Transmissible Disease and Antibiotic Use in Salmon Aquaculture:
Private Incentives and Public Policy

By

THOMAS MOYLAN ANDERSON JR.

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Agricultural and Resource Economics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

James N. Sanchirico (co-chair)

Matthew N. Reimer (co-chair)

Katrina K. Jessoe

Committee in Charge

2024

Contents

Abstract

In this dissertation, I explore and empirically evaluate the management of an endemic disease in Chile's salmon aquaculture industry. Chile is the world's second-largest producer of farmed salmon and the largest supplier of salmon to the United States. However, the industry has struggled to manage endemic, transmissible diseases that threaten industry productivity. This has driven the industry to utilize antibiotics at a rate that has raised concern among public health and environmental communities.

Given the ongoing challenges from endemic pathogens, what avenues are available for balancing disease management with antibiotic stewardship? One way to address the spread of pathogens is through coordinated action between farms. In the first chapter, I evaluate a spatially explicit disease management program that forces all farm sites within the same neighborhood to coordinate their production activities. Theory suggests that this system will be most effective when transmission between neighborhoods is low because it reduces the likelihood of immediate re-infection at the start of the coordinated production cycle. I evaluate the extent to which pathogens spread between neighborhoods, exploiting exogenous variation in the timing of production cycles induced by the disease management policy. I do not find any evidence of spillover between neighborhoods, suggesting that the spatial scale of the policy is appropriately matched to the spatial scale of transmission. To account for the possibility of a behavioral response that could mitigate observed disease prevalence, I also test for changes in the propensity and intensity of antibiotic use that might be caused by the variation in pathogen pressure. However, I find no evidence of such effects.

In the second chapter, I turn toward the private incentives for disease control, specifically antibiotic use. The incentive to apply antibiotics is determined by its effectiveness at reducing losses from disease. To simulate counterfactual disease control policies, it is therefore necessary to accurately estimate the effectiveness of antibiotic treatments in a farm setting. However, several econometric challenges are associated with measuring the effectiveness of treatment using observational data. Using an epidemiological model with endogenous antibiotic applications, I illustrate how self-selection into treatment adds bias to an empirical estimate of a structural parameter critical for counterfactual simulation of disease dynamics. This bias is driven by the behavioral

feedback between treatment and stochastic features in the epidemiological model that are unobserved by the econometrician. When producers self-select into treatment based on these unobserved features, conventional models tend to underestimate treatment effectiveness. I illustrate an alternative estimator based on the control function approach and conduct a Monte Carlo experiment to illustrate its efficacy. I then apply the estimator to data from the Chilean salmon farming industry to illustrate how the estimator can be used in practice.

In the final chapter, I simulate the impacts of targeted reductions to antibiotic use to explore the tradeoffs between disease management and antibiotic stewardship. To do this, I solve for the dynamically optimal antibiotic treatment schedule in the presence of disease transmission between individuals on a farm and between farms. To capture the effect of the neighborhood coordination policy, I parameterize the disease pressure from other farms in a time-varying manner consistent with coordinated production cycles. I then simulate the impacts of various antibiotic reduction policies that feature varying levels of effectiveness and expected implementation costs. Each of the restrictions causes a non-uniform shift in the timing of treatment in the production cycle, an effect that is ultimately consequential for the disease-related externality. I find that the most cost-effective instruments for private individuals also tend to generate the largest externalities.

Each chapter is supported by data obtained from the Servicio Nacional de Pesca y Acuicultura (SERNAPESCA), the national regulatory authority for Chile's aquaculture industry. These data are collected as part of a mandatory reporting requirement and provide georeferenced, weekly observations on the production status and veterinary health of all of the industry's fish farms. Critically, these data also include prescriptions for antibiotics applied for a wide range of purposes, including the treatment of Salmonid Rickettsial Syndrome (SRS), which is caused by the endemic pathogen *Piscirickettsia salmonis*. Although these data have been used in past epidemiological studies, this dissertation is the first to utilize these data for economic analysis. In some cases, I pair this database with supplementary data, drawing upon publicly available information on policy borders and coastlines that are central to the empirical investigation in the first chapter.

Acknowledgments

Completing this dissertation has been a challenging yet incredibly rewarding journey. There are so many important people to acknowledge.

First, I want to thank my advisors, Jim Sanchirico, Matt Reimer, and Katrina Jessoe, for their guidance, patience, and mentorship throughout this process. Your insights and constructive criticisms have significantly shaped my research and academic growth. I am also deeply thankful that you gave me the freedom and support to pursue a dissertation topic that I find endlessly interesting. I hope to continue to learn from you and to pay your generosity forward to future students.

I have also received so much support from many whose names are not on the coversheet but who were, nonetheless, critical for the completion of this dissertation. I am incredibly grateful to Jim Wilen, Dale Squires, Frank Asche, and (Uncle) Jim Anderson. I'm not sure that I would have found my way into this program or pursued this line of research had it not been for your support.

I want to give a heartfelt acknowledgment to Rolando Ibarra for providing his technical expertise and good humor to this entire project. This research would not have been possible without your input and your patient responses to my endless stream of questions.

The core of this dissertation utilizes data from the Chilean salmon industry. I want to acknowledge the efforts of the Servicio Nacional de Pesca y Acuicultura (SERNAPESCA) in collecting these data and making them available for research. I also would like to thank the UC Davis Chile Center for their key role in brokering a relationship with SERNAPESCA and offering their hospitality during my visit to Chile.

I am deeply grateful to my wife, Michelle, for her unwavering support, patience, and encouragement throughout this entire process. You kept me motivated even during the most challenging times. Thank you for listening as I tried to talk through tough conceptual problems, for challenging me to become a better writer, and for pushing me away from the desk at the end of the workday.

To my parents, Nancy and Tom, your endless love and faith in me been invaluable. From the very beginning, you instilled in me the values of hard work and perseverance. Thank you for your faith in my ability to get through this program and finish this dissertation.

