

UC Berkeley

Econ 221 - Industrial Organization Seminar

Title

Do All Markets Ultimately Tip? Experimental Evidence

Permalink

<https://escholarship.org/uc/item/4nw230qt>

Authors

Morgan, John
Hossain, Tanjim
Minor, Dylan

Publication Date

2009-05-01

Do All Markets Ultimately Tip? Experimental Evidence*

Tanjim Hossain
Department of Economics
Hong Kong University of Science & Technology

Dylan Minor
Haas School of Business
University of California, Berkeley

John Morgan
Haas School of Business and Department of Economics
University of California, Berkeley

May 2009

Abstract

Platform competition is ubiquitous, yet its outcome is little understood. Theory models typically suffer from equilibrium multiplicity—platforms might coexist or the market might tip to either platform. We use controlled laboratory experiments to study the dynamics of tipping in a class of games that included both markets with homogeneous and differentiated platforms. When platforms are primarily vertically differentiated, we find that even when platform coexistence is theoretically possible, markets inevitably tip to the efficient platform. When platforms are primarily horizontally differentiated, so there is no single efficient platform, we find strong evidence of equilibrium coexistence.

*We thank Robin Chark and Eric Set for their excellent research assistance. We also thank Syed Nageeb Ali, Mark Armstrong, Vincent Crawford, Michael Ostrovsky, Lise Vesterlund as well as seminar participants at various venues. The first author gratefully acknowledges the financial support of Hong Kong Research Grants Council. The third author gratefully acknowledges the financial support of the National Science Foundation. Please direct all correspondence to Tanjim Hossain at tanjim@ust.hk.

1 Introduction

Markets involving competition by platforms has become increasingly economically important over the last decade. The role of a platform is to match market participants of various types. Familiar platforms include the online auction site eBay and the online dating site, Match.com. However, platforms need not only match buyers to sellers or men to women (in the context of heterosexual dating sites). Video gaming consoles, such as the Wii, can also be thought of as a platform for matching game developers to gamers. The search site Google can be thought of as a platform that matches searchers with, among other things, relevant ad content provided by sellers. Credit cards, operating systems, and stock exchanges are yet other examples of platform competition.¹

There is great potential for a single dominant platform to emerge in these two-sided markets. To see why, consider an online auction context. Clearly, the more buyers that are attracted to a platform, the more valuable the platform is to sellers and, consequently, the more sellers are attracted to it. Of course, this is a virtuous circle with increasingly many buyers and sellers being attracted. This intuition, which is easily formalized, suggests that tipping (i.e. all players selecting the same platform) occurs as an equilibrium in these markets.

Yet casual observation of these markets suggests that tipping is not inevitable. Consumers enjoy more than one credit card “platform” and users seeking dates have many options besides Match.com. Theory models offer one possible solution to this apparent puzzle. In these models, multiple platforms can coexist in a given market despite the presence of positive network externalities. The reason is that users are dissuaded from switching platforms by virtue of the change in the competitive environment on the alternate platform after they switch. This can be most easily seen in an online auction context. Consider the situation of a seller on the smaller platform. By switching to the larger platform, the seller gains access to a larger number of buyers; however, owing to the seller’s presence on the larger platform, competition

¹See Armstrong (2006) and Evans and Schmalensee (2007) for many other examples of platforms or market makers.

among sellers is now slightly more fierce. The latter effect, in theory, can offset the former thus allowing both platforms to coexist.

While most of the theory models describing coexistence are static, the very notion of tipping is dynamic. One might worry about the performance of static models in real-world contexts where path dependence may be important. One might turn to field data in an attempt to examine this concern. However, studying tipping in the field is complicated by the fact tipping is a retrospective phenomenon—one only knows that a market has tipped after it has tipped.

Laboratory experiments, however, offer a unique opportunity to study the dynamics of platform competition over the “life cycle” of a market. They have the advantage that, by controlling the payoff parameters, one can turn (theoretical) coexistence on and off. They also have the advantage of allowing for a “level playing field” for the platforms; thus removing the potential confounding effect of first-mover advantages that a platform might enjoy.

In our experiments, we can easily vary the platforms’ *matching efficiency* in addition to varying the entry fees, *whereas* most theory models analyze platforms with identical matching technologies. Variations in the quality level of platforms is common in real life. The search engine Google has become a leader in bringing Internet users and advertisers to their websites in a relatively short time because of its superior search ability. The dating site eHarmony.com advertises that it uses its “Relationship Questionnaire” to create highly compatible matches based on a rigorous 29-dimension scale, thus differentiating its matching technology from those of competitors. We can also easily vary the degree of horizontal differentiation between platforms, an aspect of competition that has mainly been neglected in theory models of platform competition, but obviously an important feature of real world markets.

We offer a class of platform competition games and derive some simple theoretical properties. We then examine the performance of these games in the lab. Specifically, we conduct a series of experiments in platform competition in which subjects repeatedly participate in two-sided markets over time. Subjects choose one of two competing platforms which differ from one another in fees charged to users and may

differ in the matching technologies. In some treatments, equilibrium coexistence is possible whereas in others, only tipped equilibria arise.

Our main experimental findings are:

1. When platforms are primarily vertically differentiated, even though equilibrium coexistence is theoretically possible, in the lab, platform competition always leads to tipping.
2. Moreover, while theory is (mainly) silent as to which platform the market will tip, in the lab, the market consistently converges to the Pareto dominant platform. That is, we find little evidence of path dependent outcomes where the market gets forever locked into the “wrong” platform.
3. When platforms are primarily horizontally differentiated, so there is no Pareto dominant platform, tipping is unlikely. The markets mainly converge to the outcome predicted under equilibrium coexistence.

The paper proceeds as follows. The remainder of this section reviews the related literature on platform competition. Section 2 presents results from a simple theory model of platform competition which forms the basis for the games played in the experiment. Section 3 presents our experimental design. Section 4 presents the results of the experiments and compares these to the theory benchmarks. Section 5 focuses on how individual level dynamic decisions lead to tipping. Section 6 adds horizontal differentiation to the model and explores how this affects market dynamics. Section 7 concludes. Proofs of all theoretical results are relegated to the appendix.

1.1 Related Literature

The theoretical literature on market tipping variously explores the influence of four competing factors on the sustaining of equilibria. Although in the literature these factors bear different names, we refer to them as scale, market impact, market size, and platform fees. The “scale” effect is the increased payoff that both participant types enjoy as an equal number of both participant types are added to the platform.

This effect is simply the positive network effect of any two sided market—more of both types of participants increases the chance of successful matches (i.e., transactions), increasing expected total surplus. The “market size” effect reflects the increased payoff to an agent of one type when the number of her type is held constant, but an additional complementary agent is added to her platform. Conversely, the “market impact” effect is the reduced payoff to an agent from adding an additional agent of her own type to her platform while holding the number of agents of her complementary type constant. The former effect can be thought of as the seller in an auction market benefiting from increased demand (i.e., an additional buyer entering her platform) and the latter effect as her being penalized from additional supply (i.e., another seller joining her platform). Finally, “platform fees” are either an up-front fee or the per matching fee a platform charges its participants. These four factors together affect the expected net surplus on a particular platform, causing participants to locate on one platform over another, and thus determine equilibria.

In their pioneering work, Caillaud and Jullien (2001, 2003) explore the outcomes of competing platforms with identical matching technologies but (possibly) different platform fees. Although tipping is an equilibrium, they establish interior equilibria that are sustained by the offsetting negative effect of platform fees and the positive effect of the market size effect. Rochet and Tirole (2003) extend Caillaud and Jullien’s analysis to include heterogeneous agents, but again use the conflicting effects of platform fees and market size to find interior equilibria. Armstrong (2006) also studies platform competition with heterogeneous agents.

