1.$$

For the converse, note that $E[\kappa_+(\mathbf{X})\kappa_-(\mathbf{X})] > 1$ and the monotone convergence theorem imply that $r_m > 1$ for some m sufficiently large. For this m, Proposition 3.4.16 gives that r_m is the spectral radius of $T_{\kappa_m}^-$ and $T_{\kappa_m}^+$, and also that $\rho(\kappa_m) > 0$. Lemma 3.4.15 now gives that $1 < \rho(\kappa_m) \nearrow \rho(\kappa)$ in probabiity as $m \to \infty$. ■

Bibliography

[1] Emile HL Aarts, Jan HM Korst, and Peter JM van Laarhoven. “A quantitative analysis of the simulated annealing algorithm: A case study for the traveling salesman problem”. In: *Journal of Statistical Physics* 50.1-2 (1988), pp. 187–206.

[2] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. “Improved algorithms for linear stochastic bandits”. In: *Advances in Neural Information Processing Systems*. 2011, pp. 2312–2320.

[3] Rediet Abebe, Lada A Adamic, and Jon Kleinberg. “Mitigating overexposure in viral marketing”. In: *Thirty-Second AAAI Conference on Artificial Intelligence*. 2018.

[4] Deepak Agarwal et al. “Online models for content optimization”. In: *Advances in Neural Information Processing Systems*. 2009, pp. 17–24.

[5] Shipra Agrawal et al. “MNL-bandit: a dynamic learning approach to assortment selection”. In: *arXiv preprint: 1706.03880* (2017).

[6] Shipra Agrawal et al. “Thompson sampling for the MNL-bandit”. In: *arXiv preprint: 1706.00977* (2017).

[7] David L Applegate et al. *The traveling salesman problem: a computational study*. Princeton university press, 2011.

[8] Krishna B Athreya and Peter E Ney. *Branching Processes*. Dover, 1972.

[9] Peter Auer. “Using confidence bounds for exploitation-exploration trade-offs”. In: *Journal of Machine Learning Research* 3.Nov (2002), pp. 397–422.

[10] Gary Barnes and Peter Langworthy. “The per-mile costs of operating automobiles and trucks”. In: (2003).

[11] Hamsa Bastani et al. “Sequential Learning of Product Recommendations with Customer Disengagement”. In: *Available at SSRN 3240970* (2018).

[12] Mindaugas Bloznelis, Friedrich Götze, and Jerzy Jaworski. “Birth of a strongly connected giant in an inhomogeneous random digraph”. In: *Journal of Applied Probability* 49.03 (2012), pp. 601–611.

[13] Béla Bollobás. “Random graphs”. In: *Modern Graph Theory*. Springer, 1998, pp. 215–252.

[14] Béla Bollobás, Svante Janson, and Oliver Riordan. "The phase transition in inhomogeneous random graphs". In: *Random Structures & Algorithms* 31.1 (2007), pp. 3–122.

[15] Tom Britton, Maria Deijfen, and Anders Martin-Löf. "Generating simple random graphs with prescribed degree distribution". In: *Journal of Statistical Physics* 124.6 (2006), pp. 1377–1397.

[16] John W Byers, Michael Mitzenmacher, and Georgios Zervas. "The groupon effect on yelp ratings: a root cause analysis". In: *Proceedings of the 13th ACM conference on electronic commerce*. ACM. 2012, pp. 248–265.

[17] Felipe Caro and Jérémie Gallien. "Dynamic assortment with demand learning for seasonal consumer goods". In: *Management Science* 53.2 (2007), pp. 276–292.

[18] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.

[19] Ningyuan Chen and Mariana Olvera-Cravioto. "Directed random graphs with given degree distributions". In: *Stochastic Systems* 3.1 (2013), pp. 147–186.

[20] Wei Chen, Yajun Wang, and Yang Yuan. "Combinatorial multi-armed bandit: General framework and applications". In: *International Conference on Machine Learning*. 2013, pp. 151–159.

[21] Jiesi Cheng, Aaron Sun, and Daniel Zeng. "Information overload and viral marketing: countermeasures and strategies". In: *International Conference on Social Computing, Behavioral Modeling, and Prediction*. Springer. 2010, pp. 108–117.

[22] Wang Chi Cheung and David Simchi-Levi. "Thompson Sampling for Online Personalized Assortment Optimization Problems with Multinomial Logit Choice Models". In: (2017).

[23] Wang Chi Cheung, Vincent YF Tan, and Zixin Zhong. "Thompson Sampling Algorithms for Cascading Bandits". In: *arXiv preprint arXiv:1810.01187* (2018).

[24] Wei Chu et al. "Contextual bandits with linear payoff functions". In: *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*. 2011, pp. 208–214.

[25] Aleksandr Chuklin, Ilya Markov, and Maarten de Rijke. "Click models for web search". In: *Synthesis Lectures on Information Concepts, Retrieval, and Services* 7.3 (2015), pp. 1–115.

[26] Fan Chung and Linyuan Lu. "Connected components in random graphs with given expected degree sequences". In: *Annals of combinatorics* 6.2 (2002), pp. 125–145.

[27] Fan Chung and Linyuan Lu. "The average distances in random graphs with given expected degrees". In: *Proceedings of the National Academy of Sciences* 99.25 (2002), pp. 15879–15882.

[28] Fan Chung, Linyuan Lu, and Van Vu. “Spectra of random graphs with given expected degrees”. In: *Proceedings of the National Academy of Sciences* 100.11 (2003), pp. 6313–6318.

[29] Matthew P Clay and Nándor J Simányi. “Rényi’s parking problem revisited”. In: *Stochastics and Dynamics* 16.02 (2016), p. 1660006.

[30] Edward G Coffman Jr et al. “Packing random intervals on-line”. In: *Algorithmica* 22.4 (1998), pp. 448–476.

[31] Richard Combes et al. “Learning to rank: Regret lower bounds and efficient algorithms”. In: *ACM SIGMETRICS Performance Evaluation Review* 43.1 (2015), pp. 231–244.

[32] William J Cook. *In pursuit of the traveling salesman: mathematics at the limits of computation*. Princeton University Press, 2011.

[33] Nick Craswell et al. “An experimental comparison of click position-bias models”. In: *Proceedings of the 2008 international conference on web search and data mining*. ACM. 2008, pp. 87–94.

[34] Carlos F Daganzo. *Logistics systems analysis*. Springer Science & Business Media, 2005.

[35] Carlos F Daganzo. “The distance traveled to visit N points with a maximum of C stops per vehicle: An analytic model and an application”. In: *Transportation Science* 18.4 (1984), pp. 331–350.

[36] Guy Degla. “An overview of semi-continuity results on the spectral radius and positivity”. In: *Journal of Mathematical Analysis and Applications* 338.1 (2008), pp. 101–110.

[37] Rick Durrett. *Probability: theory and examples*. Cambridge university press, 2010.

[38] Paul Erdos and Alfréd Rényi. “On the evolution of random graphs”. In: *Publ. Math. Inst. Hung. Acad. Sci* 5.1 (1960), pp. 17–60.

[39] Paul Erdős and Alfréd Rényi. “On random graphs I”. In: *Publ. Math. Debrecen* 6 (1959), pp. 290–297.

[40] Paul Erdős and Alfréd Rényi. “On the strength of connectedness of a random graph”. In: *Acta Mathematica Academiae Scientiarum Hungarica* 12.1-2 (1964), pp. 261–267.

[41] Sarah Filippi et al. “Parametric bandits: The generalized linear case”. In: *Advances in Neural Information Processing Systems*. 2010, pp. 586–594.

[42] Alan Frieze and Michał Karoński. *Introduction to random graphs*. Cambridge University Press, 2015.

[43] Ricardo Fukasawa et al. “Robust branch-and-cut-and-price for the capacitated vehicle routing problem”. In: *Mathematical programming* 106.3 (2006), pp. 491–511.

[44] Luis A Goddyn. "Quantizers and the worst-case Euclidean traveling salesman problem". In: *Journal of Combinatorial Theory, Series B* 50.1 (1990), pp. 65–81.

[45] Bruce L Golden, Subramanian Raghavan, and Edward A Wasil. *The vehicle routing problem: latest advances and new challenges*. Vol. 43. Springer Science & Business Media, 2008.

[46] Jan Grandell. *Mixed poisson processes*. Vol. 77. CRC Press, 1997.

[47] Sushil K Gupta, Jerzy Kyparisis, and Chi-Ming Ip. "Note—Project selection and sequencing to maximize net present value of the total return". In: *Management Science* 38.5 (1992), pp. 751–752.

[48] MARK Haimovich and AHG Rinnooy Kan. "Bounds and heuristics for capacitated routing problems". In: *Mathematics of operations Research* 10.4 (1985), pp. 527–542.

[49] Jonathan V Hall and Alan B Krueger. "An analysis of the labor market for Uber's driver-partners in the United States". In: *ILR Review* (2015), p. 0019793917717222.

[50] Jonathan Hall, Cory Kendrick, and Chris Nosko. "The effects of Uber's surge pricing: A case study". In: *The University of Chicago Booth School of Business* (2015).

[51] Karla L Hoffman, Manfred Padberg, and Giovanni Rinaldi. "Traveling salesman problem". In: *Encyclopedia of operations research and management science*. Springer, 2013, pp. 1573–1578.

[52] Chuck Holland et al. "UPS optimizes delivery routes". In: *Interfaces* 47.1 (2017), pp. 8–23.

[53] Pim van der Hoorn and Mariana Olvera-Cravioto. "Typical distances in the directed configuration model". In: *The Annals of Applied Probability* 28.3 (2018), pp. 1739–1792.

[54] Nabin Kafle, Bo Zou, and Jane Lin. "Design and modeling of a crowdsource-enabled system for urban parcel relay and delivery". In: *Transportation research part B: methodological* 99 (2017), pp. 62–82.

[55] Howard J Karloff. "How long can a Euclidean traveling salesman tour be?" In: *SIAM Journal on Discrete Mathematics* 2.1 (1989), pp. 91–99.

[56] Richard M Karp. "The transitive closure of a random digraph". In: *Random Structures & Algorithms* 1.1 (1990), pp. 73–93.

[57] A Gürhan Kök, Marshall L Fisher, and Ramnath Vaidyanathan. "Assortment planning: Review of literature and industry practice". In: *Retail supply chain management*. Springer, 2008, pp. 99–153.

[58] Junpei Komiyama, Junya Honda, and Akiko Takeda. "Position-based multiple-play bandit problem with unknown position bias". In: *Advances in Neural Information Processing Systems*. 2017, pp. 4998–5008.

[59] Ivan Kryven. “Emergence of the giant weak component in directed random graphs with arbitrary degree distributions”. In: *Physical Review E* 94.1 (2016), p. 012315.

[60] Ivan Kryven. “Finite connected components in infinite directed and multiplex networks with arbitrary degree distributions”. In: *Physical Review E* 96.5 (2017), p. 052304.

[61] Branislav Kveton et al. “Cascading bandits: Learning to rank in the cascade model”. In: *International Conference on Machine Learning*. 2015, pp. 767–776.

[62] Branislav Kveton et al. “Combinatorial cascading bandits”. In: *Advances in Neural Information Processing Systems*. 2015, pp. 1450–1458.

[63] George J Kyparisis, Sushil K Gupta, and Chi-Ming Ip. “Project selection with discounted returns and multiple constraints”. In: *European Journal of Operational Research* 94.1 (1996), pp. 87–96.

[64] Paul Lagrée, Claire Vernade, and Olivier Cappé. “Multiple-play bandits in the position-based model”. In: *Advances in Neural Information Processing Systems*. 2016, pp. 1597–1605.

[65] M Lammert. *Twelve-month evaluation of UPS diesel hybrid electric delivery vans*. Tech. rep. National Renewable Energy Laboratory (NREL), Golden, CO., 2009.

[66] Eugene L Lawler et al. *The traveling salesman problem: a guided tour of combinatorial optimization*. Vol. 3. Wiley New York, 1985.

[67] Baoxiang Li et al. “The share-a-ride problem: People and parcels sharing taxis”. In: *European Journal of Operational Research* 238.1 (2014), pp. 31–40.