Finally, I would like to thank the National Oceanographic and Atmospheric Administration (NOAA) and the National Marine Fisheries Services (NMFS) for funding this research. This dissertation was supported by the California Sea Grant College Program under grant no. NA21OAR4170242, project no. M/RE-10. The statements, findings, conclusions, and recommendations are those of the author(s) and do not necessarily reflect the views of California Sea Grant, state agencies, NOAA or the U.S. Dept. of Commerce.

Introduction

Disease in crops and livestock has impacted food production since the advent of agriculture. Although ancient peoples could not see the microscopic pathogens present in their farms and cities, they were familiar with their impact. Writing in the early years of the Roman Empire, Marcus Vergilius Maro (Vergil) warned his contemporaries to watch their herds closely for animals that "listlessly nibble the top of the grass, lag in the rear, or sink while grazing in the midst of the field." If disease did arise, he urged his readers to "check the offence...before the dread disease spreads through the unwary throng [of animals]." Vergil lived centuries before the advent of differential calculus or modern epidemiology, but his recommendations reflect an intuitive understanding of infectious disease. The trajectory of a disease outbreak can be vexing and non-linear, but it is also sensitive to human behavior: disease induces the shepherd to take action, and the action influences the trajectory of the outbreak.

Our understanding of pathogens has improved dramatically since ancient Rome, yet Vergil's heuristics for vigilance and active disease management remain relevant. The global food system invests billions of dollars each year in biosecurity, monitoring, and disease prevention to preserve agricultural productivity. When disease does arrive, modern veterinarians and agronomists can draw upon a large toolbox of therapeutic products to address the problem and avoid catastrophic losses. These advances have been critical to the rapid intensification of food production that supports an ever-growing human population.

Nevertheless, despite extraordinary technological advances, pests and disease remain a significant challenge. To some extent, this results from larger and more intensive farms that create the ideal conditions for disease transmission and the emergence of new pathogens (Van Boeckel et al., 2015). Recent and ongoing experiences with highly pathogenic avian influenza, bovine spongiform encephalopathy, foot and mouth disease, and swine flu suggest that livestock are particularly vulnerable to outbreaks with serious economic consequences (Thornton, 2010). However, the feedback

mechanisms between human behavior and the spread of infectious diseases are both consequential and challenging to characterize. The private farmer faces a complex set of incentives to control the spread of disease. As Vergil recommends, the farmer should address an outbreak as soon as possible to avoid spreading to other animals; however, it is reasonable to expect this incentive to change at the farm boundaries. If the farmer's self-interest does not extend beyond their own property, then we should expect them to ignore the costs of disease spread to their neighbors. This externality can reduce the profitability of neighboring farms, suggesting a role for public policy.

A disease that moves across a landscape exhibits spatial-dynamic properties. In this case, the patterns of spread might be determined both by the physical characteristics of the landscape and by the decisions of the economic agents operating upon the landscape (Wilen, 2007). This adds complexity to the disease-induced externality. A pathogen may be transported by wind or water, where the spread is exogenous. But transmission will also be influenced by the private disease-management choices of farm owners, creating a tight linkage between disease control and the state of the system. These feedback loops are common to many environmental and resource management issues and relevant to a variety of applications, including invasive species (Epanchin-Niell and Wilen, 2015), capture fisheries (Sanchirico and Wilen, 1999), and wildlife disease (Horan and Wolf, 2005). As many have pointed out, designing and evaluating public policy in this setting can be quite challenging, and often requires deep structural knowledge of the system (Ferraro et al., 2019).

In this dissertation, I explore and evaluate the design of disease management policy in Chile's salmon aquaculture industry. Chile is the world's second-largest producer of farmed salmon and is the largest supplier of salmon to the United States. But Chile's salmon producers have struggled to manage outbreaks caused by endemic pathogens (SERNAPESCA, 2022). The main cause of this is an endemic pathogen, *Piscirickettsia salmonis*, the bacterium that causes Salmonid Rickettsial Syndrome (SRS). This pathogen infects the majority of production cycles in the country and is responsible for most of the disease-related deaths in Chile's farmed salmon. The production technology used by Chile's salmon industry leaves fish particularly vulnerable to pathogens. Although the fish are constrained to the pen, water and organic matter can move freely into and out of the system, making it impossible to exclude pathogens and parasites from entering the pen. Endemic

parasites and pathogens, including *P. salmonis*, can infect a farm and rapidly multiply (Rozas and Enríquez, 2014). It is also widely believed that SRS is transmissible between farms via tides and ocean currents, leading to the development of a costly production externality (Bravo et al., 2020). The principle method of control for SRS is antibiotics, and Chilean producers consequently utilize antibiotics at a rate that far exceeds most other salmon-producing countries (Avendaño-Herrera et al., 2023).

Given the ongoing challenges from endemic pathogens, what avenues are available for balancing disease management with antibiotic stewardship? One way to address this spatial-dynamic externality is through coordinated action. In the first chapter, I evaluate a spatially explicit disease management program that forces adjacent farm sites to coordinate their production activities. These coordinated groups are called *barrios*, or neighborhoods, and are required to observe a common minimum three-month fallow period and (approximately) coordinated production cycles. Theory suggests that this system will be most effective when transmission between neighborhoods is low because it reduces the likelihood of immediate re-infection at the start of the coordinated production cycle. However, the spatial structure of this program was not directly informed by connectivity between groups of farms. Therefore, I evaluate the extent to which pathogens (in our case, SRS) spread between neighborhoods, exploiting exogenous variation in the timing of production cycles induced by the disease management policy. I do not find evidence of spillover between neighborhoods, suggesting that the spatial scale of the policy is appropriately matched to the spatial scale of transmission. To account for the possibility of a behavioral response that could mitigate observed disease prevalence, I also test for changes in the propensity and intensity of antibiotic use that might be caused by the variation in pathogen pressure. However, I find no evidence of such effects.