Ellison and Fudenberg (2003) depart from the mostly symmetric predictions of Rochet and Tirole and Caillaud and Jullien by introducing the notion of the market impact effect not present in the previous studies. This negative effect is then sufficient to support profoundly different sized platforms coexisting in equilibrium. Further, they do not assume any platform fees, which would only give their results another degree of freedom. Ellison, Fudenberg, and Möbius (2004) apply the ideas in Ellison and Fudenberg to competing online auction platforms such as eBay.² These papers

²Agents can choose only one platform (single-home) in these models.

provide a theoretical framework to potentially explain almost all platform market compositions we witness as an equilibrium. Nevertheless, they put extensive focus on interior equilibria over tipping equilibria.

There have been relatively few empirical studies about platform competition and tipping. Clemons and Weber (1996) run experiments with both students and floor traders from NYSE who chose between two stock exchanges. When one exchange was clearly more efficient, subjects coordinated on that exchange. In a recent empirical paper, Cantillon and Yin (2008) identify the factors driving German long-term government bond futures to tip from a London-based exchange to an entrant, the Frankfurt-based exchange DTB. We, however, are the first to experimentally study market dynamics and equilibrium selection in two-sided markets differing both in terms of efficiency and fee structure.

Our paper is somewhat related to the large experimental literature on coordination games with a single player type. Tension in these games arises from coordinating on Pareto dominant versus risk-dominant outcomes. In contrast to our results, the Pareto dominant equilibrium typically fares poorly in these games. For example, Van Huyck, Battalio and Beil (1990) find convergence to the risk-dominant outcome when there are a large number of players in minimum action games. Cooper, DeJong, Forsythe and Ross (1990) find little support for the Pareto dominant prediction in 2-player coordination games. In median action games, however, Van Huyck, Battalio and Beil (1991) find relatively more coordination on the Pareto dominant outcome.

2 Theory

In this section, we describe a class of platform competition games and derive a number of results describing some equilibrium properties. The class of games is flexible enough to encompass the main effects described in the extant literature on platform competition. It also provides a useful framework for the experiments—indeed all our treatments merely represent examples in this class of games. The proofs of all the results are contained in the Appendix.

Consider a platform competition game where there are N agents of each of two

types, which we refer to as squares and triangles to maintain neutrality of the setting of the market. Agents simultaneously choose to locate on one of two platforms, labeled A and B . Thus, they cannot choose to participate in multiple platforms. If an agent chooses to locate at platform i , she has to pay an up-front access fee of p_i . She earns a gross payoff of $u_i(n_1, n_2)$ where n_1 and n_2 respectively denote the number of agents of her own type and of the opposite type locating on platform i . An agent's net payoff from choosing platform i is then $u_i(n_1, n_2) - p_i$. Agents are ex-ante homogeneous and the two types are symmetric. Payoffs depend only on the platform an agent selects and numbers of her own and the complementary type that co-locate on that platform. The access fees are exogenously given and neither access fees nor gross payoffs depend directly on the agent's type.

We restrict attention to games with generic payoffs. Specifically, suppose that $p_A > p_B$, $u_i(N, N) > p_i$ and it is not the case that for all i, j, n_1 and n_2 , $u_i(n_1, n_2) - p_i = u_j(n_1, n_2) - p_j$. Finally, we make the following assumptions on gross payoff functions:

Assumption 1 (market size effect): Gross payoffs are increasing in the number of players of the opposite type. For all $n_1, n_2 \in \{1, 2, \dots, N\}$, $u_i(n_1, n_2 + 1) > u_i(n_1, n_2)$.

Assumption 2 (market impact effect): Gross payoffs are decreasing in the number of players of own type. For all $n_1, n_2 \in \{1, 2, \dots, N\}$, $u_i(n_1, n_2) > u_i(n_1 + 1, n_2)$.

Assumption 3 (scale effect): Gross payoff increase when the number of players of both types on the platform increase equally. For all $n_1, n_2 \in \{1, 2, \dots, N\}$, $u_i(n_1 + 1, n_2 + 1) > u_i(n_1, n_2)$.

Assumption 4: For all i, j , $u_j(1, 0) - p_j < u_i(N, N) - p_i$.

Assumption 4 merely rules out the possibility that an agent would prefer to be alone on a platform rather than being on a platform in which all other agents are located. With these assumptions in place, one can show the following useful property of any Nash equilibrium for this class of games.

Lemma 1 *In any equilibrium, the same number of both types select a given platform.*

This result comes from the symmetric nature of the square and triangle types.

To see this, suppose more triangle types than square types join platform A in an equilibrium. This implies that the price difference $p_B - p_A$ is large enough to offset any gain in payoff a triangle type enjoys from switching from platform A to platform B , which has more square types than triangle types. However, then a square type on platform B will benefit from switching to platform A as $p_A - p_B$ is relatively small.

Proposition 1 *Tipping is always a Nash equilibrium. Furthermore, if there exists $0 < n < N$ such that*

$$p_A - p_B \in [u_A(n+1, n) - u_B(N-n, N-n), u_A(n, n) - u_B(N-n+1, N-n)],$$

then it is a Nash equilibrium for n players of each type to choose platform A with the remainder choosing platform B .

Tipping comprises an equilibrium for the usual reasons in this model. Somewhat more interesting is the possibility of interior equilibria. These can arise for largely the same reasons as in Ellison and Fudenberg (2003). Provided the market impact effect is sufficiently strong, then agents on the smaller platform cannot benefit sufficiently from scale effects to profitably deviate.

One might worry that interior equilibria arising in this model are “knife-edge” in the sense that any small perturbation in agent strategies leads to tipping. This is not the case. For generic parameter values, there may exist many n where $p_A - p_B \in (u_A(n+1, n) - u_B(N-n, N-n), u_A(n, n) - u_B(N-n+1, N-n))$. Here, the interior equilibrium is a *strict Nash equilibrium* and hence is robust to small perturbations. In the experiments, we choose parameter values such that any interior equilibrium is strict.

Proposition 2 *There is a unique Pareto dominant equilibrium. It consists of tipping to platform i where $u_i(N, N) - p_i > u_j(N, N) - p_j$.*

While equilibrium multiplicity is generally a problem with games of platform competition, Proposition 2 shows that, by applying the Pareto refinement, one always obtains a unique prediction. Of course, there are many coordination games where the

unique Pareto dominant prediction performs poorly. In these games, applying a risk dominance refinement is often a better predictor. For the class of games we study, one can show that the risk dominance refinement excludes interior equilibria but offers no general results beyond this without imposing further restrictions on the gross payoff functions. When platforms have identical matching technology, both Pareto and risk dominance lead to the same prediction, but that may not be true when the platforms have different gross payoff functions. As we show in the next section, for the particular parameter values we select for the experiments, risk dominance does offer a unique prediction, which we use as an additional benchmark.

3 Experimental Design

We designed the experiments to operationalize the notion of different participant types choosing between platforms with varying access fees and levels of efficiency. While the theory model is static, platform competition in practice is dynamic. Individuals repeatedly choose on which platform to locate, so a platform’s market share can change over time. To gain some insight about these dynamics, we had the same set of individuals repeatedly interact in choosing platforms.