[68] Lihong Li, Yu Lu, and Dengyong Zhou. “Provably optimal algorithms for generalized linear contextual bandits”. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org. 2017, pp. 2071–2080.

[69] Lihong Li et al. “A contextual-bandit approach to personalized news article recommendation”. In: *Proceedings of the 19th international conference on World wide web*. ACM. 2010, pp. 661–670.

[70] Lihong Li et al. “An unbiased offline evaluation of contextual bandit algorithms with generalized linear models”. In: *Proceedings of the Workshop on On-line Trading of Exploration and Exploitation 2*. 2012, pp. 19–36.

[71] Shuai Li and Shengyu Zhang. “Online clustering of contextual cascading bandits”. In: *Thirty-Second AAAI Conference on Artificial Intelligence*. 2018.

[72] Shuai Li et al. “Contextual Combinatorial Cascading Bandits.” In: *Proceedings of the 33th International Conference on Machine Learning*. Vol. 16. 2016, pp. 1245–1253.

[73] R Duncan Luce. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2012.

[74] Tomasz Łuczak. “The phase transition in the evolution of random digraphs”. In: *Journal of graph theory* 14.2 (1990), pp. 217–223.

[75] Tomasz Łuczak and Joel E Cohen. "Giant components in three-parameter random directed graphs". In: *Advances in applied probability* 24.4 (1992), pp. 845–857.

[76] Tomasz Łuczak, Boris Pittel, and John C Wierman. "The structure of a random graph at the point of the phase transition". In: *Transactions of the American Mathematical Society* 341.2 (1994), pp. 721–748.

[77] JK Mackenzie. "Sequential filling of a line by intervals placed at random and its application to linear adsorption". In: *The Journal of Chemical Physics* 37.4 (1962), pp. 723–728.

[78] Joel C Miller, Anja C Slim, and Erik M Volz. "Edge-based compartmental modelling for infectious disease spread". In: *Journal of the Royal Society Interface* 9.70 (2012), pp. 890–906.

[79] Charles J Mode. *Multitype branching processes: theory and applications*. Vol. 34. American Elsevier Pub. Co., 1971.

[80] Michael Molloy and Bruce Reed. "A critical point for random graphs with a given degree sequence". In: *Random structures & algorithms* 6.2-3 (1995), pp. 161–180.

[81] Yuichi Nagata. "New EAX crossover for large TSP instances". In: *PPSN*. Springer. 2006, pp. 372–381.

[82] Ilkka Norros and Hannu Reittu. "On a conditionally Poissonian graph process". In: *Advances in Applied Probability* 38.01 (2006), pp. 59–75.

[83] Mingdong Ou et al. "Multinomial Logit Bandit with Linear Utility Functions". In: *arXiv preprint: 1805.02971* (2018).

[84] ES Page. "The distribution of vacancies on a line". In: *Journal of the Royal Statistical Society. Series B (Methodological)* (1959), pp. 364–374.

[85] Mathew D Penrose. "The strong giant in a random digraph". In: *Journal of Applied Probability* 53.1 (2016), pp. 57–70.

[86] Ross G Pinsky. *Problems from the Discrete to the Continuous*. Springer, 2014.

[87] Jean-Yves Potvin. "Genetic algorithms for the traveling salesman problem". In: *Annals of Operations Research* 63.3 (1996), pp. 337–370.

[88] Wei Qi et al. "Shared Mobility for Last-Mile Delivery: Design, Operational Prescriptions and Environmental Impact". In: (2017).

[89] Sheng Qiang and Mohsen Bayati. "Dynamic pricing with demand covariates". In: (2016).

[90] Filip Radlinski, Robert Kleinberg, and Thorsten Joachims. "Learning diverse rankings with multi-armed bandits". In: *Proceedings of the 25th international conference on Machine learning*. ACM. 2008, pp. 784–791.

[91] Alfréd Rényi. "On a one-dimensional problem concerning random space filling". In: *Publ. Math. Inst. Hung. Acad. Sci* 3.109-127 (1958), pp. 30–36.

[92] Sidney I Resnick. *Heavy-tail phenomena: probabilistic and statistical modeling*. Springer Science & Business Media, 2007.

[93] Sidney Resnick and Gennady Samorodnitsky. "Tauberian theory for multivariate regularly varying distributions with application to preferential attachment networks". In: *Extremes* 18.3 (2015), pp. 349–367.

[94] Wansoo T Rhee and Michel Talagrand. "On the long edges in the shortest tour throughN random points". In: *Combinatorica* 12.3 (1992), pp. 323–330.

[95] Oliver Riordan. "The small giant component in scale-free random graphs". In: *Combinatorics, Probability and Computing* 14.5-6 (2005), pp. 897–938.

[96] Herbert Robbins. "Some aspects of the sequential design of experiments". In: *Herbert Robbins Selected Papers*. Springer, 1985, pp. 169–177.

[97] Mireia Roca-Riu, Miquel Estrada, and Elena Fernández. "An evaluation of urban consolidation centers through continuous analysis with non-equal market share companies". In: *Transportation Research Procedia* 12 (2016), pp. 370–382.

[98] Paat Rusmevichientong, Zuo-Jun Max Shen, and David B Shmoys. "Dynamic assortment optimization with a multinomial logit choice model and capacity constraint". In: *Operations research* 58.6 (2010), pp. 1666–1680.

[99] Paat Rusmevichientong and John N Tsitsiklis. "Linearly parameterized bandits". In: *Mathematics of Operations Research* 35.2 (2010), pp. 395–411.

[100] Gennady Samorodnitsky et al. "Nonstandard regular variation of in-degree and out-degree in the preferential attachment model". In: *Journal of Applied Probability* 53.1 (2016), pp. 146–161.

[101] Denis Sauré and Assaf Zeevi. "Optimal dynamic assortment planning with demand learning". In: *Manufacturing & Service Operations Management* 15.3 (2013), pp. 387–404.

[102] David G Schaeffer and John W Cain. *Ordinary Differential Equations: Basics and Beyond*. Vol. 65. Springer, 2016.

[103] Sven Schmit and Ramesh Johari. "Learning with Abandonment". In: *International Conference on Machine Learning*. 2018, pp. 4516–4524.

[104] Eugene Seneta. *Non-negative matrices and Markov chains*. Springer Science & Business Media, 2006.

[105] Zuo-Jun Max Shen and Lian Qi. "Incorporating inventory and routing costs in strategic location models". In: *European journal of operational research* 179.2 (2007), pp. 372–389.

[106] Indrajit Sinha and Thomas Foscht. "Over-marketing and brand suicide". In: *Reverse Psychology Marketing*. Springer, 2007, pp. 23–50.

[107] Aleksandrs Slivkins, Filip Radlinski, and Sreenivas Gollapudi. "Ranked bandits in metric spaces: learning diverse rankings over large document collections". In: *Journal of Machine Learning Research* 14.Feb (2013), pp. 399–436.

[108] Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*. MIT press, 1998.

[109] Kalyan Talluri and Garrett Van Ryzin. "Revenue management under a general discrete choice model of consumer behavior". In: *Management Science* 50.1 (2004), pp. 15–33.

[110] Gabor Timár et al. "Mapping the structure of directed networks: Beyond the bow-tie diagram". In: *Physical review letters* 118.7 (2017), p. 078301.

[111] Paolo Toth and Daniele Vigo. *The vehicle routing problem*. SIAM, 2002.

[112] Kenneth E Train. *Discrete choice methods with simulation*. Cambridge university press, 2009.

[113] Joel A Tropp. *User-friendly tail bounds for matrix martingales*. Tech. rep. California Institute of Tech Pasadena, 2011.

[114] Remco Van Der Hofstad. *Random graphs and complex networks*. 2016.

[115] Yana Volkovich, Nelly Litvak, and Bert Zwart. "Extremal dependencies and rank correlations in power law networks". In: *International Conference on Complex Sciences*. Springer. 2009, pp. 1642–1653.

[116] Shi Zong et al. "Cascading bandits for large-scale recommendation problems". In: *arXiv preprint arXiv:1603.05359* (2016).

Appendix A

CVRP and TSP related results

The following result taken from [48] (see also [34]) provides upper and lower bounds for the length of an optimal CVRP route to deliver n packages with arbitrary destinations $\mathbf{x}^{(n)} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\} \subseteq \mathbb{R}^2$.

Theorem 8 *For any set $\mathbf{x}^{(n)}$ of demand points serviced by a fleet of vehicles with capacity V that originate from a single depot, we have*

$$\max\left\{\frac{2n\bar{r}(n)}{V}, L(TSP(\mathbf{x}^{(n)}))\right\} \leq L(CVRP(\mathbf{x}^{(n)})) \leq 2\left(\frac{n}{V}+1\right)\bar{r}(n) + L(TSP(\mathbf{x}^{(n)})),$$

where $\bar{r}(n)$ is the empirical average distance, i.e., $\bar{r}(n) = \frac{1}{n}\sum_{i=1}^{n} r_i$.

The following result taken from [55] (see also [44]) gives an asymptotic upper bound for the length of an optimal TSP route through the points $\mathbf{x}^{(n)}$.

Theorem 9 *For any sequence of points $\mathbf{x}^{(n)}$ contained in a compact region $R \subseteq \mathbb{R}^2$, and any of the L_p distances on $\mathbb{R}^2$, there exists a constant $\alpha_{TSP}(p)$ satisfying*

$$\limsup_{n\to\infty} \frac{L(TSP(\mathbf{X}^{(n)}))}{\sqrt{nA}} \leq \alpha_{TSP}(p),$$

where A is the area of R.

Note that the statement appearing in [44, 55] is stated only for $p = 2$, however, the general case can be obtained from $\alpha_{\mathrm{TSP}}(2)$ by noting that $d_q(\mathbf{x},\mathbf{y}) \leq d_2(\mathbf{x},\mathbf{y}) \leq 2^{1/2-1/p} d_p(\mathbf{x},\mathbf{y})$ for $1 \leq p \leq 2 \leq q$.

Appendix B

Proofs in Section 2

B.1 Notation.

- N: the total number of messages
- X: the set of N messages
- M: the maximum possible number of messages sent to a single user
- $\pounds$: the maximum possible number of messages sent to a single user per time
- W: user's tolerance towards irrelevant content in terms of the number of rejected messages
- f_j: the probability of abandonment after rejecting the j^{th} messages, conditional on a user's tolerance is at least j
- F: the cumulative density function of W
- $\mathbf{S}^l$: a sequence of messages sent to user l
- $\mathbf{S}^*$: the optimal sequence
- u_i: the attraction probability of message i
- $U(\mathbf{S}, \mathbf{u}, F)$: the payoff from a sequence of messages $\mathbf{S}$ with the underlying parameters as $\mathbf{u}$ and F
- $I_l(j)$: the index of the j^{th} message sent to user l
- O_l: the set of messages which has been examined by user l
- $T_i(t)$: the total number of users who have examined message i by time t

- $\tilde{T}_j(t)$: the total number of users who have rejected the j^{th} unsatisfactory messages by time t
- $Q_i(t)$: the total number of users who accepted message i by time t
- $n_j^a(t)$: the number of users who have abandoned the platform upon receiving the j^{th} unsatisfactory message by time t
- $n_j^e(t)$:the number of users who did not abandon after rejecting j messages by time t
- m_l: the number of messages examined by user l
- $\mathfrak{M}_l$: the timestamps that user l receives messages from the platform
- ρ_l^j: the timestamp of sending the j^{th} message to user l
- $\chi_{i,t}$: the total number of times that message i is sent to users at time t
- Ω_j: the set of timestamps when the user faces the choice of abandonment after rejecting j messages
- $g(t, j)$: the index of user who rejects the j^{th} message at time t
- $\mathcal{A}_t$: the indices of the messages sent at time t

B.2 Model.

Example B.2.1 (An illustrative example to explain the optimization problem) *Let's say the job processing order is $3 \to 2 \to 5 \to 1$. After processing job 3 at time 1, it generates reward w_{13} and all the reward collected after job 3 will be discounted at rate p_{13}, and so on. It corresponds to placing message 3, 2, 5, 1 as the first, second, third and forth message in s sequence.*

B.3 Proofs in Section 2.3.