In the second chapter, I turn toward the private incentives for disease control, specifically antibiotic use. The incentive to apply antibiotics is determined, in large part, by its effectiveness at reducing losses from disease. To simulate counterfactual disease control policies, it is therefore necessary to accurately estimate the effectiveness of antibiotic treatments in a farm setting. However, several econometric challenges are associated with measuring the effectiveness of treatment using observational data. Employing an epidemiological model with endogenous antibiotic

applications, I illustrate how self-selection into treatment adds bias to an empirical estimate of a structural parameter critical for counterfactual simulation of disease dynamics. This bias is driven by the behavioral feedback between treatment and stochastic features in the epidemiological model that are unobserved by the econometrician. When producers self-select into treatment based on these unobserved features, conventional models tend to underestimate treatment effectiveness. I illustrate an alternative estimator based on the control function (Wooldridge, 2015) and conduct a Monte Carlo experiment to illustrate its efficacy. I then apply the estimator to data from Chilean salmon farming industry to illustrate how it might be used in practice.

The private incentives for disease control are particularly relevant for ongoing policy debates over the consumption of antibiotics in the salmon industry. In response to growing sustainability concerns, many firms in Chile's salmon industry have made public commitments to reduce their use of antibiotics. In the final chapter, I explore the cost-effectiveness of alternative disease control strategies, focusing specifically on the tradeoffs between disease management and antibiotic stewardship. To inform this research, I partnered with Monterey Bay Aquarium Seafood Watch, an environmental non-governmental organization that is heavily involved in the efforts to abate antibiotic use. I solve for the dynamically optimal antibiotic treatment schedule in the presence of disease transmission between individuals on a farm and between farms. To capture the effect of the neighborhood coordination policy, I parameterize the disease pressure from other farms in a time-varying manner consistent with coordinated production cycles. I then simulate the impacts of various antibiotic reduction policies, including a cap on the volume of antibiotics applied, a cap on the number of antibiotic applications, and a limit on the maximum weight at which fish can be treated. Each of the restrictions causes a non-uniform shift in the timing of treatment in the production cycle, an effect that is ultimately consequential for disease-related externality. I find that the most cost-effective instruments for private individuals tend to generate the largest externalities.

My dissertation chapters are supported by confidential data from the Servicio Nacional de Pesca y Acuicultura (SERNAPESCA), the national regulatory authority for Chile's aquaculture industry. These data are collected as part of a mandatory reporting requirement and provide georeferenced, weekly observations on the production status and veterinary health of all of the industry's fish farms (a full description of the data are included in the proceeding chapter). Critically, these data also

include prescriptions for antibiotics applied for a wide range of purposes, including the treatment of SRS. Although these data have been used in past epidemiological studies, this dissertation is the first to utilize these data for economic analysis. In some cases, I pair this database with supplementary data, drawing upon publicly available information on policy borders and coastlines that are central to the empirical evaluation of the neighborhood policy in the first chapter.

CHAPTER 1

Policy Analysis in a Spatial-Dynamic Environment: Chile's Neighborhood Policy

1.1. Introduction

Spatial-dynamic processes are at the heart of many challenging policy problems. Human disease, agricultural pests, invasive species, and forest fires are all defined by movement or diffusion across a landscape. However, the spatial scale that is relevant for the natural system does not often reflect the spatial scale of the property rights, political borders, and other institutions that are important for policy-makers. As a result, ideal control of a spatial-dynamic problem may require coordination across decision-makers whose self-interest does not inherently extend to the entire landscape (Wilen, 2007). While efficient levels of coordination can emerge spontaneously under decentralized management (Coase, 1960; Ostrom, 2015), policy-makers often choose to induce coordination to prevent the emergence of costly externalities (Oates, 1972).

Coordination policies, however, are not always beneficial. Landscapes are often heterogeneous across both economic and ecological dimensions, such that the nature of the externality depends upon an individual's location within the landscape. In a dynamic setting, the externality may change across space and over time. This point has been demonstrated theoretically in a variety of settings including commercial fisheries (Smith et al., 2009), infectious disease transmission (Horan and Wolf, 2005; Rowthorn et al., 2009), agricultural pests (Atallah et al., 2017; Fuller et al., 2017), and invasive species control (Albers et al., 2010; Epanchin-Niell and Wilen, 2015). Policy-induced coordination that fails to appropriately incorporate spatial heterogeneity may generate unexpected or even harmful outcomes, resulting from the underlying dynamics of the bio-physical system (Ives and Settle, 1997; Sanchirico and Wilen, 2005). If policy compliance is costly and the externality is unmitigated, then industry participants may be worse off than before the policy is implemented.

In general, the literature illustrates that the optimal approach can vary widely and will depend upon the idiosyncratic nature of each system. Thus, applying these insights in practice is challenging because it requires extensive knowledge of both the physical characteristics of the landscape and the economic agents operating within it. It is ultimately an empirical problem to evaluate whether a policy is appropriately designed for a particular context.

In this paper, we evaluate the effectiveness of a spatial policy for managing transmissible disease in aquaculture. In our case study, Chile's salmon aquaculture industry, endemic pathogens are transmissible through the water, allowing diseases to spread readily between farm sites. The resulting spatial-dynamic externalities are significant, as pathogens induce millions of dollars in damage annually to Chilean fish farmers (Dresdner et al., 2019). To address this problem, the spatial management policy requires groups of neighboring farms to observe coordinated fallowing (i.e., "rest") periods and production cycles. Given that the key pathogens are thought to be transmissible over long distances, the spatial scale of the coordinated neighborhoods should, ideally, match the scale of pathogen dispersal. If the neighborhoods are too small, then pathogens will spread between neighborhoods and undermine the benefits of the coordinated fallowing policy. If they are too large, then the policy may unnecessarily increase the production costs for the industry. Using data on veterinary outcomes and antimicrobial-use decisions for the entire Chilean salmon industry, we empirically test for spillovers between coordinated groups of farms and evaluate whether Chile's spatial management policy is appropriately scaled to address the spatial externality.