Specifically, we conducted 20 sessions of the experiment between May 2006 and March 2007. Three-hundred fifty-two undergraduate students from Hong Kong University of Science and Technology participated with none participating in more than one session. Each session took about 90 minutes including reading instructions and paying subjects. On average, a subject earned almost HKD 170 (about \$22) from participating in a session—an amount considerably above most subjects’ outside options. The experiments were programmed and conducted with the software *z-Tree* developed by Fischbacher (2007).

Each session consisted of four sets, consisting of 15 periods.³ At the beginning of a set, a participant was randomly assigned a *type* of either a “square” or a “triangle,” and randomly matched with three other players. These four players, two of each type,

³In “Homogeneous” sessions, sets consisted of 10 periods.

comprised a *market*.⁴ During each period, players in a market simultaneously chose which of two platforms, named “firm %” and “firm #,” to locate on. We informed subjects about the access fee for each platform and how much they would earn as a function of how many of each type located on each platform. These gross payoffs were presented in the form of payoff matrices. After each period, subjects learned how many of each type located on each platform, and how many points they earned. At the end of a set, each subject was randomly reassigned a new type, randomly re-matched into a new market, and shown a new set of payoffs. At the conclusion of a session, each subject was compensated based on cumulative points earned. The Appendix provides the instructions used in one of the sessions and payoff matrices used in all the sessions.

Treatments

Within each session, sets alternated as No Tip (**N**) or Tip (**T**). While tipping to either platform were Nash equilibria in all treatments, the payoffs in **N** sets additionally supported a strict Nash equilibrium in the interior. To control for presentation effects, half of the sessions began with an **N** set (referred to as an **NTNT** session) while the other half began with a **T** set (referred to as a **TNTN** session).

Platforms were either homogeneous or vertically differentiated in a given session. In homogeneous sessions, platforms had identical payoffs but different access fees. In differentiated sessions, platforms differed both in payoffs and access fees. Table 1 summarizes the treatments as well as several theoretical benchmarks. We label the platforms *A* and *B* in the remainder of the paper, where *B* denotes the platform with the cheaper access fee.

4 Market Level Results

In this section, we treat behavior at the market level as the unit of observation and analyze the evolution of market share for each platform. Our two main findings are:

⁴“Homogeneous-Large” sessions followed the same procedure but had eight-person markets with four players of each type.

Table 1: Summary of Treatments

Treatment	<i>NTNT</i> Sessions	<i>TNTN</i> Sessions	Cheap Heuristic Prediction	Risk Dominance Prediction	Pareto Dominance Prediction
Homogeneous	1, 3, and 5	2, 4, and 6	Tip to Platform <i>B</i>	Tip to Platform <i>B</i>	Tip to Platform <i>B</i>
Homogeneous-Large	7	8	Tip to Platform <i>B</i>	Tip to Platform <i>B</i>	Tip to Platform <i>B</i>
Differentiated	9 and 11	10 and 12	Tip to Platform <i>B</i>	Tip to Platform <i>B</i>	Tip to Platform <i>B</i>
Differentiated-Cheap	13 and 15	14 and 16	Tip to Platform <i>B</i>	Tip to Platform <i>A</i>	Tip to Platform <i>A</i>
Differentiated-RD	17 and 19	18 and 20	Tip to Platform <i>B</i>	Tip to Platform <i>B</i>	Tip to Platform <i>A</i>

Finding 1. Tipping to the Pareto dominant platform is pervasive.


Finding 2. Non-tipped equilibria have little impact. Markets never converge to these equilibria. However, the existence of such equilibria or vertical differentiation between platforms may reduce the speed of convergence to the Pareto dominant platform.

The remainder of the section analyzes each treatment and shows that the two findings are robust to market size and platform differentiation.

4.1 Homogeneous platforms

We first consider the case where platforms are homogeneous—equally efficient in matching agents. These are the experimental analogs to the theory models of Caillaud and Jullien (2003), Rochet and Tirole (2003), Ellison and Fudenberg (2003). For homogeneous treatments, the payoff structure as a function of the subject’s choice and the proportions of each type locating on the subject’s platform was identical for the two platforms; that is $u_i(n_1, n_2) = u_j(n_1, n_2)$ for all n_1, n_2 . However, the platforms did differ in their access fees. Both Pareto dominance and risk dominance offer the same prediction, both predict tipping to the platform with the lower access fee. A simple behavioral heuristic of choosing the cheaper platform (the so-called “Cheap Heuristic”) also predicts tipping.


Homogeneous In the homogeneous treatment, a market consists of four players—two squares and two triangles. This treatment serves as our baseline treatment. Although we are mostly interested in the market level results, we start by looking at




entire sessions first.⁵ Figure 1 presents a time series of the percentages of players choosing the cheaper platform in all the **NTNT** and **TNTN** sessions. Once a market converges to the cheaper platform, the market stays tipped there throughout the session. There is little evidence of a presentation effect, which is confirmed in the empirical analysis in Section 5.

Figure 2 displays the fraction of all markets that tipped by the end of each 10-period set, as well as to where they tipped. We say that a market has tipped to a particular platform by the end of a set if *all subjects* in that market choose that specific platform in each of the last three periods of that set. Since we ran six sessions with four markets per session, each of the bars in the figure represents twenty-four markets. Tipping is prevalent (occurring more than 90% of the time in each set) and systematic—markets only tipped to the platform with the cheaper access fee. Existence of a non-tipped equilibrium had virtually no effect on behavior. First, there were only three markets where tipping did not occur, and two of these were in Tip (**T**) sets. In other words, the frequency of tipping was (slightly) higher in the presence of an interior equilibrium. One might argue that tipping occurred because the markets were small and hence coordination was easy. Our next set of treatments complicates the coordination problem by doubling the size of the market.

⁵Recall that, four separate markets operated at the same time in each session.



Homogeneous-Large For these treatments, there were eight participants comprising a market. We also increased the length of a set to 15 periods anticipating the coordination difficulties of a larger group. Since the session-wise dynamics of platform choice are similar to the homogeneous treatment, we only present market-level behavior in the last three periods of each set. Figure 3 reproduces the analysis of Figures 2 for the Homogeneous-Large treatment and shows that every market tipped to the cheaper platform. This was not due to extending set length—even by the 10th period, all markets had tipped. Once again, the non-tipping treatment had no effect.




We were surprised to find more tipping when we increased the size of the market. Unlike Van Huyck, Battalio and Beil (1990), our experiments reach the Pareto dominant outcome even when we raise the group-size. This suggests that ease of coordination in smaller markets was not driving tipping. Of course, one might argue tipping occurred because of the focality of the “better” platform in the homogeneous case. When platforms differ in their efficiency and access fees, identifying the “better” platform is more of a challenge. To study this possibility, we next investigate markets with vertically differentiated platforms.

4.2 Differentiated Platforms


When a given number of own and other type agents receive different gross payoffs for the two platforms, we say that platforms are differentiated. A simple way in which this might occur is if one platform had a superior matching technology to the other. We model this by choosing payoffs such that $u_A(n_1, n_2) > u_B(n_1, n_2)$ for almost all (n_1, n_2) pairs with $n_1, n_2 > 0$. As before, platforms differ in their access fees. Here we were able to test whether adding a second dimension, platform quality, changes market outcomes.

Differentiated As shown in Table 1, the market tipping to the cheaper platform B is still both Pareto and risk dominant equilibrium in this treatment. Figures 4 shows that the subjects overwhelmingly chose the net payoff-dominant platform B over the platform A . Nevertheless, adding the quality dimension to platform competition slowed convergence, at least initially. After the first set, only 81% of markets converged compared with 94-100% convergence when platforms are homogeneous. From the second set onwards, however, 100% of markets converged. In every instance throughout the sessions, when a market converged it tipped to the Pareto dominant platform. Indeed, there is no evidence of platform coexistence, even when parameter values are such that an interior equilibrium exists.