Proof. Proof of Theorem 4

We prove the result by contradiction. Assume the optimal sequence is $\mathbf{S}^* = (S_1, \cdots, S_m)$. Recall that I is the index function, where $I(i) = k$ if and only if $S_i = \{k\}$. If the optimal sequence is not ordered by their attraction probabilities, then there exists a neighboring pair $I(i)$ and $I(i+1)$ such that $u_{I(i)} < u_{I(i+1)}$. For ease of notation, let $x_1 = u_{I(i)}$, $x_2 = u_{I(i+1)}$,

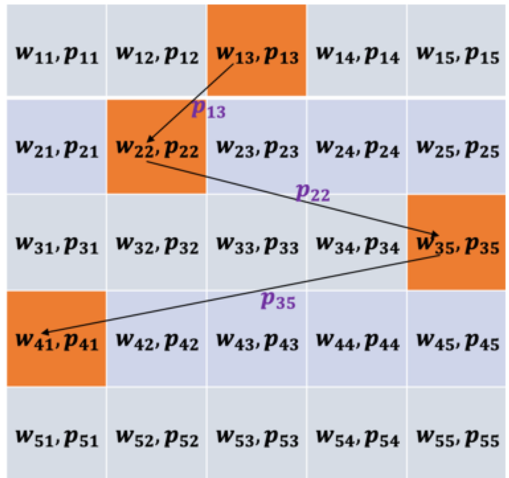

Figure B.1: An illustrative example on message selection.

$m = |\mathbf{S}^*|$. We first have

$$\begin{aligned}
&E[U((S_i, S_{i+1}, \cdots, S_m), \mathbf{u}, (f_i, \cdots, f_m))] \\
&= x_1 - c(1-x_1)f_i + (1-x_1)(1-f_i)E[U((S_{i+2}, \cdots, S_m), \mathbf{u}, (f_{i+1}, \cdots, f_m))] \\
&= x_1 - c(1-x_1)f_i + (1-x_1)(1-f_i)\left(x_2 - c(1-x_2)f_{i+1}\right. \\
&\quad \left. + (1-x_2)(1-f_{i+1})E[U((S_{i+3}, \cdots, S_m), \mathbf{u}, (f_{i+3}, \cdots, f_m))]\right),
\end{aligned}$$

and

$$\begin{aligned}
&E[U((S_{i+1}, S_i, \cdots, S_m), \mathbf{u}, (f_i, f_{i+1}, \cdots, f_m))] \\
&= x_2 - c(1-x_2)f_i + (1-x_2)(1-f_i)E[U((S_{i+2}, \cdots, S_m), \mathbf{u}, (f_{i+1}, \cdots, f_m))] \\
&= x_2 - c(1-x_2)f_i + (1-x_2)(1-f_i)\left(x_1 - c(1-x_1)f_{i+1}\right. \\
&\quad \left. + (1-x_1)(1-f_{i+1})E[U((S_{i+3}, \cdots, S_m), \mathbf{u}, (f_{i+3}, \cdots, f_m))]\right),
\end{aligned}$$

which implies that

$$\begin{aligned}
&E[U((S_i, S_{i+1}, \cdots, S_m), \mathbf{u}, (f_i, \cdots, f_m))] - E[U((S_{i+1}, S_i, \cdots, S_m), \mathbf{u}, (f_i, f_{i+1}, \cdots, f_m))] \\
&= (1 + cf_i)(x_1 - x_2) + (1-x_1)(1-f_i)x_2 - (1-x_2)(1-f_i)x_1 \\
&= f_i(1+c)(x_1 - x_2) < 0.
\end{aligned}$$

The last statement shows that swapping message $I(i)$ and $I(i+1)$ will increase the profit, which is a contradiction to the fact that $\mathbf{S}^*$ is an optimal sequence. We can similarly prove that if there exists a pair of messages i and j such that $i \in \mathbf{S}^*$, $j \notin \mathbf{S}^*$, but $u_i < u_j$, substituting i with j can improve the expected payoff. Therefore, the optimal order is to rank messages according to their attraction probabilities. ■

Proof. Proof of Theorem 5 We prove the following statement by induction: Until step j, Algorithm 1 finds the optimal sequence under the constraint for the sequence length larger than or equal to $j-1$. We first verify the first step. If $E[U(S_N, \mathbf{u}, f_N)] < 0$, then we have $E[U((S_1, \cdots, S_N), \mathbf{u}, (f_1, \cdots, f_N))] < E[U((S_1, \cdots, S_{N-1}), \mathbf{u}, (f_1, \cdots, f_{N-1}))]$, which implies that $(S_1, \cdots, S_{N-1})$ is the optimal sequence with length larger than or equal to $j-1$. Now assuming $(S_1, \cdots, S_l)$ is the optimal sequence with length larger than or equal to k where $l \geq k$, we prove for $k-1$. We claim that if $E[U((S_k, \cdots, S_l), \mathbf{u}, (f_k, \cdots, f_l))] < 0$, the optimal sequence is $(S_1, \cdots, S_{k-1})$ among all sequences with length larger than or equal to $k-1$. Otherwise, if the optimal sequence is $(S_1, \cdots, S_{m'})$ where $m' \geq k$, we can conclude that $m' = l$, otherwise it is a contradiction to the fact that $(S_1, \cdots, S_m)$ is optimal with length larger than or equal to k. However, since $E[U(S_k, \cdots, S_l), \mathbf{u}, (f_k, \cdots, f_l)] < 0$, we can conclude that $E[U((S_1, \cdots, S_{k-1}), \mathbf{u}, (f_1, \cdots, f_{k-1}))] < E[U((S_1, \cdots, S_l), \mathbf{u}, (f_1, \cdots, f_l))]$. It is a contradiction. Therefore, $(S_1, \cdots, S_{k-1})$ is the optimal sequence among all sequences with length larger than or equal to $k-1$. We can similarly prove for the scenario where $E[U(S_k, \cdots, S_l), \mathbf{u}, (f_k, \cdots, f_l)] \geq 0$. Following the induction step, we have proved that Algorithm 1 finds the optimal sequence. ∎

Proof. Proof of Proposition 2.3.2 By definition, $W_1 \gtrsim_{s.t.} W_2$ means that

$$P(W_1 > x) \geq P(W_2 > x) \text{ for all } x \in \mathbb{R}.$$

For any sequence $\mathbf{S}$ and a given abandonment distribution W, the probability of choosing message i as the l^{th} message from $\mathbf{S}$ is given by

$$\mathbb{P}_i(\mathbf{S}; W) = \prod_{k=1}^{l-1}(1 - u_{I(k)})u_i P(W \geq l).$$

Thus, we have $\mathbb{P}_i(\mathbf{S}; W_1) \geq \mathbb{P}_i(\mathbf{S}; W_2)$ for any $i \in \mathbf{S}$ when $W_1 \gtrsim_{s.t.} W_2$.

Define $\mathbb{P}_0$ as the probability that a user who finishes viewing an entire sequence does not select any message or abandon the platform, where

$$\mathbb{P}_0(\mathbf{S}; W) = \prod_{k=1}^{|\mathbf{S}|}(1 - u_{I(k)})P(W \geq |\mathbf{S}|).$$

Clearly, $\mathbb{P}_0(\mathbf{S}; W_1) \geq \mathbb{P}_0(\mathbf{S}; W_2)$ when $W_1 \gtrsim_{s.t.} W_2$. Since the probability of abandonment $\mathbb{P}_a(\mathbf{S}; W) = 1 - \sum_{i \in X} \mathbb{P}_i(\mathbf{S}; W) - \mathbb{P}_0(\mathbf{S}; W)$, we have $\mathbb{P}_a(\mathbf{S}; W_1) \leq \mathbb{P}_a(\mathbf{S}; W_2)$.

Denote $\mathbf{S}'$ and $\mathbf{S}''$ are the optimal sequences given the abandonment distribution W_1 and W_2 respectively. Because

$$E[U(\mathbf{S}, \mathbf{u}, F_W)] = \sum_{i \in X} \mathbb{P}_i(\mathbf{S}; W) - c\mathbb{P}_a(\mathbf{S}; W),$$

we have

$$E[U(\mathbf{S}', \mathbf{u}, F_{W_1})] \geq E[U(\mathbf{S}'', \mathbf{u}, F_{W_1})] \geq E[U(\mathbf{S}'', \mathbf{u}, F_{W_2})].$$

∎

B.4 Proofs in Section 2.4

Proof. Proof of Lemma 2.4.2 Lemma 2.4.2 follows immediately from the fact that each sample collected for estimating u_i is i.i.d. and follows Bernoulli distribution with mean u_i. The similar argument holds for $\hat{f}_j(t)$. ■

Lemma B.4.1 (Concentration bound) *For any $T_i(t)$ and $\tilde{T}_j(t)$, we have*

$$P\left(u_{i,t}^{OP} - \sqrt{8\frac{\log t}{T_i(t)}} < u_i < u_{i,t}^{OP}\right) \geq 1 - \frac{2}{t^4} \text{ and } P\left(f_{j,t}^{OP} < f_j < f_{j,t}^{OP} + \sqrt{8\frac{\log t}{\tilde{T}_j(t)}}\right) \geq 1 - \frac{2}{t^4}.$$

Proof. Proof of Lemma B.4.1 The upper confidence bounds for attraction probabilities **u** and abandonment probabilities are defined as follows,

$$u_{i,t}^{OP} = \min\{\hat{u}_i(t) + \sqrt{2\log t/T_i(t)}, 1\}$$

and

$$\tilde{f}_{j,t}^{OP} = \max\left\{\hat{f}_j(t) - \sqrt{2\log t/\tilde{T}_j(t)}, 0\right\}, \quad f_{j,t}^{OP} = \max_{k\leq j} \quad \tilde{f}_{k,t}^{OP}.$$

Using Hoeffding's inequality, we have

$$\begin{aligned}
&P\left(u_{i,t}^{OP} < u_i\right) + P\left(u_{i,t}^{OP} > u_i + 2\sqrt{2\log t/T_i(t)}\right) \\
&= P\left(\hat{u}_i(t) + \sqrt{2\log t/T_i(t)} < u_i\right) + P\left(\hat{u}_i(t) > u_i + \sqrt{2\log t/T_i(t)}\right) \\
&= P\left(|\hat{u}_i(t) - u_i| > \sqrt{2\log t/T_i(t)}\right) \leq 2\exp(-4\log t) = \frac{2}{t^4},
\end{aligned}$$

which implies that

$$P\left(u_{i,t}^{OP} - \sqrt{\frac{8\log t}{T_i(t)}} < u_i < u_{i,t}^{OP}\right) \geq 1 - \frac{2}{t^4}.$$

Define $k = \text{argmax}_{k'} \tilde{f}_{k',t}^{OP}$, then we have

$$\begin{aligned}
&P\left(f_{j,t}^{OP} > f_j\right) + P\left(f_{j,t}^{OP} < f_j - 2\sqrt{2\log t/\tilde{T}_j(t)}\right) \\
&\leq P\left(\tilde{f}_{k,t}^{OP} > f_k\right) + P\left(\tilde{f}_{j,t}^{OP} < f_j - 2\sqrt{2\log t/\tilde{T}_j(t)}\right) \\
&= P\left(\hat{f}_k(t) - \sqrt{2\log t/\tilde{T}_k(t)} > f_k\right) + P\left(\hat{f}_j(t) < f_j - \sqrt{2\log t/\tilde{T}_j(t)}\right) \\
&\leq 2\exp(-4\log t) = \frac{2}{t^4},
\end{aligned}$$

which implies that

$$P\left(f_{j,t}^{OP} < f_j < f_{j,t}^{OP} + \sqrt{\frac{8\log t}{\tilde{T}_j(t)}}\right) \geq 1 - \frac{2}{t^4}.$$