Empirical analyses of marine disease have primarily come from the epidemiology literature and have convincingly illustrated that pathogen levels are correlated across time and space, suggesting water-born transmission over long distances (Bravo et al., 2020; Kristoffersen et al., 2013; Rees et al., 2014). The focus of such literature is on predictive rather than causal inference. To evaluate the extent to which Chile's spatial policy is appropriately scaled to dispersal, we need to measure whether pathogen prevalence in one neighborhood causes pathogen levels to change in another. This requires us to carefully consider the mechanisms driving disease transmission. We argue that observed pathogen levels are simultaneously determined across farms linked by dispersal. This point is not emphasized in the epidemiology literature associated with marine disease because of the focus on predictive inference, but it is a fundamental implication of the theoretical bioeconomic literature

on spatial-dynamic systems. Feedback between observational units across space and through time leads to simultaneity-induced endogeneity that commonly undermines causal identification (Ferraro et al., 2019).

To avoid simultaneity bias, we exploit recurring, quasi-experimental variation in pathogen prevalence induced by coordinated fallowing to measure the extent to which pathogens in one neighborhood spill over onto farms in adjacent neighborhoods. Using a stylized model of pathogen dispersal, we motivate two complementary empirical strategies. Our main specification, a difference-in-differences model with dynamic treatment effects (Sun and Abraham, 2021), imposes few assumptions over the structure of the underlying connectivity patterns. However, the variation that underpins the difference-in-differences design is most credible at times when spillover effects are likely the smallest. To address this, we complement the analysis with a difference generalized method of moments model (Arellano and Bond, 1991), which allows us to utilize more information (including the periods in which spillovers are most likely), although it requires more restrictive assumptions regarding the spatial-temporal patterns of pathogen dispersal.

We find little evidence of spillovers between coordinated neighborhoods. This finding is robust to a variety of structural assumptions over the nature of connectivity between farms. We also do not find evidence that the rate or intensity of antimicrobial treatments change in response to neighboring pathogen pressure. These results suggest that the spatial scale of dispersal does not, on average, exceed the spatial scale of the neighborhood policy and that the spatial scale of the coordinated neighborhoods is sufficiently large. This is surprising given that the borders of the policy are known to reflect a compromise between competing interests and are thought to be largely independent of underlying connectivity patterns. Nevertheless, our results suggest that this second-best policy is large enough to contain pathogen spread across space, limiting the extent of the externality.

This paper makes several contributions to the literature. First, it adds to the growing body of quasi-experimental evidence related to regulatory spillovers in environmental settings. Spatial-spillovers have been studied in a variety of empirical contexts including surface water pollution (Lipscomb and Mobarak, 2017), groundwater pumping (Pfeiffer and Lin, 2012; Rouhi Rad et al., 2021), conservation (Alix-Garcia et al., 2012), deforestation (Robalino et al., 2017), and marine

protected areas (Reimer and Haynie, 2018). However, policy-induced spillover in our setting directly informs the extent to which the policy matches the spatial scale of the externality—in our case, water-born pathogen dispersal. This is of practical importance to policymakers and offers a rare opportunity to empirically bring the extensive theoretical literature on landscape heterogeneity to an applied context.

This paper also studies a relatively novel policy design. Chile's neighborhood system is a top-down approach to controlling production externalities that relies on both spatially and temporally explicit restrictions. In this sense, the policy resembles time-area closures meant to protect non-target aquatic species from bycatch in commercial fisheries (Hazen et al., 2018) or migratory species with conservation value (Ando et al., 1998; Reynolds et al., 2017). As in our setting, the effectiveness of these policies depends critically upon the functional connectivity of the landscape and the movement of organisms over space. However, there are few empirical applications in which the effectiveness of a spatial-dynamic policy is actually evaluated.

Third, our paper informs our understanding of the economics of pest and disease control in food systems. Chile's neighborhood system bears close similarities to area-wide pest management programs in terrestrial agricultural settings. However, area-wide pest management policies are challenging to implement. Voluntary programs can be undermined by free-riding and strategic uncertainty (Singerman and Useche, 2019), while mandatory programs are often politically intractable (Olmstead and Rhode, 2004). Chile's spatial management policy offers a unique opportunity to study a program that is both federally mandated[1] and large in economic scope. The program simultaneously coordinates the operation of hundreds of farm sites in southern Chile, representing billions of dollars worth of fish. Other salmon-producing countries, including Norway and Scotland, have also deployed large-scale spatial management programs to solve similar problems (Murray and Peeler, 2005). However, few have empirically evaluated the extent to which the spatial scale was appropriate for biological dispersal or effective at reducing any pathogen. This paper is the first to evaluate the role of public policy in controlling Salmonid Rickettsial Syndrome (SRS), an endemic

[1]This became politically feasible in the wake of a disease-induced collapse and restructuring of the industry in 2007-2009 (Iizuka, 2016)

marine disease in Chile's coastal waters. SRS is responsible for nearly all of the industry's antibiotic use, a subject of growing concern among veterinarians and public health experts (Smith and Mardones, 2020).

The remainder of this paper proceeds in the following manner. Section 1.2 describes the economics of endemic diseases. Section 1.3 describes the Chilean salmon aquaculture industry and the relevant spatial management policies. Section 1.5 discusses are empirical strategy and describes our data. Sections 1.6 and 1.7 present the model estimates of our two complementary models. Sections 1.8 discuss the results and 1.9 concludes.

1.2. Economics of Endemic Disease

Communicable diseases are a long-standing challenge for livestock industries around the world. While individual producers have a basic incentive to maintain healthy and productive stock, the level of investment depends upon the costs of surveillance, prevention, and control of pathogens relative to their benefits. However, the key challenges related to disease management are intertwined with the production technology and disease ecology that characterize the system. In this section, we discuss the incentives of private producers facing an endemic disease in an open production environment and consider possible policy options to better understand the Chilean experience.