Overall, we find that if a market converges to any outcome, it tips to the Payoff dominant (in net terms) platform. While we have been interpreting the results of the experiments as supporting the Pareto or risk dominant predictions with strategic players, the data is also consistent with non-strategic players who merely locate on the platform with the cheaper access fee. Our next section seeks to distinguish between these two hypotheses.


Differentiated-Cheap By varying the difference in the access fees as well as the degree of vertical differentiation, there are parameter values where the Pareto dominant platform is not the cheaper one. Thus, we can distinguish strategic behavior from the “cheap” heuristic. In these sessions we chose the gross payoffs and platform subscription fees such that market tipping to the more expensive platform is the Pareto dominant equilibrium.




The session-wise dynamics for this treatment is shown in Figure 5. Interestingly, in the first set of the **NTNT** sessions, around 75% of the subjects chose the Pareto dominant platform and 25% chose the cheap platform giving the overall market a “non-tipped” look. It is, however, instructive to examine each of the 4-player “markets” separately, as shown in Figure 6. In the first set, we find 75% of markets tipped to the Pareto dominant platform and 6% tipped to the cheap platform. Thus, at least initially, there is some evidence of market tipping to the less efficient (in net terms) platform. From set two onwards, however, 100% of markets tipped to the Pareto dominant, but more expensive, platform. Interestingly, 3 out of the 4 players from the market tipping to the cheaper platform in the first set chose the Pareto dominant platform for the beginning of the second set, after having been randomly reassigned to a new market group. As with all the previous treatments, there is no evidence of platform coexistence implied by the presence of an interior equilibrium.


None of the treatments offered so far have the flavor of “stag hunt” type games—the Pareto prediction corresponds exactly to the risk dominant prediction. Both theory and experiments suggest that when these two predictions diverge, the risk dominant prediction often prevails.⁶ Our final set of sessions seeks to differentiate between these two predictions.

⁶For example, see van Huyck, Battalio and Beil (1990) and Young (1993).



Differentiated-Risk Dominant A simple way to separate the Pareto and risk dominant predictions without disturbing the rest of the structure of the game is to increase the “upside” from mistakes on the Pareto inferior platform. To operationalize this, we simply change a single (off equilibrium) payoff cell to increase the market size effect for this platform. Since the risk dominance prediction is influenced by payoffs from mistakes while the Pareto refinement is not, this change has the effect of separating the two. In our experiments, tipping to the more expensive platform is the Pareto dominant equilibrium, while tipping to the cheaper platform is the risk dominant equilibrium.





The results are much more nuanced in this treatment. The session-wise dynamics, as seen in Figure 7, do not suggest convergence. Nevertheless, a much higher percentage of subjects chose the Pareto dominant platform at the end of each session than at the beginning. When we look at 4-player markets separately in Figure 8, we see that a majority of markets did, in fact, converge. In the first set, the majority of tipped markets converged to the risk-dominant platform. However, as subjects gained experience, tipping increasingly favored the Pareto dominant platform. By set four, 92% of markets had tipped, and, of these, 69% tipped to the Pareto dominant platform. It appears that subjects learned to select the Pareto dominant platform as they gained experience in the game. For the first time in the experiment, the market converged to a coexisting outcome one square and one triangle agent in each platform: once in an **N** set and once in a **T** set (where this outcome was not an equilibrium).

Overall, more markets tipped to the Pareto dominant platform over the four sets. Using a Pearson Chi-Squared test, we can test the null hypothesis that conditional on market tipping, there is an equal chance of a market tipping to the Pareto dominant and the risk dominant prediction. Although, we cannot reject this null hypothesis for the first three sets, we can reject it with a p -value of .07 for set four. In other words, there is modest statistical support that Pareto dominance is a better predictor

of (experienced) market tipping behavior.

Our finding that small markets within a large market can tip to different platforms is suggestive of a number of real-world situations where platform users are segmented in some way. For instance, this may be a useful description of US online dating markets which, at an aggregate level, appear to offer support for platform coexistence, but might, at a more local level, reflect mainly locally tipped markets. The US credit cards market is similar. Visa/Master cards have a larger share of the consumer credit cards while American Express dominates the corporate segment.

5 Individual Level Results

Taken together, the market level results suggest that tipping is inevitable even when theory offers the prospect of equilibrium coexistence. In general, the market systematically tips to the platform offering the greater collective surplus. Of course, the “market” is simply an aggregate of individual choices. What determines the path that these choices take? To investigate the factors driving the dynamics of tipping, we study individual choice data.

Specifically, we regress whether an individual chose the Pareto dominant platform on various explanatory variables using a longitudinal random-effects logistic regression.⁷ Our explanatory variables are: *Last Choice*, a measure of inertia, takes the value 1 if the previous period’s choice was the Pareto dominant platform and 0 otherwise. *Initial Market Share*, a measure of path dependence, is the market share of the Pareto dominant platform in a given market in the first period of a set. *Best Response*, a measure of strategy based on fictitious play dynamics, takes on a value 1 if an individual’s platform choice is the best response to actions of other players in the same market during the previous period. *Tipped Session* is simply a dummy variable which equals 1 when there is no interior equilibrium. *Sub Period*, a measure of

⁷We also used a conditional fixed effects logit, allowing a time invariant fixed effect for each individual. However, any individual that made the same choice for the duration of a set is dropped in a conditional fixed effect model (i.e., because of no choice variation). Hence, for many of the sessions up to half of the observations are lost via this method. Nonetheless, even under this radically reduced sample the coefficients are essentially the same as our random effects model and all share the same significant coefficients.

Table 2: Summary Statistics of Some Relevant Variables

		Homogeneous	Homogeneous- Large	Differentiated	Differentiated- Cheap	Differentiated- RD
Choice (PD=1)	Mean	0.97	0.98	0.97	0.95	0.50
	Std. Dev.	0.18	0.13	0.16	0.21	0.50
	Observations	3840	3840	3840	3840	3840
Initial Mkt Share	Mean	0.88	0.90	0.94	0.88	0.39
	Std. Dev.	0.20	0.16	0.15	0.20	0.25
	Observations	96	32	64	64	64
Best Response	Mean	0.94	0.99	0.96	0.94	0.50
	Std. Dev.	0.23	0.11	0.20	0.23	0.50
	Observations	3456	3584	3584	3584	3584

learning within a market, is simply a counter indicating the choice period (1 through 15) for a given set. *Set*, coded 1 through 4, is a measure of learning across markets. We interact *Tipped Session* with *Sub Period* to test whether learning is affected by the existence of an interior equilibrium. Table 2 reports summary statistics for some of the explanatory variables and the dependent variable *Choice* which equals 1 in the periods the subject chose the Pareto dominant platform.

The coefficient estimates are reported in Odds Ratio format; that is, the ratio of the probabilities of choosing versus not choosing the Pareto dominant platform. Thus, if an individual chooses the Pareto dominant platform two-thirds of the time, the odds ratio coefficient is 2—indicating 2 to 1 odds of choosing the Pareto platform. The coefficient estimates describe the change in the odds ratio from a one unit change in the regressor. Table 3 reports the coefficient estimates for each treatment. Standard errors for the actual coefficient values are in parentheses.⁸ We also denote whether the coefficients are significant using *Z*-tests.

⁸We also controlled for the presentation effect and what type the subject was assigned to. As these variables never had statistically significant coefficients, we present the regressions without those variables.