■

Proof. Proof of Lemma 2.4.4 Without loss of generality, assume $u_1 \geq u_2 \geq \cdots \geq u_N$ and $\mathbf{S}^* = (1, \cdots, m^*)$. If $I_l(j) \leq j$, the inequality holds because $u_{I_l(j)} \geq u_j$, which implies that $E[U(I_l(j), \mathbf{u}, f_j)] \geq E[U(j, \mathbf{u}, f_j)]$. Otherwise, if $I_l(j) > j$, then $u_{I_l(j)} \leq u_j$. Note that $u_{I_l(j),t}^{OP}$ is at least the j^{th} largest among $\mathbf{u}_t^{OP}$, otherwise $I_l(j)$ will not be chosen. With $\mathbf{u}_t^{OP} \geq \mathbf{u}$, we have $u_j \leq u_{I_l(j),t}^{OP}$ because the j^{th} largest value in sequence $\mathbf{u}_t^{OP}$ is larger than or equal to the j^{th} largest value in $\mathbf{u}$. Therefore, we have $u_{I_l(j)} \leq u_j \leq u_{I_l(j),t}^{OP}$. It implies that $u_j - u_{I_l(j)} \leq u_{I_l(j),t}^{OP} - u_{I_l(j)}$. Thus, we have reached the desired result. ■

Proof. Proof of Theorem 6 Define the optimal length of message as m^*. Assume the sequence offered to user l (entering at time l) is $\mathbf{S}^l$ with total message number m_l. We want to quantify the difference between the expected payoff generated by $\mathbf{S}^l$ and $\mathbf{S}^*$. Without loss of generality, we assume $\mathbf{S}^* = (1, 2, \cdots, m^*)$, i.e., $u_1 \geq u_2 \geq \cdots \geq u_N$. Define events

$$B_{i,t} = \{u_{i,t}^{OP} - \sqrt{8\log t/T_i(t)} < u_i < u_{i,t}^{OP}\} \quad \text{and} \quad E_{j,t} = \left\{f_{j,t}^{OP} < f_j < f_{j,t}^{OP} + \sqrt{8\log t/\tilde{T}_j(t)}\right\}.$$

Define $J_t = \bigcap_{i\in X} B_{i,t} \bigcap_{1\leq j\leq M} E_{j,t}$. First we note that on event $\bigcap_{j=1}^{m_l} J_{\rho_l^j-1}$, the length of $\mathbf{S}^l$ is longer than that of $\mathbf{S}^*$ based on the optimality of Algorithm 1. Therefore,

$$\begin{aligned}
&E_\pi\left[\left(U(\mathbf{S}^*, \mathbf{u}, F) - U(\mathbf{S}^l, \mathbf{u}, F)\right)\prod_{j=1}^{m_l} 1(J_{\rho_l^j-1})\right] \\
&= E_\pi\left[\left(U(\mathbf{S}^*, \mathbf{u}, F) - U((S_1^l, \cdots, S_{m^*}^l), \mathbf{u}, F)\right)\prod_{j=1}^{m_l} 1(J_{\rho_l^j-1})\right] \\
&\quad - E_\pi\left[U((S_{m^*+1}^l, \cdots, S_{m_l}^l), \mathbf{u}, (f_{m^*+1}, \cdots, f_{m_l}))\prod_{j=1}^{m_l} 1(J_{\rho_l^j-1})\right]. \qquad \text{(B.4.1)}
\end{aligned}$$

Note that $\mathbf{S}^l$ is the sequence proposed by our algorithm. To bound the difference in the expected payoff achieved under $\mathbf{S}^*$ and $\mathbf{S}^l$, we will utilize the coupling method. We couple the recommending process of $\mathbf{S}^*$ (call this process 1) and $\mathbf{S}^l$ (call this process 2) with the total number of messages m_l. For the j^{th} recommending message at time t (for $t = \rho_l^j$) to user l, set $a_1 = \min\{u_j, u_{I_l(j)}\}$ and $a_2 = \max\{u_j, u_{I_l(j)}\}$. Generate two independent uniform random variables $w_1 \sim unif[0,1]$ and $w_2 \sim unif[0,1]$. The event $w_1 < a_1$ denotes that both processes lead to a success when recommending the j^{th} message. If $w_1 \geq a_2$, both fail when recommending the j^{th} message. When $u_{I_l(j)} < u_j$, the event $a_1 \leq w_1 < a_2$ means that the j^{th} message is accepted in process 1 but gets rejected in process 2, and vice versa. If $w_2 \leq f_j$, then both abandon the platform after rejecting the current message. Otherwise, they will both get the next message unless the whole sequence has run out. Define the stopping time $\tilde{\tau}_l$ as the time that the coupling breaks when the recommendation in $\mathbf{S}^*$ with parameters $\mathbf{u}$ and F is a success (user clicks) but that in $\mathbf{S}^l$ with parameters $\mathbf{u}$ and F is a failure (user rejects). When the coupling breaks, the difference of reward between two processes is at most $1+c$. Then we have

$$E_\pi\left[(U(\mathbf{S}^*,\mathbf{u},F)-U((S_1^l,\cdots,S_{m^*}^l),\mathbf{u},F))\prod_{j=1}^{m_l}1(J_{\rho_l^j-1})\right]\le(1+c)E_\pi\left[\sum_{j=1}^{m^*}1(\tilde{\tau}_l=j)\prod_{j=1}^{m_l}1(J_{\rho_l^j-1})\right].$$

Now we consider another recommending process $\mathbf{S}^l$ with attraction probabilities $u^{OP}_{I_l(j),t-1}$ where $t=\rho_l^j$ for $j=1,\cdots,m_l$. Use the same process to couple $\mathbf{S}^l$ with parameter $u^{OP}_{I_l(j),t-1}$ and $\mathbf{S}^l$ with parameter $\mathbf{u}$ for the first m^* messages. Define τ'_l as the stopping time when the recommendation in $\mathbf{S}^l$ with $\mathbf{u}^{OP}_{t-1}$ is a success but with parameter $\mathbf{u}$ is a failure. According to Lemma 2.4.4, on the event that $\mathbf{u}^{OP}_{\rho_l^j-1}\ge\mathbf{u}$ for $j=1,\cdots,m^*$, we have

$$\begin{aligned}&E_\pi[E[(U(j,\mathbf{u},f_j)-U(I_l(j),\mathbf{u},f_j))1(\mathbf{u}^{OP}_{\rho_l^j-1}\ge\mathbf{u})|\mathcal{F}_{\rho_l^j-1}]]\\&\le(1+c)E_\pi\left[E\left[\left(u^{OP}_{I_l(j),\rho_l^j-1}-u_{I_l(j)}\right)1(\mathbf{u}^{OP}_{\rho_l^j-1}\ge\mathbf{u})\Big|\,\mathcal{F}_{\rho_l^j-1}\right]\right],\end{aligned}$$

which implies that

$$\begin{aligned}&E_\pi\left[\sum_{k=1}^{m^*}1(\tilde{\tau}_l=k)\prod_{j=1}^{m_l}1(J_{\rho_l^j-1})\right]\le E_\pi\left[\sum_{k=1}^{m^*}1(\tau'_l=k)\prod_{j=1}^{m_l}1(J_{\rho_l^j-1})\right]\\&\le E_\pi\left[\sum_{j=1}^{m^*}\left(\sum_{i=1}^{N}1(i\in S_j^l)\left(u^{OP}_{i,\rho_l^j-1}-u_i\right)\right)1(J_{\rho_l^j-1})\right].\end{aligned}$$

To prove the bound for the second part of sequence, we first note that for $j\le m_l$, we have

$$u^{OP}_{I_l(j),t-1}-c\left(1-u^{OP}_{I_l(j),t-1}\right)f^{OP}_{j,t-1}\ge0 \tag{B.4.2}$$

for $t=\rho_l^j$. Otherwise, if the above inequality does not hold, since $f^{OP}_{j,t-1}$ is increasing monotonically with j and the optimal sequence is sorted by $\mathbf{u}^{OP}_{t-1}$, we have

$$u^{OP}_{I_l^t(k),t-1}-c\left(1-u^{OP}_{I_l^t(k),t-1}\right)f^{OP}_{k,t-1}<0$$

for any $k>j$, where $I_l^t(k)$ is the index of the k^{th} message scheduled to sent to user l at time t. We want to emphasize the difference between two notation $I_l(j)$ and $I_l^t(k)$ here: $I_l(j)$ is the j^{th} message sent to the user l while $I_l^t(k)$ is the k^{th} message planned to sent to the user at time t. Note that this sequence may be updated at a later time. Therefore with attractiveness $\mathbf{u}^{OP}_{t-1}$ and abandonment distribution F^{OP}_{t-1}, we have

$$E[U((S_{I_l^t(j)},\cdots,S_{I_l^t(m_l)}),\mathbf{u}^{OP}_{t-1},(f^{OP}_{j,t-1},\cdots,f^{OP}_{m_l,t-1})))]<0,$$

In the proof for Theorem 5, we have proved that if $E[U((S_k, \cdots, S_m), \mathbf{u}, (f_k, \cdots, f_m))] < 0$, messages $S_k, S_{k+1}, \cdots, S_m$ will not be included in the optimal sequence. Thus, the recommendation process terminates at the j^{th} message, which is a contradiction, so we reach the conclusion of inequality (B.4.2). Therefore, for the j^{th} message where $j > m^*$, we have

$$\begin{aligned}
&E_\pi \left[E \left[(-U(I_l(j), \mathbf{u}, f_j)) 1(J_{\rho_l^j - 1}) | \mathcal{F}_{\rho_l^j - 1} \right] \right] \\
&\leq (1+c) E_\pi \left[\left(u^{OP}_{I_l(j), \rho_l^j - 1} - c \left(1 - u^{OP}_{I_l(j), \rho_l^j - 1} \right) f^{OP}_{j, \rho_l^j - 1} \right) - \left(u_{I_l(j)} - c \left(1 - u_{I_l(j)} \right) f_j \right) 1(J_{\rho_l^j - 1}) \right] \\
&\leq (1+c) E_\pi \left[\left(u^{OP}_{I_l(j), \rho_l^j - 1} - u_{I_l(j)} + f_j - f^{OP}_{j, \rho_l^j - 1} \right) 1(J_{\rho_l^j - 1}) \right].
\end{aligned}$$

It implies that

$$\begin{aligned}
&- E_\pi \left[U((S^l_{m^*+1}, \cdots, S^l_{m_l}), \mathbf{u}, F) \prod_{j=1}^{m_l} 1(J_{\rho_l^j - 1}) \right] \\
&\leq (1+c) E_\pi \left[\sum_{j=m^*+1}^{m_l} \left(\sum_{i=1}^{N} 1(i \in S^l_j)(u^{OP}_{i, \rho_l^j - 1} - u_i) + (f_j - f^{OP}_{j, \rho_l^j - 1}) \right) 1(J_{\rho_l^j - 1}) \right].
\end{aligned}$$

Thus, we have

$$\begin{aligned}
&E_\pi \left[(U(\mathbf{S}^*, \mathbf{u}, F) - U(\mathbf{S}^l, \mathbf{u}, F)) \prod_{j=1}^{m_l} 1(J_{\rho_l^j - 1}) \right] \\
&\leq (1+c) \left(E_\pi \left[\sum_{k=1}^{m^*} 1(\tau_l' = k) \prod_{j=1}^{m_l} 1(J_{\rho_l^j - 1}) \right] \right. \\
&\quad \left. + E_\pi \left[\sum_{j=m^*+1}^{m_l} \left(\sum_{i=1}^{N} 1(i \in S^l_j)(u^{OP}_{i, \rho_l^j - 1} - u_i) + (f_j - f^{OP}_{j, \rho_l^j - 1}) \right) 1(J_{\rho_l^j - 1}) \right] \right) \\
&\leq (1+c) E_\pi \left[\sum_{j=1}^{m_l} \left(\sum_{i=1}^{N} 1(i \in S^l_j) \left(u^{OP}_{i, \rho_l^j - 1} - u_i \right) + (f_j - f^{OP}_{j, \rho_l^j - 1}) \right) 1(J_{\rho_l^j - 1}) \right] \\
&\leq (1+c) E_\pi \left[\sum_{j=1}^{m_l} \sum_{i=1}^{N} 1(i \in S^l_j) \sqrt{8 \frac{\log(\rho_l^j - 1)}{T_i(\rho_l^j - 1)}} + \sum_{j=1}^{m_l} \sqrt{8 \frac{\log(\rho_l^j - 1)}{\tilde{T}_j(\rho_l^j - 1)}} \right].
\end{aligned}$$

Define $\chi_{i,t}$ as the total number of times that message i is sent to active users at time t.