Diseases vary widely in their virulence, transmissibility, and distribution. From the perspective of a farm manager, exotic diseases are those that are not currently present in the production environment but could be introduced. By contrast, endemic diseases are driven by pathogens already present in the production environment. In comparison to exotic diseases, endemic diseases are much more likely to require consistent management efforts because of the constant risk of infection (Hennessy, 2007). Intensive production of plants and animals increases the density of possible disease hosts, providing the right conditions for a low-prevalence, endemic disease to rapidly infect a large proportion of the farm. Even if the endemic disease is non-lethal, this can impose significant costs on the producer.

If diseases are not transmissible between farms, individual producers can control their exposure to pathogens and capture the full benefits associated with investments in disease management. Even if the production technology does not allow for perfect protection against diseases, they do

not create a production externality that might affect other producers. In open production systems, however, transmission between producers is possible. When farms are epidemiologically linked, the disease management choices of one producer will affect the disease status (and production costs) of others in the industry. This condition makes the shared production environment a common pool resource and disease transmission a (harmful) production externality.

In the simple model of an open production system illustrated in Hennessy and Marsh (2021), equilibrium production in the commons exceeds the surplus maximizing level, and disease-related losses reduce industry-wide productivity. Under certain conditions, cooperative solutions to production externalities can emerge without public policy. Previous theoretical work has shown the potential benefits of private solutions in transmissible agriculture disease (Atallah et al., 2017; Fuller et al., 2017) as well as in other spatial-dynamic externalities such as invasive species (Epanchin-Niell and Wilen, 2015) and nuisance wildlife populations (Bhat and Huffaker, 2007). In some cases, cooperative solutions can also emerge via horizontal industry consolidation as the larger firms may be incentivized to internalize the full costs of disease transmission (Fischer et al., 2017).

However, cooperative solutions often fail to emerge, suggesting a role for public policy. Ideally, a policy would target the specific transmission pathways through which the externality is affecting other producers. In the case of agricultural disease, this might be the amount of the pathogen an infected farm transmits to its neighbors. But the challenges of measuring the pathogen prevalence and characterizing the connectivity between neighboring farms tend to drive more practical, 'second-best' alternatives. These include policies designed to reduce the extent to which the shared production environment is used for the externality-producing activity (e.g., resource taxes or entry-restrictions) as well as spatiotemporal regulations designed to enhance the effectiveness of private externality abatement efforts (e.g., coordinated fallowing periods, coordinated disease treatments, all-in all-out production) (Hennessy and Marsh, 2021). Simulation studies suggest that coordinated disease management, in particular, can reduce the abundance of pathogens, even with complex connectivity patterns and transmission between clusters of farm sites (Murray and Salama, 2016).

Policies that regulate the externality-generating activity instead of the externality itself face additional challenges. For example, producers typically respond to restrictive policies across many margins. As Hennessy and Marsh (2021) show, entry restrictions without consideration of other

margins for adjustment will tend to drive incumbent users to intensify their use of the resource, undermining the intent of the policy. This general outcome is echoed by Oglend and Soini (2020), who show that permit limitations in the Norwegian salmon aquaculture industry may not be effective at controlling parasitic sea lice. Furthermore, the effectiveness of spatially explicit policies depends upon the connectivity patterns in the environment. For example, Werkman et al. (2011) show that the effectiveness of coordinated disease management efforts depends upon the connectivity between groups of coordinated sites. The benefits of coordinated fallowing periods, for example, are far larger under lower infection pressure from uncoordinated sites. This is important because enforced coordination reduces a firm's marketing and production flexibility, possibly inducing significant private costs. Therefore, to maximize net benefits, policy-makers should consider the underlying geophysical features driving connectivity when forcing coordination.

1.3. The Chilean Salmon Industry

Chile is the world's second-largest salmon producer, behind Norway. In 2021, the industry had nearly one million tons of biomass under production, directly employed more than 20,000 workers in southern Chile, and exported 5.1 billion USD worth of salmon products (SERNAPESCA, 2022). Currently, the industry produces three salmonid species: Atlantic Salmon, Coho Salmon, and Rainbow Trout.

Although Chilean salmon competes in a global seafood market with wild-caught fish, most modern salmon farms bear a closer resemblance to terrestrial livestock operations. Eggs from captive broodstock are fertilized and raised in freshwater land-based hatcheries. When the juveniles are large enough to survive in saltwater (50-300 grams), they are transferred into saltwater net pens to grow to market size (4-6 kilograms). In this final phase, farmed salmon are contained in the net pen until harvest, but water and organic matter can move freely between the pen and the outside environment. The farmer can feed the fish and conduct maintenance but cannot control the environmental conditions or the exposure to pathogens (Olson and Criddle, 2008).

The net pen or "marine fattening" phase is the focus of this paper. Fish remain in the net pen for 10-20 months, depending upon the species of the fish. The net pen remains in a single location during the entire production cycle. Since 1979, the Chilean government has leased sections of coastal

water (concessions) to private individuals and corporations for salmon farming. The concessions are fixed by GPS coordinates and cannot be moved. Once granted, the concession gives a firm the option to operate a salmon farm with the concession boundaries, to transfer or lease the operating right to another firm, or not to use the concession at all. New leases were offered on a perpetual basis until 2010, at which point the terms were reduced to 25 years with the possibility of renewal (Iizuka, 2016).

The vast majority of concessions have been granted in Región de Los Lagos (Region X) and Región Aysén del General Carlos Ibáñez del Campoin (Region XI), the southern part of the country, where farm sites enjoy close proximity to processing facilities and transportation infrastructure. In recent years, however, producers have begun to obtain concessions further south in Región de Magallanes y de la Antártica Chilena (Region XII).

1.3.1. Salmonid Rickettsial Syndrome (SRS). The most significant production risk during the fattening phase is infectious disease. Marine net pens are an open production system. Producers cannot regulate the environmental conditions on the farm, nor can they prevent exposure to pathogens in the water where the pen is located. Although many diseases and parasites threaten Chile's salmon industry, this analysis focuses on one of the most economically important: Salmonid Rickettsial Syndrome (SRS) or *Piscirickettsiosis*.