Table 3: The Determinates of Individual Platform Choice					
Dependent Variable: Choosing the Pareto Dominant Platform					
OR Coefficient	Homogeneous	Homogeneous- Large	Differentiated	Differentiated- Cheap	Differentiated- RD
Last Choice	2.29**	1.52	5.03***	6.57***	10.02***
	(.82)	(.70)	(2.08)	(2.54)	(1.13)
Initial Mkt Shr	17.88 ***	.84	20.21**	147.644***	7.86***
	(12.74)	(1.24)	(25.76)	(160.78)	(2.15)
Best Response	3.50***	2.50*	4.74***	6.62***	8.02***
	(1.06)	(1.20)	(1.71)	(2.26)	(.88)
Tipped Session	1.44	1.09	.28	.33*	.69*
	(.83)	(.92)	(.19)	(.22)	(.16)
SubPeriod	1.49***	1.45***	1.15***	1.07	1.00
	(.16)	(.13)	(.06)	(.05)	(.02)
Set	1.94***	15.44***	3.65***	3.70***	1.48***
	(.40)	(10.68)	(1.18)	(1.53)	(.08)
SubPeriod*Tip	.83	1.22	1.36**	1.29**	1.06**
	(.10)	(.24)	(.16)	(.14)	(.03)
Observations	3456	3584	3584	3584	3584
***1% Significance **5% Significance *10% Significance					

Table 3 reveals three key determinants of choice behavior: path dependence, experience, and adaptive learning.

Path Dependence

Notice that in all treatments save for Homogeneous-Large, the initial market share of a platform significantly determines subsequent individual choice.⁹ For instance, if the initial market share of the Pareto dominant platform falls by 10% in the Homogeneous treatment, Table 3 shows that the odds of choosing that platform subsequently fall by a factor of 1.8. This would seem to suggest that these markets exhibit strong first-mover advantages. A platform jumping out to an early lead enjoys persisting market share gains. Note, however, that for the bulk of the treatments, the initial market share strongly favored the eventual winning (Pareto dominant) platform; thus, making it difficult to draw conclusions as to the impact of an initial market share advantage by the weaker platform. In related work (Hossain and Morgan, 2009), we

⁹Even for the Homogeneous-Large, initial market share strongly determines subsequent choices, but is imprecisely estimated. Recall that, after set 1, every market converges to the Pareto dominant platform immediately; thus there is no variation in initial market share.

studied first-mover advantage directly in experiments similar to those reported here. We found no evidence of the persistence of initial market share on subsequent choices when the Pareto inferior platform was the first-mover.

Subject choices appear to exhibit considerable inertia save for the Homogeneous-Large treatment. The odds ratio coefficients for Last Choice in these treatments are all highly statistically and economically significant—they indicate that, in the differentiated treatments, the odds of choosing the Pareto platform increase at least fivefold when this platform is chosen in the previous period. Thus, even though we saw little evidence of the so-called QWERTY effect—tipping to the inferior platform—path dependence seems to be an important driver of subject choice.

Experience

Notice that, for all treatments, additional experience with the game strongly influences an individual's propensity to choose the Pareto dominant platform. As Table 3 shows, the odds ratio coefficients for Set are highly significant and greater than one (i.e. choice tilts towards the Pareto platform) for all treatments. Moreover, the coefficients for Sub Period are highly significant for both treatments with homogeneous platforms and the Differentiated treatment.

Adaptive Learning

From Figure 6, we found evidence that subjects were behaving strategically in the Differentiated-Cheap session. The significance of the Best Response coefficient estimates suggests that this is true for all treatments. This suggests that subjects behave “strategically” in the sense of fictitious play adaptive learning. In all differentiated sessions, the odds of choosing the Pareto dominant platform increase about fivefold when this is a best response to the play in the previous period.

Other Effects

The effects of the other regressors vary with the treatment. At the market level, we saw no evidence that the presence of an interior equilibrium had any attractive power. At the individual level, however, the speed of learning is higher in the tipped sessions of all three differentiated treatments as indicated by the coefficient estimate of the interaction term of Sub Period and Tipped Session. Thus, a coexisting equilibrium

seems to slow down the speed of convergence to the Pareto dominant platform within a set in some treatments.

6 Is Tipping Inevitable?

Our previous results suggest that tipping is an inevitable consequence of platform competition. Regardless of whether markets are large or small, whether platforms are homogeneous or differentiated, or whether there is a coexisting equilibrium or not, platform competition eventually gave way to tipping—mainly to the Pareto dominant platform.


In practice, platforms differ from one another not only vertically, but also horizontally. The “right” platform may well differ from user to user. For example, the platform Jdate.com matches individuals seeking dates. It is fairly easy to use, has reasonable rates for access, and enjoys reasonable market share. Yet there is little reason to think that the online dating market will eventually tip to Jdate.com for one simple reason—JDate.com only matches individuals who happen to be Jewish. Similarly, ChristianMingle.com specializes in matching committed heterosexual Christian singles.

From a theory standpoint, horizontal differentiation admits a new possibility—it may be that neither platform is Pareto dominant when tipped. To investigate how horizontal differentiation affects tipping, we conducted 4 additional experimental sessions with 64 undergraduate subjects from Hong Kong University of Science and Technology in March, 2009. We amended our original experimental design as follows: In each market, a pair of agents, one of each type, received a discount for choosing platform #, while the other pair received a discount for choosing platform %. The discounts reflect the idea of horizontal differentiation—each pair of types prefers to coordinate on the discounted platform.


We chose parameters such that an interior equilibrium always existed. In half the sets, the parameters were such that a tipped equilibrium was Pareto dominant. In the other half, there was no Pareto dominant tipped equilibrium. Sets alternated between these treatments. The payoff matrices used for these sessions are listed in

the appendix.

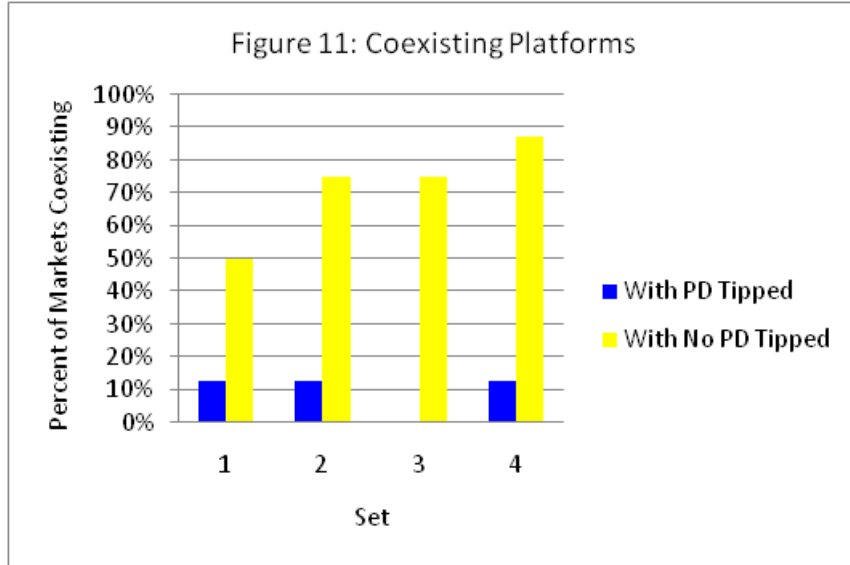
To begin, we examine the impact of horizontal differentiation when there is a Pareto dominant tipped equilibria. That is, the discounts players receive for their preferred platform do not dominate payoff difference between platforms on the vertical dimension. Figure 9 displays the results. As the figure makes clear, merely adding horizontal differentiation does not alter the broad tendency of these markets to tip to the Pareto dominant platform. In set 1, six of the eight markets converged to the Pareto dominant platform, while in sets 2-4, seven of eight converged. Below, we will account for the non-converging markets.