If none of item i is recommended at time t, $\chi_{i,t} = 0$. Summing over all users, we have

$$
\begin{aligned}
&\frac{1}{1+c}\sum_{l=1}^{T} E_\pi\left[U(\mathbf{S}^*,\mathbf{u},F) - U(\mathbf{S}^l,\mathbf{u},F)\right] \\
&\le \sum_{l=1}^{T} E_\pi\left[\sum_{j=1}^{m_l}\sum_{i=1}^{N} 1(i\in S_j^l)\sqrt{8\frac{\log(\rho_l^j-1)}{T_i(\rho_l^j-1)}} + \sum_{j=1}^{m_l}\sqrt{8\frac{\log(\rho_l^j-1)}{\tilde{T}_j(\rho_l^j-1)}}\right] + \sum_{l=1}^{T} E_\pi\left[\sum_{j=1}^{m_l} 1(J^c_{\rho_l^j-1})\right] \\
&\le C_1\sqrt{\log T}\sum_{i=1}^{N} E_\pi\left[\sum_{t=1}^{T}\chi_{i,t}\sqrt{\frac{1}{T_i(t-1)}}\right] + C_1\sqrt{\log T}\sum_{j=1}^{M} E_\pi\left[\sum_{t=1}^{T} 1(\rho_l^j = t \text{ for some } l)\sqrt{\frac{1}{\tilde{T}_j(t-1)}}\right] \\
&\quad + \sum_{l=1}^{T} E_\pi\left[\sum_{j=1}^{m_l} 1(J^c_{\rho_l^j-1})\right],
\end{aligned}
$$

where C_1 is some constant. Note that $\chi_{i,t} \le M$ since the total number of users waiting for messages at time t is at most M. Thus, we have

$$T_i(t) = T_i(t-1) + \chi_{i,t}, \forall i.$$

It indicates that

$$E_\pi\left[\sum_{t=1}^{T}\chi_{i,t}\sqrt{\frac{1}{T_i(t-1)}}\right] \le M + E_\pi\left[\sum_{k=1}^{\lceil T_i(T)/M\rceil} M\sqrt{\frac{1}{kM}}\right] \le C_2 E_\pi[\sqrt{T_i(T)}].$$

Similarly we have

$$E_\pi\left[\sum_{t=1}^{T} 1(\rho_l^j = t \text{ for some } l)\sqrt{\frac{1}{\tilde{T}_j(t-1)}}\right] \le C_3 E_\pi\left[\sqrt{\tilde{T}_j(T)}\right].$$

Since $\sum_{i\in X} T_i(T) \le MT$ and $\sum_{j=1}^{M}\tilde{T}_j(T) \le MT$, we have

$$
\begin{aligned}
&C_1\sqrt{\log T}E_\pi\left[\sum_{i\in X}\sqrt{T_i(T)} + \sum_{j=1}^{M}\sqrt{\tilde{T}_j(T)}\right] \le C_4\sqrt{\log T}(\sqrt{NMT} + M\sqrt{T}) \\
&\le C_5\sqrt{NMT\log T},
\end{aligned}
$$

where the last inequality holds because $M \le N$. Applying Lemma B.4.1, we have

$$\sum_{l=1}^{T} E_\pi\left[\sum_{j=1}^{m_l} 1(J^c_{\rho_l^j-1})\right] \le ME_\pi\left[\sum_{t=1}^{T} 1(J_t^c)\right] \le M\sum_{t=1}^{T}\left(\sum_{i\in N} E_\pi\left[B^c_{i,t}\right] + \sum_{j=1}^{M} E_\pi[E^c_{j,t}]\right) \le C'.$$

Combining all the results above, we have

$$\begin{aligned}
&\sum_{l=1}^{T} E_\pi[U(\mathbf{S}^*, \mathbf{u}, F)] - E_\pi[U(\mathbf{S}^l, \mathbf{u}, F)] \\
&\leq \sum_{l=1}^{T} E_\pi \left[(U(\mathbf{S}^*, \mathbf{u}, F) - U(\mathbf{S}^l, \mathbf{u}, F)) \prod_{j=1}^{m_l} 1(J_{\rho_l^j - 1}) \right] + (1+c) \sum_{l=1}^{T} E_\pi \left[\sum_{j=1}^{m_l} 1(J^c_{\rho_l^j - 1}) \right] \\
&\leq C\sqrt{NMT \log T},
\end{aligned}$$

for some constant C. ■

B.5 Proofs in Section 2.5

Recall that $Y_{l,i}$ is the response of user l when message $I_l(i)$ is displayed. The response of the user, when message $I_l(i)$ with feature $\mathbf{z}_{I_l(i)}$ is displayed, is

$$Y_{l,i} = \mu(\mathbf{z}'_{I_l(i)}\beta^*) + \epsilon_{l,i}, \tag{B.5.1}$$

where $\epsilon_{l,i} = 1 - \mu(\mathbf{z}'_{I_l(i)}\beta^*)$ with probability $\mu(\mathbf{z}'_{I_l(i)}\beta^*)$, and $\epsilon_{l,i} = -\mu(\mathbf{z}'_{I_l(i)}\beta^*)$ with probability $1 - \mu(\mathbf{z}'_{I_l(i)}\beta^*)$. Meanwhile, $\hat{Y}_{l,j}$ is the response on abandonment of user l with feature $\mathbf{x}_l$ when she rejects message $I_l(j)$. Similarly, it can be written as

$$\hat{Y}_{l,j} = \mu(\mathbf{x}'_l\alpha^*_j) + \hat{\epsilon}_{l,j}, \tag{B.5.2}$$

where $\hat{\epsilon}_{l,j} = 1 - \mu(\mathbf{x}'_l\alpha^*)$ with probability $\mu(\mathbf{x}'_l\alpha^*)$, and $\hat{\epsilon}_{l,j} = -\mu(\mathbf{x}'_l\alpha^*)$ with probability $1 - \mu(\mathbf{x}'_l\alpha^*)$. The following Lemma proves that both $\epsilon_{l,j}$ and $\hat{\epsilon}_{l,j}$ are sub-Gaussian with parameter $\sigma = 1$.

Lemma B.5.1 *The noise $\epsilon_{l,j}$ and $\hat{\epsilon}_{l,j}$ defined in Equation* (B.5.1) *and Equation* (B.5.2) *are sub-Gaussian, i.e., for all l and any $\kappa \geq 0$ we have*

$$E[e^{\kappa\epsilon_{l,j}} | \mathcal{F}_{\rho_l^j - 1}] \leq e^{\kappa^2/2} \quad \text{and} \quad E[e^{\kappa\hat{\epsilon}_{l,j}} | \mathcal{F}_{\rho_l^j - 1}] \leq e^{\kappa^2/2},$$

where $\mathcal{F}_t$ is the filtration associated with the policy π up to time t.

Proof. Proof of Lemma B.5.1 According to the problem setup, $\epsilon_{l,j} = 1 - \mu(\mathbf{y}'_{I_l(j)}\beta^*)$ with probability $\mu(\mathbf{y}'_{I_l(j)}\beta^*)$, and $\epsilon_{l,j} = -\mu(\mathbf{y}'_{I_l(j)}\beta^*)$ with probability $1 - \mu(\mathbf{y}'_{I_l(j)}\beta^*)$, so $E[\epsilon_{l,j} | \mathcal{F}_{\rho_l^j - 1}] = 0$ and $-1 \leq \epsilon_{l,j} \leq 1$ almost surely. From Hoeffding's inequality (Lemma A.1. in [18]), we have

$$E[e^{\kappa\epsilon_{l,j}} | \mathcal{F}_{\rho_l^j - 1}] \leq e^{\kappa^2/2}, \quad \text{for all } \kappa \geq 0.$$

Similarly, we have

$$E[e^{\kappa\hat{\epsilon}_{l,j}} | \mathcal{F}_{\rho_l^j - 1}] \leq e^{\kappa^2/2}, \quad \text{for all } \kappa \geq 0.$$

■

Proof. Proof for Lemma 2.5.6 Since we randomly select the feature vectors in the first ξ_Z rounds, we have

$$\lambda_{min}\left(E\left[\sum_{l=1}^{\xi_Z} \mathbf{z}_{I_l(1)}\mathbf{z}'_{I_l(1)}\middle|\mathcal{F}_{l-1}\right]\right) = \lambda_{\min}\left(\sum_{l=1}^{\xi_Z} E[\mathbf{z}\mathbf{z}']\right) \geq \lambda_0\xi_Z.$$

Applying Lemma B.6.2, for $t \geq \xi_Z$, we have

$$\begin{aligned}
&P\left(\lambda_{min}\left(V_t\right) \leq \frac{1}{2}\lambda_0\xi_Z\right)\\
&\leq P\left(\lambda_{min}\left(\sum_{l=1}^{\xi_Z}\mathbf{z}_{I_l(1)}\mathbf{z}'_{I_l(1)}\right) \leq \frac{1}{2}\lambda_0\xi_Z \text{ and } \lambda_{min}\left(\sum_{l=1}^{\xi_Z} E[\mathbf{z}_{I_l(1)}\mathbf{z}'_{I_l(1)}|\mathcal{F}_{l-1}]\right) \geq \lambda_0\xi_Z\right)\\
&\leq d_Z\left(\frac{e}{2}\right)^{-\lambda_0\xi_Z/(2R)} = d_Z\left(\frac{e}{2}\right)^{-\lambda_0\sqrt{d_Z T}/(2R)}.
\end{aligned}$$

Fix $\delta = 1/\sqrt{T}$. Since T satisfies that

$$\frac{1}{2}\lambda_0\sqrt{d_Z T} \geq \frac{512M_\mu^2\sigma^2}{\kappa^4}\left(d_Z^2 + \frac{1}{2}\log T\right),$$

then by Theorem B.6.1 (in Appendix B.6), with probability at least $1-3\delta$, the maximum likelihood estimator satisfies, for any $\mathbf{z} \in \mathbb{B}^{d_Z}$, we have

$$|\mathbf{z}'(\hat{\beta}_t - \beta^*)| \leq \frac{3\sigma}{\kappa}\sqrt{\log(1/\delta)}\|\mathbf{z}\|_{V_t^{-1}}.$$

Applying union bound, we have

$$P\left(|\mathbf{z}_i'(\hat{\beta}_t - \beta^*)| \leq \frac{3\sigma}{\kappa}\sqrt{\log(N/\delta)}\|\mathbf{z}_i\|_{V_t^{-1}}, \forall i \in [N]\right) \geq 1 - 3\delta.$$

Now we prove for $\mathcal{E}_X^j$. Define the set of timestamps corresponding to $M_{j,t}$ as

$$\Omega_j = \{t|\tilde{T}_j(t) \leq \xi_X \text{ and } \tilde{T}_j(t) = \tilde{T}_j(t-1) + 1\},$$

which denotes the timestamps when the user faces the choice of abandonment after rejecting j messages. Define $g(t,j)$ as the index of user who rejects the j^{th} message at time t. To prove the inequality related to $\lambda_{min}(M_{j,t})$, we note that

$$\lambda_{min}\left(\sum_{t\in\Omega_j} E\left[\mathbf{x}_{g(t,j)}\mathbf{x}'_{g(t,j)}|\mathcal{F}_{t-1}\right]\right) \geq \lambda_{\min}\left(\sum_{k=1}^{\xi_X} E[\mathbf{x}_k\mathbf{x}_k' 1(W(\mathbf{x}_k) \geq j-1)]\right) \geq \lambda_0\xi_X$$

where the above inequality holds by Assumption 2.5.2.