SRS has been a major challenge for the Chilean salmon industry since the disease was first described in 1989 (Bravo and Campos, 1989). SRS is caused by *Piscirickettsia salmonis*, a bacterial pathogen, and is currently the leading cause of disease-related mortality on Chilean salmon farms (SERNAPESCA, 2022). The disease has become endemic in Chile and is present in nearly 60% of all production cycles. However, SRS outbreaks are concentrated in Regions X and XI, the most northern and heavily utilized area of Chile's salmon-producing regions. SRS has also been observed on salmon farms in North America and Europe, but the disease is not believed to be endemic as observed outbreaks have been mild and isolated (Smith and Mardones, 2020).

Vaccines for SRS are commercially available and are widely used, but none provide long-term protection against infection (Maisey et al., 2017). Once infection arrives, producers use antimicrobial products administered via medicated feed pellets or injection (Smith and Mardones, 2020). Prescriptions targeted at SRS account for more than 90% of the industry's total antibiotic use, but

they often fail to eliminate infections (Price et al., 2016). It is important to note that producers cannot use antibiotics on the farm before disease arrival (i.e., for prophylaxis). However, once the disease has arrived on the farm, producers are permitted to treat as many cages as necessary.

While farmers have a private incentive to control outbreaks, their efforts are rarely sufficient to prevent transmission. Infected fish shed bacteria into the water column, where they can be transported via surface currents. Lannan and Fryer (1994) show that *P. salmonis* can survive for two weeks in seawater, depending upon salinity and temperature. Observational and simulation-based studies seem to corroborate this, highlighting spatiotemporal correlation in SRS mortalities between farms up to 10 km apart (Rees et al., 2014) and the role of regional hydrodynamic connectivity patterns in SRS risk (Bravo et al., 2020).

FIGURE 1.1. Salmon Farming Concessions and Neighborhoods in Regions X and XI

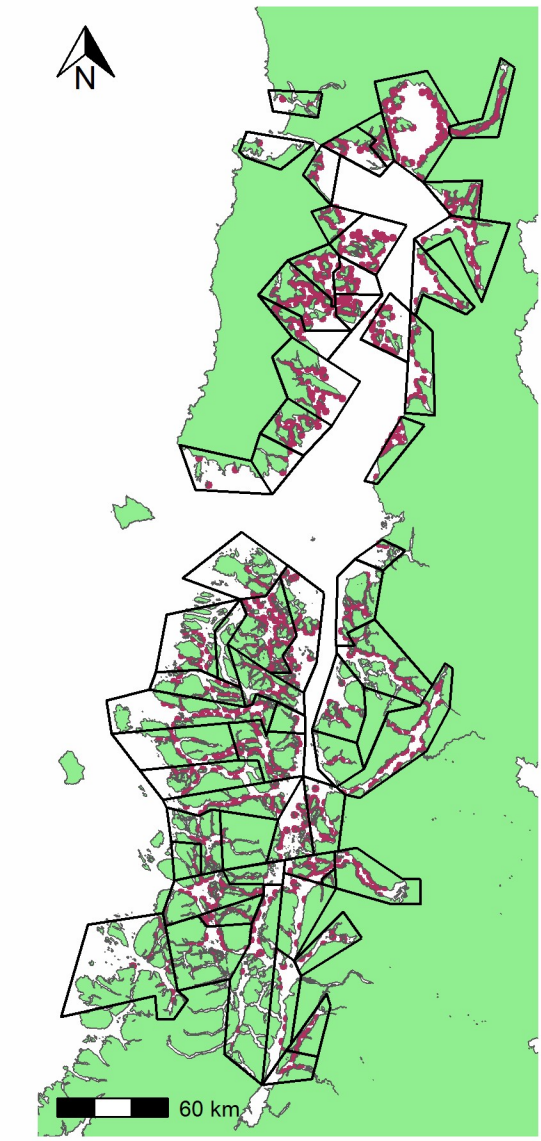

Note: The red dots are all of the concessions granted by SERNAPESCA for salmon-farming activities since 1979. These reflect all of the possible locations where salmon farming could occur in Regions X and XI. The black polygons illustrate the spatial extent of the neighborhoods.

1.3.2. Spatial Management. Before 2010, salmon farming companies in Chile were allowed to utilize their concessions with relatively few restrictions. Farm sites were required to be one mile apart from any other farm site, and concessions were only granted in areas deemed well-suited

for aquaculture activities to avoid conflict with other maritime industries. However, a major disease outbreak between 2007-2009 led to a regulatory overhaul (Iizuka, 2016). To address ongoing challenges with both endemic and exotic pathogens, the Chilean government passed the modified Aquaculture Law (Ley No. 20.434) in 2010. Among other provisions, the policy increased the minimum distance required between individual sites to 1.5 miles and established a network of production *barrios* (neighborhoods). These are depicted on the map in Figure 1.1. The neighborhoods serve as a coordinating mechanism between adjacent farm sites. Each neighborhood is assigned a schedule of mandatory, three-month fallowing (rest) periods in which all fish, nets, and other farm equipment must be removed from the water. By stipulating when fish cannot be in the water, the fallowing policy also determines when fish can be in the water, effectively coordinating the production cycles of all farms in the neighborhood.

The months between the fallowing periods in which farms are allowed to operate (hereafter "productive periods") are either 21 or 24 months in length in Regions X and XI. However, farms rarely utilize the entire productive period, often extending their fallowing period beyond the required three months. This occurs because the allowable productive period exceeds the time required to grow salmonids to market size but is typically too short to complete multiple production cycles.[2] As a result, neighborhoods collectively also tend to exhibit "effective" coordinated fallowing periods exceeding three months and coordinated productive periods that are shorter than the policy requires.

[2]In rare cases, companies can fit two production cycles of Rainbow Trout or Coho salmon into a single production period.