If we increase the degree of horizontal differentiation to the point where it dominates the vertical differentiation, this leads to a situation where there is no platform that is universally preferred. Figure 10 below displays the results for this treatment. While tipping was the norm in Figure 9, it is the exception in Figure 10. Strikingly, by the fourth set, none of the markets tipped. When neither platform Pareto dominates, the tipped equilibria lose much of their attractive power.



What happened when markets did not tip? One possibility, suggested by the results above under the RD treatment, is that these markets simply never converged at all. Another possibility is that they converged to the coexisting equilibrium. Figure 11 displays the frequency with which this occurred. Out of the five markets that did not tip to the Pareto dominant platform (in the treatment where there was such a platform), three of these converged to a coexisting equilibrium while the remaining two did not converge at all. When there was no Pareto dominant platform, most markets converged to the coexisting equilibrium. By set 4, seven out of eight markets converged to this outcome. Thus, with sufficient horizontal differentiation, tipping is not the inevitable outcome of platform competition. Instead, coexistence is the most likely outcome.



Turning to the individual level data allows us to more closely study the impact of a Pareto dominant platform on choices. Table 4, below, reproduces the analysis contained in Table 3 for sessions where platforms are horizontally differentiated. Since there is no longer a Pareto dominant platform in all treatments, our dependent variable is the choice of platform $\#$. We add two new right-hand side variables, *Pareto Dominant Platform* is a dummy which equals 1 when tipping to platform $\#$ is a Pareto dominant equilibrium. *Discounted Platform* is a dummy which equals one when platform $\#$ is player i 's discounted platform.

Table 4: Individual Choice with Horizontal Differentiation

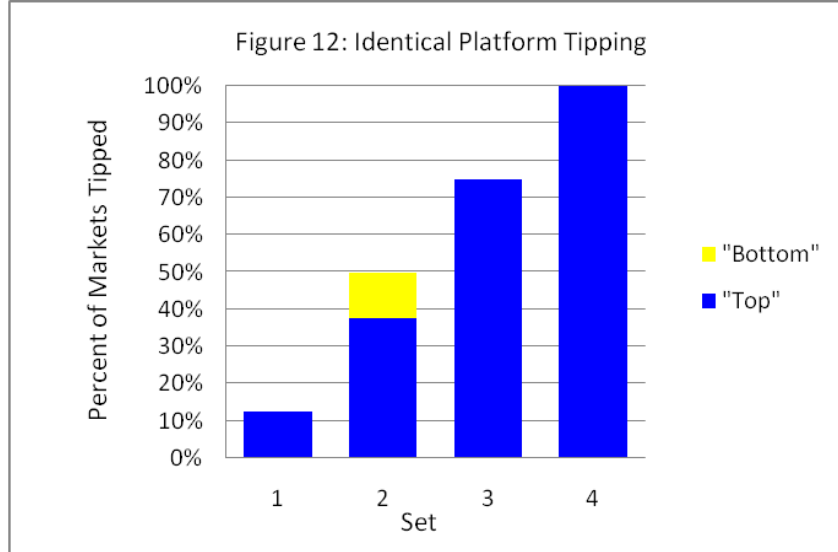
OR Coefficient	Choosing Platform #
Last Choice	22.05***(4.27)
Initial Market Share	18.83***(10.94)
Best Response	17.04***(3.04)
SubPeriod	1.01 (.02)
Set	1.28***(.11)
Discounted Platform	7.03***(1.69)
PD Tipped Platform	3.36***(.76)
Observations	3584

As before, path dependence, experience, and adaptive learning are all important features in explaining subject choices. Horizontal differentiation also plays a key role. In fact, subjects' odds ratio of choosing their discounted platform is increased by seven over the opposite platform. Similarly, the presence of a Pareto dominant platform increases the odds of its being chosen (i.e., when there is a PD tipped platform, it is platform #).

Identical Platforms

Above we considered the extreme where there was no Pareto dominant platform and found that this was sufficient to stop tipping. What happens under the other extreme—when both platforms are Pareto dominant. While generically, this cannot happen under vertical differentiation, it does provide a useful benchmark. To examine this question, we replicated the Homogeneous-Large treatment, but with equal access fees for the two platforms. As a result, the erstwhile **N** game had three coexisting equilibria while the **T** game had one coexisting equilibrium. We randomized the order in which we displayed the radio buttons for platform choice. In one session, platform "#" is on top, while in the other platform "%" is on top.

Our results may be easily summarized: Despite the existence of multiple interior equilibria, markets never converged to these outcomes. Instead, tipping was the most likely outcome. Specifically, as subjects gained experience, they learned to coordinate on whichever platform was displayed on the top of the screen. Figure 12 illustrates the pattern of tipping.



To summarize, having multiple Pareto dominant tipped equilibria does not lead to equilibrium coexistence. Rather, it leads subjects to coordinate on other features of the game to select the “winning” platform.

7 Discussion and Conclusion

A source of continuing fascination to economists is the possibility that markets will tip to an inefficient platform. Anecdotes along these lines abound, ranging from the QWERTY keyboard to the VHS format for videocassettes (see Katz and Shapiro, 1994). Underlying this worry is the simple observation that, in the presence of scale effects, tipping to either platform comprises a Nash equilibrium. This is true in our experiments as well. As we showed, however, a refinement of Nash equilibrium—the Pareto criterion—makes a unique prediction about tipping. This prediction rules out QWERTY type results. Our experiments suggest that outcomes where users get locked into the inferior platform are fairly rare and typically remedied over time. Indeed, we find no instances of convergence to the inefficient platform in the cases where the Pareto dominant platform is also risk dominant. When the two concepts diverge, some markets initially converge to the inferior (risk dominant) platform; however, this convergence disappears as subjects become more experienced. In short, our labora-

tory experiments suggest that markets inevitably tip to the efficient platform, when there is one.

Nonetheless, motivated by the observation many online platforms seem to be co-existing, recent theory models establish the possibility of equilibrium coexistence. The “glue” holding these equilibria together is the countervailing market impact effect of a user switching platforms offsetting market size and scale gains, making local deviations (strictly) unprofitable. This same “glue” is present in our experiments. Surprisingly, however, such interior equilibria are a poor description of subject behavior. Indeed, the presence or absence of an interior equilibrium has little effect on the propensity of these markets to tip to a single platform. A key difference between the theory models and the real world (as well as our experiments) is that the former are mainly static while the latter are dynamic. Dynamic considerations can allow individuals to “escape” from the inefficient interior equilibrium: while the payoffs from a one-period deviation from the interior equilibrium are negative, if an individual believes her deviation can teach the rest of the market to play the Pareto dominant equilibrium in the future, bearing such costs might be worthwhile. In other words, dynamic considerations appear to act as a “solvent” for the glue holding together the interior equilibrium in the one-shot context, leading to pervasive market tipping.

One might think that, absent QWERTY type results, our laboratory data differs strongly from the real world in terms of path dependence, which is viewed as the key force driving these outcomes. This is not the case. We find that a platform’s initial market share strongly influences subsequent platform choice decisions. When markets are only vertically differentiated, almost all markets ultimately tip to the Pareto dominant platform, if a platform’s initial market share is below average, the speed of tipping is considerably slowed. Along these same lines, while these markets never converge to an interior equilibrium, its presence modestly slows the tipping speed.