Therefore, for t satisfying $\tilde{T}_j(t) \geq \xi_X$, applying Lemma B.6.2, we have

$$\begin{aligned}
&P\left(\lambda_{min}\left(M_{j,t}\right) \leq \frac{1}{2}\lambda_0\xi_X\right) \\
&= P\left(\lambda_{min}\left(\sum_{s\in\Omega_j} \mathbf{x}_{g(s,j)}\mathbf{x}'_{g(s,j)}\right) \leq \frac{1}{2}\lambda_0\xi_X \text{ and } \lambda_{min}\left(\sum_{s\in\Omega_j} E[\mathbf{x}_{g(s,j)}\mathbf{x}'_{g(s,j)}|\mathcal{F}_{s-1}]\right) \geq \lambda_0\xi_X\right) \\
&\leq d_X\left(\frac{e}{2}\right)^{-\lambda_0\xi_X/(2R)} = d_X\left(\frac{e}{2}\right)^{-\lambda_0\sqrt{d_X T}/(2R)}.
\end{aligned}$$

Since T satisfies that

$$\frac{1}{2}\lambda_0\sqrt{d_X T} \geq \frac{512M_\mu^2\sigma^2}{\kappa^4}\left(d_X^2 + \frac{1}{2}\log T\right),$$

then by Theorem B.6.1, when $\lambda_{min}(M_{j,t}) \geq \frac{1}{2}\lambda_0\sqrt{d_X T} \geq \frac{512M_\mu^2\sigma^2}{\kappa^4}\left(d_X^2 + \frac{1}{2}\log T\right)$, with probability at least $1 - 3/\sqrt{T}$, the maximum likelihood estimator satisfies, for any $\mathbf{x} \in \mathbb{B}^{d_X}$, we have

$$|\mathbf{x}'(\hat{\alpha}_{j,t} - \alpha^*)| \leq \frac{3\sigma}{\kappa}\sqrt{\frac{1}{2}\log T}\|\mathbf{x}\|_{M_{j,t}^{-1}}.$$

It implies that

$$\begin{aligned}
P((\mathcal{E}_{X,t}^j)^c) = E[1((\mathcal{E}_{X,t}^j)^c)] &\leq E[1((\mathcal{E}_{X,t}^j)^c)1(\mathcal{P}_{X,t})] + E[1(\mathcal{P}_{X,t}^c)] \\
&\leq 3\delta + d_X\left(\frac{e}{2}\right)^{-\lambda_0\sqrt{d_X T}/(2R)}.
\end{aligned}$$

Hence, for any $1 \leq j \leq M$,

$$P(\mathcal{E}_{X,t}^j) \geq 1 - 3\delta - d_X\left(\frac{e}{2}\right)^{-\lambda_0\sqrt{d_X T}/(2R)}.$$

Similarly, we have

$$P(\mathcal{E}_{Z,t}) \geq 1 - 3\delta - d_Z\left(\frac{e}{2}\right)^{-\lambda_0\sqrt{d_Z T}/(2R)}.$$

■

Proof. Proof for Lemma 2.5.7 For any $t \geq r+1$, define $B = \sum_{i\in\mathcal{A}_t} V_{t-1}^{-1/2}X_{i,t}(V_{t-1}^{-1/2}X_{i,t})^T$, then we have

$$\begin{aligned}
det(V_t) &= det\left(V_{t-1} + \sum_{i\in\mathcal{A}_t} X_{i,t}X'_{i,t}\right) \\
&= det(V_{t-1})det\left(I + \sum_{i\in\mathcal{A}_t} V_{t-1}^{-1/2}X_{i,t}(V_{t-1}^{-1/2}X_{i,t})^T\right) \\
&= det(V_{t-1})\prod_j(1 + \lambda_j(B)),
\end{aligned}$$

where λ_j denotes the j^{th} largest eigenvalue. Let $tr(B)$ denote the trace of matrix B. Note that

$$\lambda_j(B) \le tr(B) = \sum_{i\in\mathcal{A}_t} tr\left(V_{t-1}^{-1/2}X_{i,t}(V_{t-1}^{-1/2}X_{i,t})^T\right) = \sum_{i\in\mathcal{A}_t} \|X_{i,t}\|_{V_{t-1}^{-1}}$$
$$\le \lambda_{min}^{-1}(V_{t-1})\sum_{i\in\mathcal{A}_t}\|X_{i,t}\|^2 \le M,$$

where the last inequality holds because $|\mathcal{A}_t| \le M$, $\|X_{i,t}\|^2 \le 1$, and the condition that $\lambda_{min}(V_{r+1}) \ge 1$. Combining with the inequality $x \le 2M\log(1+x)$ for $x \in [0, M]$, we have

$$\begin{aligned}\log(det(V_t)) &= \log(det(V_{t-1})) + \sum_j \log(1+\lambda_j(B)) \\ &\ge \log(det(V_{t-1})) + \frac{1}{2M}\sum_j \lambda_j(B) \\ &= \log(det(V_{t-1})) + \frac{1}{2M}tr(B) \\ &= \log(det(V_{t-1})) + \frac{1}{2M}\sum_{i\in\mathcal{A}_t}\|X_{i,t}\|_{V_{t-1}^{-1}}.\end{aligned}$$

It implies that

$$\sum_{i\in\mathcal{A}_t}\|X_{i,t}\|_{V_{t-1}^{-1}} \le 2M\log\left(\frac{det(V_t)}{det(V_{t-1})}\right).$$

Thus, we have

$$\sum_{t=r+1}^{r+n}\sum_{i\in\mathcal{A}_t}\|X_{i,t}\|_{V_{t-1}^{-1}} \le 2M\left(\sum_{t=r+1}^{r+n}\log\left(\frac{det(V_t)}{det(V_{t-1})}\right)\right) \le 2M\log\left(\frac{det(V_{r+n})}{det(V_r)}\right).$$

Since the trace of V_{n+r} is bounded by $(n+r)M$ from above and bounded by 1 from below. Hence,

$$det(V_{n+r}) = \prod_{i=1}^d \lambda_i \le \left(\frac{(n+r)M}{d}\right)^d,$$

and

$$det(V_r) = \prod_{i=1}^d \lambda_i \ge \lambda_{min}^d(V_r) \ge 1.$$

Applying Cauchy-Schwartz inequality yields,

$$\sum_{t=r+1}^{r+n}\sum_{i\in\mathcal{A}_t}\|X_t\|_{V_t^{-1}} \le \sqrt{nM\sum_{t=r+1}^{r+n}\sum_{i\in\mathcal{A}_t}\|X_t\|_{V_t^{-1}}^2} \le \sqrt{2nM^2\log\left(\frac{det(V_{r+n})}{det(V_r)}\right)}$$
$$\le \sqrt{2nM^2 d\log\left(\frac{(n+r)M}{d}\right)}.$$

■

Our proof for Theorem 7 utilizes one result from [68], which is stated as Lemma B.6.1 in Appendix B.6 for completeness. Lemma B.6.1 presented a finite-sample version of the classical asymptotic normality results of the maximum likelihood estimator (MLE) under the assumption of sub-Gaussian of the noise ϵ_t with parameter σ, i.e., there exists $\sigma > 0$ such that for all $\kappa \geq 0$, we have $E[e^{\kappa\epsilon_t}|\mathcal{F}_{t-1}] \leq e^{\kappa^2\sigma^2/2}$.

Now we put everything together to prove the regret bound for the contextual version.

Proof. Proof of Theorem 7 Define $\mathcal{F}_t$ as the filtration associated with the policy π up to time t. Recall that

$$\mathcal{E}_{Z,t} := \{|\mathbf{z}_i'\hat{\beta}_t - \mathbf{z}_i'\beta^*| \leq \omega_{i,t}(\mathbf{z}), \forall i \in [N]\}, \forall t \in [\xi_Z + 1, T], \text{ and}$$

$$\mathcal{E}^j_{X,t} := \{|\mathbf{x}'\hat{\alpha}_{j,t} - \mathbf{x}'\alpha^*_j| \leq \tilde{\omega}_{j,t}(\mathbf{x})\}, \forall \mathbf{x} \in \mathbb{B}^{d_X}, \forall t \text{ s.t. } \tilde{T}_j(t) \geq \xi_X + 1;$$

where $\xi_Z = \sqrt{d_Z T}$ and $\xi_X = \sqrt{d_X T}$. Assume the sequence offered to user l (entering at time l) is $\mathbf{S}^l$ with m_l messages in total. Let ρ^j_l be the timestamp of sending the j^{th} messages in $\mathbf{S}^l$. Recall that $V_{s,t} = \sum_{t\in\Psi_s(T)}\sum_{i=1}^N 1(i \in \mathcal{A}_t)\mathbf{z}_{i,t}\mathbf{z}'_{i,t}$. Define

$$\mathcal{P}^s_{Z,t} = \left\{\lambda_{min}(V_{s,t}) > \frac{1}{2}\lambda_0\xi_Z\right\}, \forall t \in [\xi_Z + 1, T], \text{ and } \mathcal{P}_{Z,t} = \bigcap_{s=1}^{\nu}\mathcal{P}^s_{Z,t},$$

where $\nu = \log T$. Recall that $g(t, j)$ is the index of user who rejects the j^{th} message at time t. Following similarly from the proof for non-contextual SC-bandit, we have

$$\begin{aligned}
&E_\pi\left[\sum_{l=1}^{T}(U(\mathbf{S}^*, \mathbf{u}, F) - U(\mathbf{S}^l, \mathbf{u}, F))\prod_{k\in\mathfrak{M}_l} 1(\mathcal{E}_{Z,k})1(\mathcal{P}_{Z,k})\prod_{j=1}^{m_l} 1(\mathcal{E}^j_{X,\rho^j_l})1(\mathcal{P}^j_{X,\rho^j_l})\right]\\
&\leq (1+c)\xi_Z + (1+c)E_\pi\left[\sum_{l=\xi_Z}^{T}\sum_{j=1}^{m_l}\left(\sum_{i=1}^{N} 1(i \in S^l_j)\left(u^{OP}_{i,\rho^j_l-1} - u_i\right)\prod_{k\in\mathfrak{M}_l} 1(\mathcal{E}_{Z,k})1(\mathcal{P}_{Z,k})\right)\right]\\
&\quad + (1+c)M\xi_X + \sum_{j=1}^{M}(1+c)E_\pi\left[\sum_{t\in\Omega_j(T)}\left(f_j(\mathbf{x}_{g(t,j)}) - f^{OP}_{j,t-1}(\mathbf{x}_{g(t,j)})\right)\right) 1(\mathcal{E}^j_{X,t})1(\mathcal{P}^j_{X,t})\right]\\
&= (1+c)\xi_Z + (1+c)ME_\pi\left[\sum_{t\in\Psi_0(T)}\left(\sum_{i=1}^{N} 1(i \in \mathcal{A}_t)\left(u^{OP}_{i,t-1} - u_i\right) 1(\mathcal{E}_{Z,t})1(\mathcal{P}_{Z,t})\right)\right]\\
&\quad + (1+c)ME_\pi\left[\sum_{s=1}^{\nu}\sum_{t\in\Psi_s(T)}\left(\sum_{i=1}^{N} 1(i \in \mathcal{A}_t)\left(u^{OP}_{i,t-1} - u_i\right) 1(\mathcal{E}_{Z,t})1(\mathcal{P}_{Z,t})\right)\right]\\
&\quad + (1+c)M\xi_X + \sum_{j=1}^{M}(1+c)E_\pi\left[\sum_{t\in\Omega_j(T)}\left(f_j(\mathbf{x}_{g(t,j)}) - f^{OP}_{j,t-1}(\mathbf{x}_{g(t,j)})\right) 1(\mathcal{E}^j_{X,t})1(\mathcal{P}^j_{X,t})\right].
\end{aligned}$$

Since

$$\sum_{t\in\Psi_s(T)}\left(\sum_{i=1}^{N}1(i\in\mathcal{A}_t)\omega_{i,t}^{(s)}1(\mathcal{P}_{Z,t}^s)\right)=\sum_{t\in\Psi_s(T)}\left(\sum_{i=1}^{N}1(i\in\mathcal{A}_t)\rho_{Z,t}\|\mathbf{z}_{i,t}\|_{V_{s,t}^{-1}}1(\mathcal{P}_{Z,t}^s)\right),$$

then by Lemma 2.5.7, we have

$$\begin{aligned}&\sum_{t\in\Psi_s(T)}\left(\sum_{i=1}^{N}1(i\in\mathcal{A}_t)\rho_{Z,t}\|\mathbf{z}_{i,t}\|_{V_{s,t}^{-1}}1(\mathcal{P}_{Z,t}^s)\right)\\&\leq\sum_{t\in\Psi_s(T)}\left(\sum_{i=1}^{N}1(i\in\mathcal{A}_t)\rho_{Z,t}\|\mathbf{z}_{i,t}\|_{V_{s,t}^{-1}}1(\lambda_{min}(V_{s,t})>1)\right)\\&\leq\rho_{Z,T}M\sqrt{2d_Z\log\left(T/d_Z\right)|\Psi_s(T)|}.\end{aligned}$$