FIGURE 1.2. Cumulative Concessions for Salmon Farming in Regions X and XI

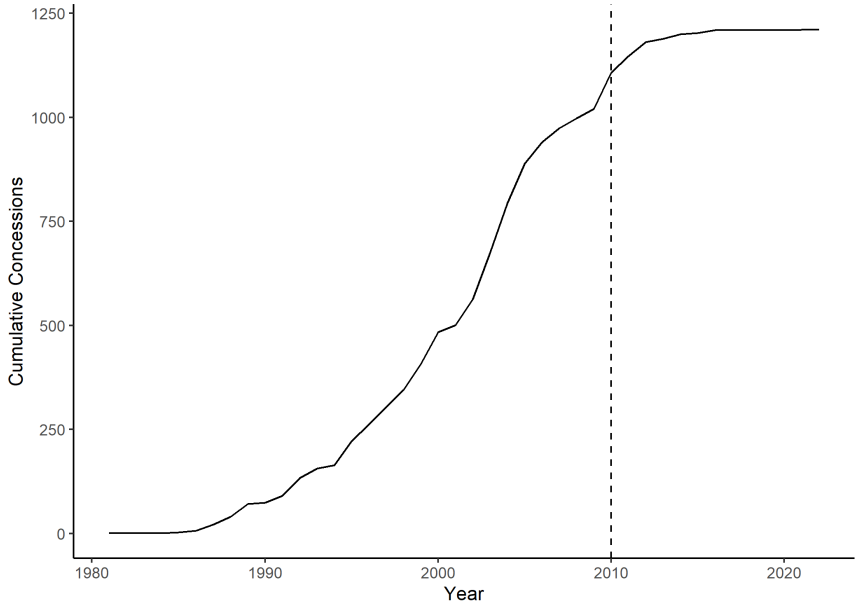

Note: This figure plots the cumulative number of concessions granted by Chile's Undersecretariat of Fisheries and Aquaculture (SUBPESCA) between 1980 and 2022. The vertical dashed line denotes the passage of the modified Aquaculture Law in 2010. Between 2011 and 2022, 64 new concessions were approved.

In the years since 2010, few new concessions have been granted in Regions X and XI. The modified Aquaculture Law placed short-term limits on new concessions following 2010. Therefore, the network of available concessions in Regions X and XI has not changed significantly since the passage of the policy. This restriction is clearly depicted in Figure 1.2. [3] At the same time, Chilean policy also caps the total biomass allowed on each concession. This limits the extent to which companies can respond to entry restrictions with growth along the intensive margin.

The location of the neighborhood boundaries is crucial to the success of the program. If pathogens can spread between neighborhoods, then farms are likely to be rapidly re-infected after fallowing, undermining the benefits of coordination (Werkman et al., 2011). *P. salmonis* is thought to be transmissible via surface currents, but transmission is likely also dependent upon physical conditions (e.g., temperature and salinity) as well as by the ambient pathogen levels in the area. An ideal system might include all of these characteristics in defining the spatial extent of the

[3]The approval process for new concessions may take years. It is likely that many of the concessions granted after 2010 reflect applications begun before the passage of the law.

neighborhoods and the schedule for fallowing. The text of the modified Aquaculture Law stipulates that groups of adjacent concessions with similar epidemiological, oceanographic, operational, or geographic characteristics should be coordinated (Ley No. 20.434, Article 1, Section 52). In practice, this allowed for the creation of neighborhoods that reflect a compromise between the operational needs of companies, the (non-disease) ecological impacts of salmon farming, and the epidemiological conditions relevant to transmissible pathogens. Therefore, it is unlikely that the existing neighborhoods reflect underlying connectivity patterns relevant specifically to SRS so as to enclose distinct, isolated groups of farms (Alvial, 2017; Tecklin, 2016).[4]

Spatial coordination policies are not unique to Chile. Norway (Guarracino et al., 2018) and Scotland (Murray and Salama, 2016) utilize a similar program designed to manage disease and parasites in their salmon industries, but with relatively less stringent levels of coordination between farms. More recently, Norway has also implemented a spatially defined output control policy, the "traffic light system," to control populations of parasitic sea lice. However, this policy focuses on mitigating parasite-induced mortality of wild fish rather than improving the productivity of the fish farming industry (Taranger et al., 2015). Many of the world's largest fish-farming countries, including China, Brazil, Indonesia, Mexico, the Philippines, and Turkey, have all implemented some kind of zoning program to manage their aquaculture industry. But, in most cases, the zoning process was not explicitly designed to reduce disease spread but to manage conflict between other users or to limit the industry's ecological impacts (Aguilar-Manjarrez et al., 2017). Furthermore, this system bears a close resemblance to area-wide pest management strategies in terrestrial agriculture. However, the spatial scale and economic scope of this policy are far larger than most area-wide management programs, reflecting the high levels of connectivity that exist in the marine environment.

1.4. Model of Pathogen Prevalence and Fallowing

In the last section, we argued that neighborhood boundaries are critical for the success of the program. The boundaries determine the extent to which farms in one neighborhood are exposed to farms in adjacent neighborhoods. When connectivity is high, pathogens are likely to spill over

[4]Informal interviews with industry participants revealed a common belief that the neighborhood borders were constructed *ad hoc*

between neighborhoods, reducing the benefits of coordination. In this section, we formalize this idea with a simple, two-patch model of pathogen dispersal to illustrate the key mechanisms in the system and to motivate an empirical strategy for measuring spillovers between neighborhoods.