Enriching the model by allowing for horizontal as well as vertical differentiation leads to more nuanced conclusions about tipping. When the vertical dimension dominates, so there is still a more efficient platform from the perspective of all users, it

is still the case that markets overwhelmingly tip to the efficient platform. However, when the horizontal dimension is large, so there is no Pareto dominant equilibrium, platform coexistence is the most likely outcome. From an antitrust perspective, our results indicate that measuring the magnitudes of horizontal versus vertical differentiation among competing platforms is crucial for assessing the likelihood of tipping and eventual market power.

Obviously, there are a number of limitations to using our study as a basis for understanding real world two-sided markets. One limitation is that, owing to space constraints in the laboratory, our experimental markets are small relative to their real-world counterparts. Small markets might seem to bias the results in favor of tipping since coordination is easier. At the same time, however, small markets might also bias the results in favor of coexistence since the competitive impact of an additional individual on a platform or the market impact effect is likely to be more pronounced. Interestingly, when we doubled the size of the experimental market, we found more evidence of tipping in the larger market. A second potential limitation of our study is the external validity of the subject pool. In our view, undergraduates are not all that dissimilar to a typical user of platforms such as video gaming consoles, online auction markets, dating sites or search engines.

In our analysis, platforms compete on an even playing field—neither platform enjoys the first-mover advantage of an existing base of users. Tipping is often attributed to a first-mover asymmetry between platforms. In the situation of pure vertical differentiation, our conclusions are substantially unaltered by introducing this type of asymmetry. Specifically, if one amends the experimental setting to allow for an incumbent platform, the introduction of competition still quickly leads to tipping to the Pareto dominant platform (see Hossain and Morgan, 2009).

Finally, platform access fees in our model are exogenous and non-discriminatory. Clearly, neither of these conditions is true in the real world. Thus, another interesting direction for extending our results would be to designate subjects as platform operators and allow them to set fees endogenously. While clearly interesting, given the complexity of pricing even in our simple model of two-sided markets, we worried

about the external validity from such a design and opted not to pursue that path.

References

- [1] Armstrong, Mark (2006): “Competition in Two-Sided Markets,” *RAND Journal of Economics*, 37(3), 668-691.
- [2] Caillaud, Bernard and Bruno Jullien (2001): “Competing Cybermediaries,” *European Economic Review*, 45, 797-808.
- [3] Caillaud, Bernard and Bruno Jullien (2003): “Chicken and Egg: Competition among Intermediation Service Providers.” *RAND Journal of Economics*, 34(2), 521-552.
- [4] Cantillon, Estelle, and Pai-ling Yin (2008): “Competition between Exchanges: Lessons from the Battle of the Bund,” MIT Sloan School of Management, Working Paper.
- [5] Clemons, Eric K., and Bruce W. Weber (1996): “Alternative Securities Trading Systems: Tests and Regulatory Implications of the Adoption of Technology,” *Information Systems Research*, 7(2), 163-188.
- [6] Cooper, Russell W., Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross (1990): “Selection Criteria in Coordination Games: Some Experimental Results,” *American Economic Review*, 80, 218-233.
- [7] Ellison, Glenn, and Drew Fudenberg (2003): “Knife-Edge or Plateau: When Do Market Models Tip?,” *Quarterly Journal of Economics*, 118(4), 1249-1278.
- [8] Ellison, Glenn, Drew Fudenberg, and Markus Möbius (2004): “Competing Auctions,” *Journal of the European Economic Association*, 2(1), 30-66.
- [9] Evans, David S., and Richard Schmalensee (2007): *Catalyst Code: The Strategies Behind the World’s Most Dynamic Companies*, Harvard Business School Press, Boston, MA.
- [10] Fischbacher, Urs (2007): “z-Tree - Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics*, 10(2), 171-178.

- [11] Katz, Michael L. and Carl Shapiro (1994): “Systems Competition and Network Effects,” *Journal of Economic Perspectives*, 8 (2), 93-115.
- [12] Rochet, Jean-Charles and Jean Tirole (2003): “Platform Competition in Two-Sided Markets,” *Journal of the European Economic Association*, 1(4), 990-1029.
- [13] Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil (1990): “Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure,” *American Economic Review*, 80(1), 234-248.
- [14] Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil (1991): “Strategic Uncertainty, Equilibrium Selection and Coordination Failure in Average Opinion Games,” *Quarterly Journal of Economics*, 106(3), 885-910.
- [15] Young, H. Peyton, (1993): “Evolution of Conventions,” *Econometrica*, 61(1), 57-84.

A Appendix

Proof of Lemma 1

Proof. Suppose in an equilibrium n_1 triangle agents and n_2 square agents with $n_1 > n_2$ locate on platform A . For the agents of triangle type in platform A not to have an incentive to deviate requires

$$\begin{aligned} u_A(n_1, n_2) - p_A &\geq u_B(N - n_1 + 1, N - n_2) - p_B \\ \Rightarrow u_A(n_1, n_2) - u_B(N - n_1 + 1, N - n_2) &\geq p_A - p_B. \end{aligned} \quad (1)$$

Since $n_1 > n_2$, it then follows that

$$u_A(n_1, n_2) \leq u_A(n_2 + 1, n_2) < u_A(n_2 + 1, n_1) \quad (2)$$

where the weak inequality follows from Assumption 2 and the strict inequality follows from Assumption 1. Moreover,

$$u_B(N - n_1 + 1, N - n_2) \geq u_B(N - n_2, N - n_2) > u_B(N - n_2, N - n_1). \quad (3)$$

where again weak inequality follows from Assumption 2 and the strict inequality follows from Assumption 1.

Therefore, combining equations (2) and (3), we have that

$$u_A(n_2 + 1, n_1) - u_B(N - n_2, N - n_1) > u_A(n_1, n_2) - u_B(N - n_1 + 1, N - n_2).$$

Then, using equation (1), we obtain

$$u_A(n_2 + 1, n_1) - u_B(N - n_2, N - n_1) > p_A - p_B$$

which may be rewritten as

$$u_A(n_2 + 1, n_1) - p_A > u_B(N - n_2, N - n_1) - p_B.$$

But this implies that a square type agent located on platform B can profit from unilaterally deviating to platform A . This is a contradiction; therefore $n_1 = n_2$ in any equilibrium. ■

Proof of Proposition 1

Proof. Consider a platform game of size $N \geq 2$. First we show that if all agents are located at the same platform, there is no incentive to deviate. Without loss of generality, let us assume all agents are located on platform A earning net payoffs of $u_A(N, N) - p_A > 0$. If an arbitrary agent instead locates at platform B , she will be the the only agent of either type on platform B and, by Assumption 4, this is not profitable. Thus, tipping to platform A is an equilibrium. An identical argument shows that tipping to platform B is an equilibrium. Now suppose there exists an interior or non-tipped equilibrium where $n < N$ squares and triangles choose platform A and $N - n$ squares and triangles choose platform B .¹⁰ Such an equilibrium will exist if the market impact effect is strong enough to deter tipping. This just requires that there exists $n < N$ such that

$$u_A(n, n) - p_A \geq u_B(N - n + 1, N - n) - p_B$$

¹⁰Lemma 1 shows that in any coexisting equilibrium, equal number of agents of both type will be located in each platform.

and

$$u_B(N - n, N - n) - p_B \geq u_A(n + 1, n) - p_A.$$

That is, players at neither platform have any incentive to unilaterally change their locations. This also implies that there is $n < N$ such that

$$p_A - p_B \in [u_A(n + 1, n) - u_B(N - n, N - n), u_A(n, n) - u_B(N - n + 1, N - n)].$$

Here the price differential is such that unilaterally relocating to a different platform does not increase net payoff for any player. ■

Proof of Proposition 2

Proof. We first show that tipping is a necessary condition for Pareto dominance. Consider some interior equilibrium where n of each type of agent visit platform A . By Assumption 3,

$$u_A(n, n) - p_A < u_A(N, N) - p_A$$

and since tipping to platform A is also an equilibrium, this contradicts the notion that the interior equilibrium is Pareto dominant.