On the other hand, by the algorithm step (line 10-15), we have

$$\sum_{t\in\Psi_s(T)}\left(\sum_{i=1}^{N}1(i\in\mathcal{A}_t)\omega_{i,t}^{(s)}1(\mathcal{P}_{Z,t}^s)\right)\geq 2^{-s}\sum_{t\in\Psi_s(T)}1(\mathcal{P}_{Z,t}^s).$$

Therefore, we have

$$\begin{aligned}\sum_{t\in\Psi_s(T)}1(\mathcal{P}_{Z,t}^s)&\leq 2^s\sum_{t\in\Psi_s(T)}\left(\sum_{i=1}^{N}1(i\in\mathcal{A}_t)\omega_{i,t}^{(s)}1(\mathcal{P}_{Z,t}^s)\right)\\&\leq 2^s\rho_{Z,T}M\sqrt{2d_Z\log(T/d_Z)|\Psi_s(T)|}.\end{aligned}$$

Similarly, for any $1\leq j\leq M$, since in our setting only one user will get the j^{th} message at one time, we have

$$E_\pi\left[\sum_{t\in\Omega_j(T)}1((\mathcal{P}_{X,t}^j)^c)\right]\leq Td_X\left(\frac{e}{2}\right)^{-\lambda_0\sqrt{d_XT}/(2R)}.$$

Thus,

$$
\begin{aligned}
&(1+c)\xi_Z + (1+c)ME_\pi\left[\sum_{t\in\Psi_0(T)}\left(\sum_{i=1}^{N}1(i\in\mathcal{A}_t)\left(u_{i,t-1}^{OP}-u_i\right)1(\mathcal{E}_{Z,t})1(\mathcal{P}_{Z,t})\right)\right]\\
&\quad+(1+c)ME_\pi\left[\sum_{s=1}^{\nu}\sum_{t\in\Psi_s(T)}\left(\sum_{i=1}^{N}1(i\in\mathcal{A}_t)\left(u_{i,t-1}^{OP}-u_i\right)1(\mathcal{E}_{Z,t})1(\mathcal{P}_{Z,t}^s)\right)\right]\\
&\quad+(1+c)M\xi_X+\sum_{j=1}^{M}(1+c)E_\pi\left[\sum_{t\in\Omega_j(T)}\left(f_j(\mathbf{x}_{g(t,j)})-f_{j,t-1}^{OP}(\mathbf{x}_{g(t,j)}))\right)1(\mathcal{E}_{X,t}^j)1(\mathcal{P}_{X,t}^j)\right]\\
&\le(1+c)\left(\sum_{t\in\Psi_0(T)}\frac{2L_\mu}{\sqrt{t}}+\sum_{s=1}^{\nu}L_\mu 2^{3-s}E_\pi\left[\sum_{t\in\Psi_s(T)}1(\mathcal{P}_{Z,t}^s)\right]\right)\\
&\quad+(1+c)\sum_{j=1}^{M}E\left[\sum_{t\in\Omega_j(T)}L_\mu\tilde{w}_{j,t}(\mathbf{x}_{g(t,j)})1(\mathcal{E}_{X,t}^j)1(\mathcal{P}_{X,t}^j)\right]+(1+c)(\xi_Z+M\xi_X)\\
&\le(1+c)\left(M\sqrt{d_XT}+\sqrt{d_ZT}+2(M+1)L_\mu\sqrt{T}+8L_\mu\rho_{Z,T}M\sum_{s=1}^{\nu}E_\pi\left[\sqrt{2d_Z\log(T/d_Z)|\Psi_s(T)|}\right]\right.\\
&\quad\left.+8L_\mu\rho_{X,T}\sum_{j=1}^{M}E_\pi\left[\sqrt{2d_X\log(T/d_X)}\right]\right)\\
&\le C_1M\left(\sqrt{\log(T/d_X)}\sqrt{d_XT\log(T)}+\sqrt{\log(T/d_Z)}\sqrt{d_ZT\log(NT)}\right).
\end{aligned}
$$

According to Lemma 2.5.6, we have

$$
\begin{aligned}
&E_\pi\left[\sum_{l=1}^{T}(U(\mathbf{S}^*,\mathbf{u},F)-U(\mathbf{S}^l,\mathbf{u},F))\left(1-\prod_{k\in\mathfrak{M}_l}1(\mathcal{E}_{Z,k})1(\mathcal{P}_{Z,k})\prod_{j=1}^{m_l}1(\mathcal{E}_{X,\rho_l^j}^j)1(\mathcal{P}_{X,\rho_l^j}^j)\right)\right]\\
&\le(1+c)E_\pi\left[\sum_{t=1}^{T}M1(\mathcal{E}_{Z,t}^c)+M\sum_{s=1}^{\nu}1((\mathcal{P}_{Z,t}^s)^c)+\sum_{j=1}^{M}1((\mathcal{E}_{X,t}^j)^c)+\sum_{j=1}^{M}1((\mathcal{P}_{X,t}^j)^c)\right]\\
&\le C_2M\sqrt{T}.
\end{aligned}
$$

Combining all the results above, we have

$$
\begin{aligned}
&\sum_{l=1}^{T} E_\pi[U(\mathbf{S}^*, \mathbf{u}, F)] - E_\pi[U(\mathbf{S}^l, \mathbf{u}, F)] \\
&\le E_\pi\left[\sum_{l=1}^{T}(U(\mathbf{S}^*, \mathbf{u}, F) - U(\mathbf{S}^l, \mathbf{u}, F)) \prod_{k\in\mathfrak{M}_l} 1(\mathcal{E}_{Z,k})1(\mathcal{P}_{Z,k}) \prod_{j=1}^{m_l} 1(\mathcal{E}^j_{X,\rho_l^j})1(\mathcal{P}^j_{X,\rho_l^j})\right] \\
&\quad + E_\pi\left[\sum_{l=1}^{T}(U(\mathbf{S}^*, \mathbf{u}, F) - U(\mathbf{S}^l, \mathbf{u}, F)) \left(1 - \prod_{k\in\mathfrak{M}_l} 1(\mathcal{E}_{Z,k})1(\mathcal{P}_{Z,k}) \prod_{j=1}^{m_l} 1(\mathcal{E}^j_{X,\rho_l^j})1(\mathcal{P}^j_{X,\rho_l^j})\right)\right] \\
&\le CM\left(\sqrt{\log(T/d_X)}\sqrt{d_X T \log T} + \sqrt{\log(T/d_Z)}\sqrt{d_Z T \log(NT)}\right).
\end{aligned}
$$

■

B.6 Supplementary results

Lemma B.6.1 (Theorem 1 in [68]) *Suppose the choice model is* $Y_t = \mu(\mathbf{x}_t'\theta^*) + \epsilon_t$ *and* ϵ_t *is sub-Gaussian with parameter* σ. *Define* $A_t = \sum_{k=1}^{t-1} \mathbf{x}_k\mathbf{x}_k'$, $\delta > 0$, *and assume*

$$\lambda_{min}(A_t) \ge \frac{512 M_\mu^2 \sigma^2}{\eta^4}\left(d^2 + \log\frac{1}{\delta}\right),$$

where $\mathbf{x}_k \in \mathbb{B}^d$ *and* M_μ, η *are parameters defined in Lemma 2.5.5 and Assumption 2.5.2. Then, with probability at least* $1 - 3\delta$, *the maximum likelihood estimator, denoted as* $\hat{\theta}_t$, *satisfies, for any* $\mathbf{x} \in \mathbb{B}^d$, *that*

$$|\mathbf{x}'(\hat{\theta}_t - \theta^*)| \le \frac{3\sigma}{\eta}\sqrt{\log(1/\delta)}\|\mathbf{x}\|_{A_t^{-1}}.$$

Lemma B.6.2 (Lemma 3.1 in [113]) *Consider a finite adapted sequence* $\{\mathbf{X}_k\}$ *of positive-semidefinite matrices with dimension d, and suppose that*

$$\lambda_{max}(\mathbf{X}_k) \le R \text{ almost surely.}$$

Define the finite series

$$\mathbf{Y} := \sum_k \mathbf{X}_k \quad \text{and} \quad \mathbf{W} := \sum_k E_{k-1}\mathbf{X}_k.$$

For all $\mu \ge 0$,

$$P(\lambda_{min}(\mathbf{Y}) \le (1-\delta)\mu \text{ and } \lambda_{min}(\mathbf{W}) \ge \mu) \le d\left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^{\mu/R} \text{ for } \delta \in [0,1), \text{ and}$$

$$P(\lambda_{max}(\mathbf{Y}) \ge (1+\delta)\mu \text{ and } \lambda_{max}(\mathbf{W}) \le \mu) \le d\left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu/R} \text{ for } \delta \ge 0.$$

B.7 Numerical experiments

B.7.1 Robustness of non-contextual SC-Bandit algorithm

In this study, we investigate the robustness of Algorithm 2, by comparing how the regret changes with respect to different values of attraction probabilities $\mathbf{u}$.

We consider a setting with $N = 30$ and the cost of abandonment $c = 0.5$. The attractiveness $\mathbf{u}$ is uniformly generated from [0,0.2]. We present four experiments in this study. For the first two experiments, W, which measures a user's tolerance towards unsatisfactory content, is drawn from Poisson distribution with mean $\lambda = 8$ and $\lambda = 15$ respectively. Under this setting,

$$P(\text{abandon after } j^{th} \text{ unsatisfactory message}) = P(W = j | W \geq j) = \frac{\lambda^j e^{-\lambda}/j!}{1 - \sum_{k=1}^{j-1} \lambda^k e^{-\lambda}/k!}.$$

We define the truncated abandonment distribution as

$$f_j = \frac{\lambda^j e^{-\lambda}/j!}{1 - \sum_{k=1}^{j-1} \lambda^k e^{-\lambda}/k!}, \text{ for } j \leq 2\lambda \quad \text{and} \quad f_j = f_{[2\lambda]}, \text{ for } j > 2\lambda.$$

Fig B.2 illustrates the abandonment probability as a function of the number of prior rejected message with respect to different values of λ. The abandonment probability f_j first increases and then remains the same. It depicts one type of abandonment behavior which "plateaus out" after its initial surge.

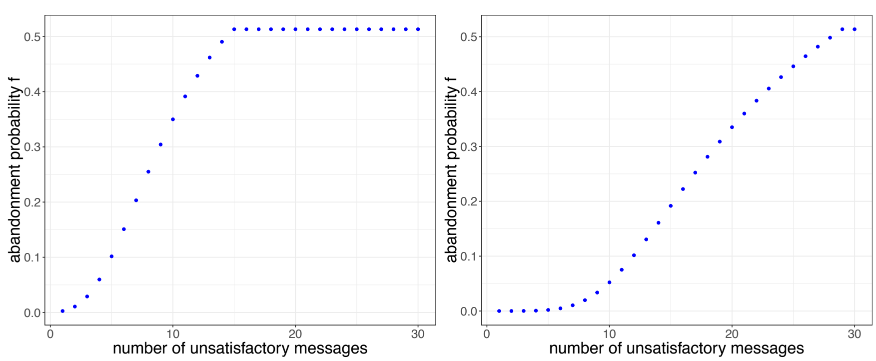

Figure B.2: Abandonment probability with respect to the number of unsatisfying message when W is drawn from truncated Poisson 1) with mean 8; 2) with mean 15.

For Experiment 3 and 4 of this study, W is drawn from geometric distribution with parameter $p = 0.2$ and 0.4 respectively. Under this setting, the abandonment probability is independent of the number of unsatisfactory messages, i.e., $f_i = p$ for all i.

Figure B.3 shows the results based on 20 independent simulations for different scenarios of $\mathbf{u}$. The average regrets are 389.26, 293.28, 323.89, and 681.93, respectively. Figure B.3

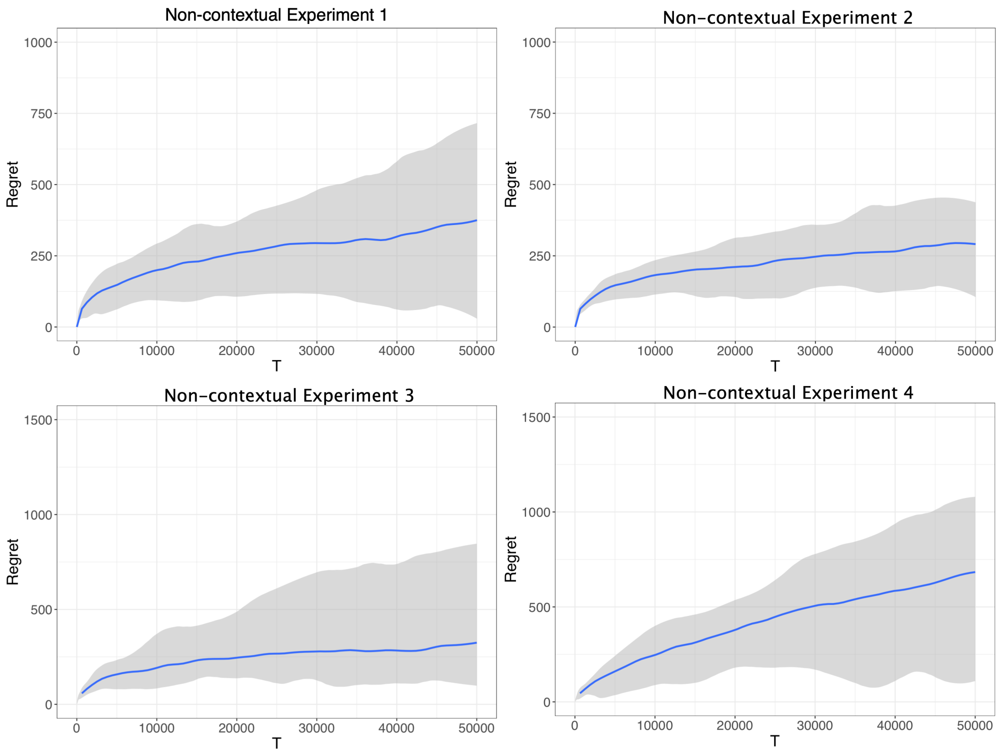

Figure B.3: Experiments on non-contextual SC-Bandit problem with Algorithm 2, with respect to different abandonment distributions.

suggests that when the abandonment's tail distribution is heavier, it is easier for the algorithm to learn the optimal ranking. The intuition is that if the user is more tolerant towards unsatisfactory messages, the platform is more likely to send out longer sequences of messages, enabling faster learning.

B.7.2 Robustness of contextual SC-Bandit algorithm

In this study, we investigate the robustness of Algorithm 3 as we perturb the coefficients α_j and β. We consider a setting with $N = 30$ and $c = 0.5$. The message contexts $\mathbf{z}_i$ include 3 features, which are uniformly generated from $[0, 1]^3$. There are also 3 features for user contexts $\mathbf{x}$, which are uniformly generated from $[0, 1]^3$.

We perform three experiments. In Experiment 1, the coefficient related to the abandonment distribution is $\alpha_j = (0.1j, 0.08j, 0.15j, -3j)$, where α_0 is the intercept. The coefficient related to the valuation of messages is $\beta = (-1, -0.8, 0.2, 0.1)$, where β_0 is the intercept. In Experiment 2, we keep β unchanged, while the coefficient $\alpha_j = (0.05j, 0.04j, 0.075j, -3j)$. In Experiment 3, α remains the same as in Experiment 1, and $\beta = (-0.5, -0.4, 0.4, 0.2)$.

Figure B.4 shows the results based on 20 independent simulations for each of the three

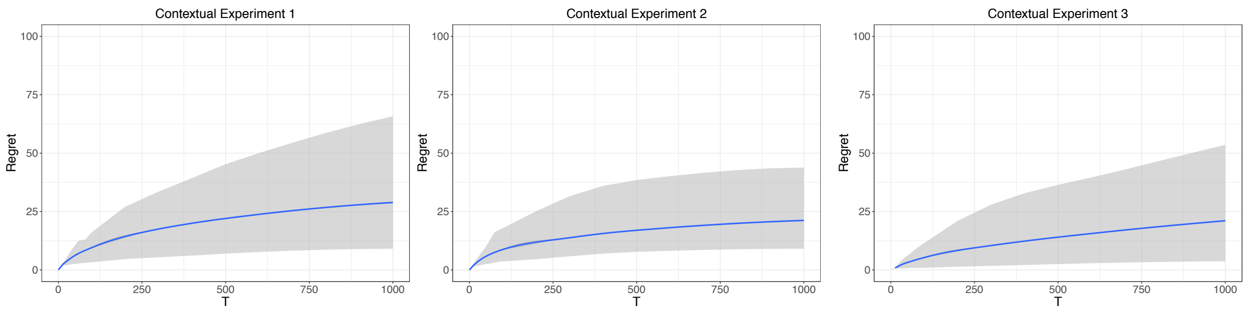

Figure B.4: Experiments on contextual SC-Bandit problem with Algorithm 3, with respect to different parameters.

experiments described above. The average regrets for the three experiments are 28.93, 21.26, and 21.04, respectively. Comparing Experiment 1 and 2, α is smaller in Experiment 2, implying that users are more tolerant of unsatisfactory messages. As a result, the platform is more likely to send out longer sequences of messages, which is translated into faster learning and a lower regret as shown in Figure B.4. By varying β in Experiment 3, the attraction probabilities of messages u_i become higher. We observe faster learning in Experiment 3, and the same phenomenon is also observed in the non-contextual experiments shown in Figure B.3.

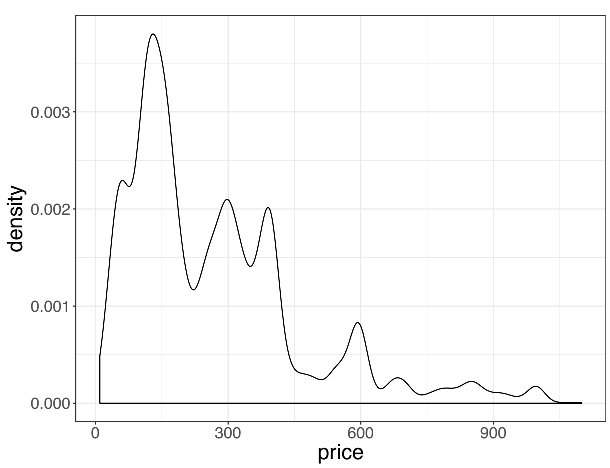

(a) Distribution of price of products in category 1665.

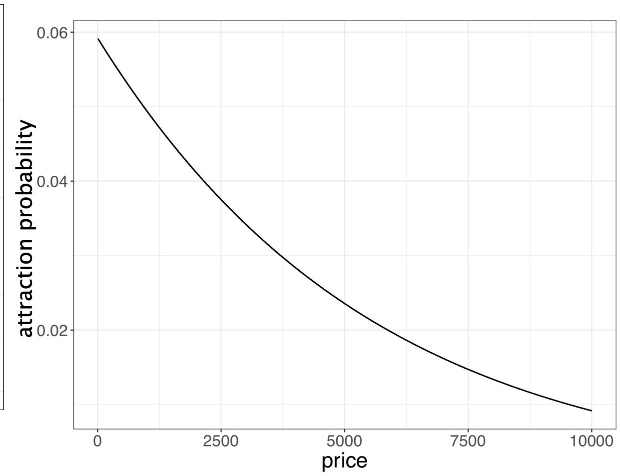

(b) The fitted attraction probability.

Figure B.5: Distribution of price and attraction probability of products.

Table B.1: Logistic Regression Result on User Abandonment Behavior Parameter $\mu(\phi(\mathbf{x}))$

	Coefficient
Gender = Male (ref = Female)	0.122***
	(0.023)
Age_level = 0 (ref = 6)	0.534
	(0.418)
Age_level = 1 (ref = 6)	−0.139
	(0.086)
Age_level = 2 (ref = 6)	−0.263***
	(0.076)
Age_level = 3 (ref = 6)	−0.318***
	(0.075)
Age_level = 4 (ref = 6)	−0.306***
	(0.075)
Age_level = 5 (ref = 6)	−0.234***
	(0.076)
Shopping_level = Shallow (ref = Deep)	0.523***
	(0.159)
Shopping_level= Medium (ref = Deep)	0.303***
	(0.046)
Intercept	−4.827***
	(0.073)
Note:	*p<0.1; **p<0.05; ***p<0.01

B.7.3 Experiments on non-contextual SC-Bandit (Real Data)

In this section, we focus on the non-contextual SC-Bandit setting. We use the estimated parameters from the data as our ground truth. We evaluate the regret for Algorithm 2 and

Table B.2: Logistic Regression Result on Product Valuation

	Coefficient
Price	$-1.918\text{e}{-04}^{***}$
	(1.547e−05)
Intercept	-2.766^{***}
	(0.005)
Note:	*p<0.1; **p<0.05; ***p<0.01

compare it with a competing algorithm which is defined next.

A benchmark algorithm for SC-Bandit Due to the novelty of our setting, there does not exist a direct benchmark in the literature. We propose a benchmark algorithm which adopts the explore-then-exploit approach, similarly to the algorithm in [101]. There is an exploration phase where every message is learned for at least $\gamma \log(t)$ times during the time period $[0, t]$, where γ is a tuning parameter. After this phase, the algorithm uses the estimated parameters to determine an optimal sequence which is offered to all subsequent users. To make the algorithm more competitive, it utilizes knowledge it has already learned to make the exploration more efficient. To be precise, suppose we want to guarantee that message i is explored during the exploration phase, we assign message i to the first tier, i.e., $S_1 = \{i\}$. Then we solve the optimal sequence problem excluding message i based on the valuations of the messages which it has already learned, and then attach them after message i. The optimization problem one needs to solve here is nearly identical to (2.3.2) with an additional constraint that $S_1 = \{i\}$.

Experiment setup We consider a setting with $N = 100$, where we select 100 products from category 1665. We use the fitted parameters β shown in Table B.1 and compute the corresponding $u_i(\mathbf{z}_i)$ as the ground truth for the product valuation. Similarly, we compute the average abandonment rate across all users as the ground truth for the tolerance, where $p = 0.64\%$. We set $c = 8.5$ in this experiment.

Experiment results Figure B.6 shows the regret comparison of Algorithm 2 and the benchmark algorithm. The average regret is 829.02 for Algorithm 2 and 1591.63 for the benchmark algorithm. Both curves grow in the sublinear shape, but the regret is almost a half in our algorithm compared to the benchmark algorithm.

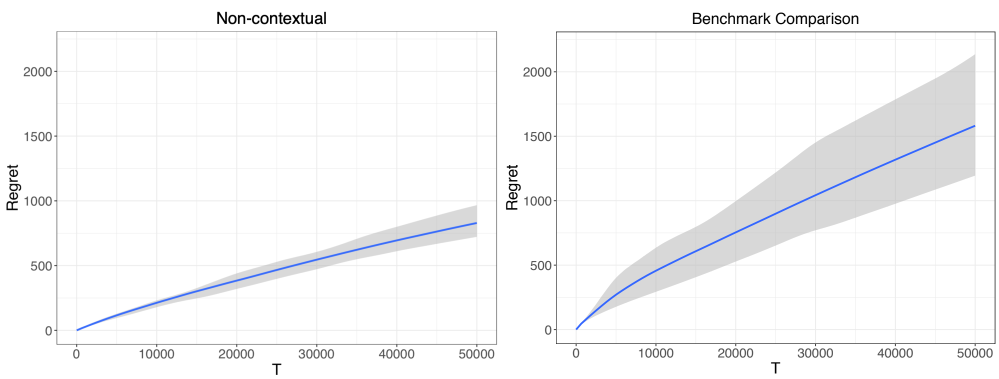

Figure B.6: Comparison of Algorithm 2 with the benchmark.