Thus, if a Pareto dominant equilibrium exists, it consists of tipping to one of the platforms. With generic payoffs suppose that for some i , $u_i(N, N) - p_i > u_j(N, N) - p_j$. Hence, tipping to platform i Pareto dominates tipping to platform j . Since this exhausts the set of equilibria, Pareto dominance always selects a unique equilibrium—tipping to platform i . ■

A Sample Instruction Sheet from a Homogeneous *NTNT* Session

Name: _____

Student ID: _____

Instructions

General Rules

This session is part of an experiment in the economics of decision making. If you follow the instructions carefully and make good decisions, you can earn a considerable amount of money. You will be paid in private and in cash at the end of the session.

There are sixteen people in this room who are participating in this session. They have all been recruited in the same way as you and are reading the same instructions as you are for the first time. It is important that you do not talk to any of the other participants in the room until the session is over.

The session will consist of 40 periods, in each of which you can earn points. At the end of the experiment you will be paid based on your total point earnings from all 40 periods. Each point is worth 50 cents. Thus, if you earn y points from the experiment then your total income will be HKD $y/2$. Notice that the more points you earn, the more cash you will receive.

Description of a Period

At the start of period 1, you will be randomly matched with exactly three other subjects in the room and will be designated as either a square or a triangle player. You and these three others form a “market” consisting of exactly two triangle players and two square players. During periods 1 through 10 you will be playing with the same three other people and retain the same type (square or triangle). At the start of period 11, you will be randomly matched with three other people in the room and randomly designated the types square or triangle and will play in a new market. The same thing will happen at the start of periods 21 and 31. Thus, the people with whom you are participating will change every ten periods and your type may also change.

In each period, you will decide between joining either one of two competing firms (labeled “firm %” and “firm #”). If you join firm #, you pay a subscription fee of 4 points and if you join the firm %, you pay a subscription fee of 2 points. The three other players in your market will also individually decide on which firm to join at the same time as you. On your screen, click on the firm (% or #) that you want to join. After you click “OK,” a new box will pop up to confirm that you are certain about your choice. If you want to stay with your choice, please click “yes” and click “no” otherwise. If you click “no,” you will go back to the initial box that allows you to choose one of the firms. When all the players in the market have made their decisions, you will learn your payoffs.

At the end of the period, for each firm, you will learn the number of players of each type that joined that firm in that period. Your net payoff depends on the numbers of players of each type in the firm that you join as well as that firm's subscription fee. Once you join a firm, before paying the subscription fee, in rounds 1-10, you will earn a gross payoff according to Table 1. The two columns present your gross payoffs when the number of players of your type (including yourself) in the firm you choose is 1 and 2 respectively. The three rows present your gross payoffs when the number of players of your opposite type in the firm you choose is 0, 1 or 2 respectively. You will be able to see the table on your screen during these periods.

Table 1. Gross payoffs before paying the subscription fee in periods 1-10 and 21-30

		Number of players of your own type (including yourself) in the firm you joined	
		1	2
Number of players of the opposite type in the firm you joined	0	5	5
	1	9	6
	2	12	11

The subscription fee is 2 points for firm % and 4 points for firm #. At the end of the period, you will see your net payoff (your gross payoff minus your firm's subscription fee) in points from that period. At the end of every 10 periods, you will see your net payoffs from all previous periods.

Differences between periods

At the start of period 11, your payoffs will change. Specifically, in rounds 11-20, you will earn gross payoffs (before paying the subscription fee) according to the following table:

Table 2. Gross payoffs before paying the subscription fee in periods 11-20 and 31-40

		Number of players of your own type (including yourself) in the firm you joined	
		1	2
Number of players of the opposite type in the firm you joined	0	5	5
	1	9	8
	2	12	11

Once again, you will be able to see the table on your screen during these periods. Also, remember that the subscription fee is 2 points for firm % and 4 points for firm #.

The payoffs in periods 21-30 are calculated in the same way as in periods 1-10 using Table 1. The payoffs in periods 31-40 are calculated in the same way as in periods 11-20 using Table 2.

Ending the session

At the end of period 40, you will see a screen displaying your total earnings for the experiment. Recall that, if you earn y points in total from the experiment, your total income from the experiment would be HKD $y/2$. You will be paid this amount in cash.

Payoff Matrices for Other Settings

For the remaining four settings, we present the gross payoffs for both **N** and **T** games using one table for conciseness. With differentiated platforms, the entry (u_A, u_B) lists the payoffs from platforms *A* and *B* respectively. For the outcomes where the gross payoffs are different for the two games, we present the **T** game payoffs inside parentheses.

Gross Payoffs for the *Homogeneous-Large* Treatment

The platform subscription fees were $p_A = 6$ and $p_B = 2$ in this treatment.

		Number of players of the player's own type			
		1	2	3	4
Number of players of the opposite type	0	7	7	7	7
	1	11	10	7 [9]	7 [8]
	2	13	12	7 [11]	7 [10]
	3	15	14	13	10 [12]
	4	17	16	15	14

Gross Payoffs for the *Differentiated* Treatment

The platform subscription fees were $p_A = 5$ and $p_B = 2$ in this treatment.

		Number of players of the player's own type	
		1	2
Number of players of the opposite type	0	(6, 3)	(6, 3)
	1	(10, 9)	(7 [9], 6 [8])
	2	(13, 12)	(12, 11)

Gross Payoffs for *Differentiated-Cheap* and *Differentiated-RD* Treatments

The platform subscription fees were $p_A = 3$ and $p_B = 2$ in these treatments.

		Number of players of the player's own type	
		1	2
Number of players of the opposite type	0	(4, 4)	(4, 4)
	1	(11, 8)	(8 [10], 6)
	2	(13, 11)	(12, 10)

For both **N** and **T** games, the gross payoff equals 22 for a player who is the only one of her type to choose platform *B* while both players of the other type choose platform *B* in the *Differentiated-RD* treatment instead of 11 as in the *Differentiated-Cheap* treatment.

Gross Payoffs for the *Heterogeneous Agents Treatment*

The platform subscription fees were: for one pair of square and triangle players, $p_A = 5$ and $p_B = 2$ and for the other pair of square and triangle players, $p_A = 3$ and $p_B = 4$.

		Number of players of the player's own type	
		1	2
Number of players of the opposite type	0	$(6, 5)$	$(6, 5)$
	1	$(10[11], 9[10])$	$(7[8], 6[7])$
	2	$(16[13], 12)$	$(15[12], 11)$

With the payoffs not in the square brackets, the market tipping to platform A is the Pareto dominant equilibrium. With the payoffs inside the square brackets, none of the equilibria is Pareto dominant. Specifically, the coexisting equilibrium where a player goes to her preferred platform is not Pareto dominated by any other platform.

Gross Payoffs for the *Identical Treatment*

We used the same gross payoff matrices as in the *Homogeneous-Large Treatment* in this treatment. However, the platform subscription fees were $p_A = p_B = 2$.