Consider a landscape with two patches, where salmon farming can occur and endemic pathogens are present. We will assume that the two patches are connected to one another via dispersal. At any point in time, pathogen levels, x, can be described using the following system of equations

$$x_{1,t+1} = \rho_1 f_{1t} x_{1t} + \rho_1^f (1 - f_{1t}) x_{1t} + \gamma_{21} x_{2t} + \delta_{21} x_{2,t+1} \tag{1.1}$$

$$x_{2,t+1} = \rho_2 f_{2t} x_{2t} + \rho_2^f (1 - f_{2t}) x_{2t} + \gamma_{12} x_{1t} + \delta_{12} x_{1,t+1} \tag{1.2}$$

where $x \geq 0$.[5] When farming occurs on a patch, $f = 1$ and a large influx of susceptible hosts enters the system. These fish become infected and pathogen levels on the patch grow at a rate determined by $\rho > 1$. At the end of the production cycle, the patch enters a fallow period and $f = 0$. During this period, all of the fish are removed, reducing the number of available hosts. Pathogen levels on the patch shrink at a rate determined by $\rho^f < 0$. We expect producers to incur costs fighting disease outbreaks and reducing pathogen loads. In doing so, farmers moderate the growth of pathogen populations by applying antimicrobial products, essentially working against local reproduction. However, for the sake of simplicity, we simply acknowledge that the patch-specific reproduction rates, ρ, reflect both biological and human mechanisms.[6]

The third and fourth terms determine the extent to which disease is transmissable between patches. If $\gamma = \delta = 0$, then the disease is not transmitted between patches and there is no spillover. In this case, pathogen levels are completely determined by the rate at which the pathogen reproduces on the patch. This can be expressed as either a recursive function or as an explicit function of time elapsed since the start or end of fallowing. For example, suppose $f_1 = 0$ from from $t = 0$ to $t = \bar{t}$ and $f_1 = 1$ from from $t = \bar{t}$ to $t = T$. It follows immediately that, when $\gamma = \delta = 0$,

$$x_{1,T} = x_{1,T-1}\rho_1 = x_{1,\bar{t}}(\rho_1)^{T-\bar{t}} \tag{1.3}$$

[5]Given our interest in endemic pathogens, we will assume some ambient pathogen prevalence exists, such that $x_0 > 0$ in both patches.

[6]To isolate spillover effects between patches, it is necessary to address treatment. We return to this in the empirical section of the paper.

where $x_{1,\bar{t}}$ is the pathogen load at the end of fallowing. In period T, $T - \bar{t}$ periods have elapsed since the end of the last fallowing period. Although it would be cumbersome to express the system of equations in terms of their initial conditions and time since the beginning or end of fallowing, it is important to emphasize that the time elapsed between the policy-induced production milestones enters explicitly into our model of pathogen prevalence. This ultimately will become important in our empirical strategy.

When $\gamma, \delta > 0$, the disease can be transmitted between patches. We assume that pathogen dispersal occurs continuously between x_1 and x_2, so both lagged and contemporaneous pathogen levels on each patch enter into the state dynamics equation of the other. Although producers can apply treatment in response to disease arrival on the farm, they can do little to avoid dispersal. Therefore, farms benefit from lower values of γ and δ. This has implications for the design of the neighborhood system. All else equal, farms in a neighborhood would benefit from lower exposure to dispersal from farms in another neighborhood.

Equations 1.1 and 1.2 point to an important empirical challenge that is common in spatial-dynamic problems. Simultaneity obscures the causal relationship between patches, making it challenging to recover unbiased estimates of the model parameters (Ferraro et al., 2019). When both patches are active ($f_1 = f_2 = 1$), x_1 and x_2 are likely to be correlated even if $\gamma = \delta = 0$ and the two patches are independent from one another. This is particularly relevant in our setting because the production cycles of adjacent farm sites are coordinated, implying that f_1 and f_2 follow similar patterns. To solve this problem, we need variation in pathogen prevalence that is plausibly independent from neighboring patches. The fallowing policy, f, generates this variation and we will use it to identify dispersal between patches.

1.4.1. Fallowing and Pathogen Dispersal. In the previous sections, we described how current fallowing policy enforces coordinated fallowing and production cycles between adjacent farm sites. In terms of our model, this corresponds to a case in which f_{1t} and f_{2t} are synchronized, such that f_1 tends to equal f_2. In this section, however, we will focus on our attention on an alternative case, one in which the patches are not syncronized.

When the patches are not syncronized, the fallowing policy generates variation in pathogen pressure that can be used as part of an identification strategy to measure spillovers. To illustrate

this point, we will simulate Equations 1.1 and 1.2 under two different fallowing scenarios. In the first scenario, patch 1 operates during the entire simulation interval, $t \in [0, T]$. Meanwhile, patch 2 fallows during the interval $[0, \bar{t}]$ and then becomes active in the interval $(\bar{t}, T]$, where $\bar{t} < T$. Denote the vector of f_{2t} values that describes fallowing in Patch 2 during first scenario as f_2^1. The first scenario is summarized in the following system of equations:

$$x_{1,t+1} = \rho_1 x_{1t} + \gamma_{21} x_{2,t} + \delta_{21} x_{2,t+1} \qquad 0 \leq t \leq T \tag{1.4}$$

and

$$x_{2,t+1} = \begin{cases} \rho_2^f x_{2t} + \gamma_{12} x_{1,t} + \delta_{12} x_{1,t+1} & 0 \leq t \leq \bar{t},\ f_1 = 1, f_2 = 0 \\ \rho_2 x_{2t} + \gamma_{12} x_{1,t} + \delta_{12} x_{1,t+1} & \bar{t} < t \leq T,\ f_1 = 1, f_2 = 1 \end{cases} \tag{1.5}$$

In the second scenario, Patch 1 is active during the entire simulation ($f_1 = 1$ for all t) while Patch 2 remains fallow for the entire simulation ($f_2 = 0$ for all t). Denote the vector of f_{2t} values that describes fallowing in Patch 2 during the first scenario as f_2^2. In both scenarios, we allow simultaneous dispersal between the two patches by setting $\gamma = \delta > 0$. Any switch between fallowing and production in either patch will change the pattern of pathogen prevalence and dispersal. When Patch 2 switches from fallow to active at $t = \bar{t}$, pathogen levels in that patch will begin to increase. This not only increases the pathogen load in Patch 2 but it also increases dispersal to Patch 1. The pathogen loads in Patch 1 will now be higher than they would be had Patch 2 remained fallow.

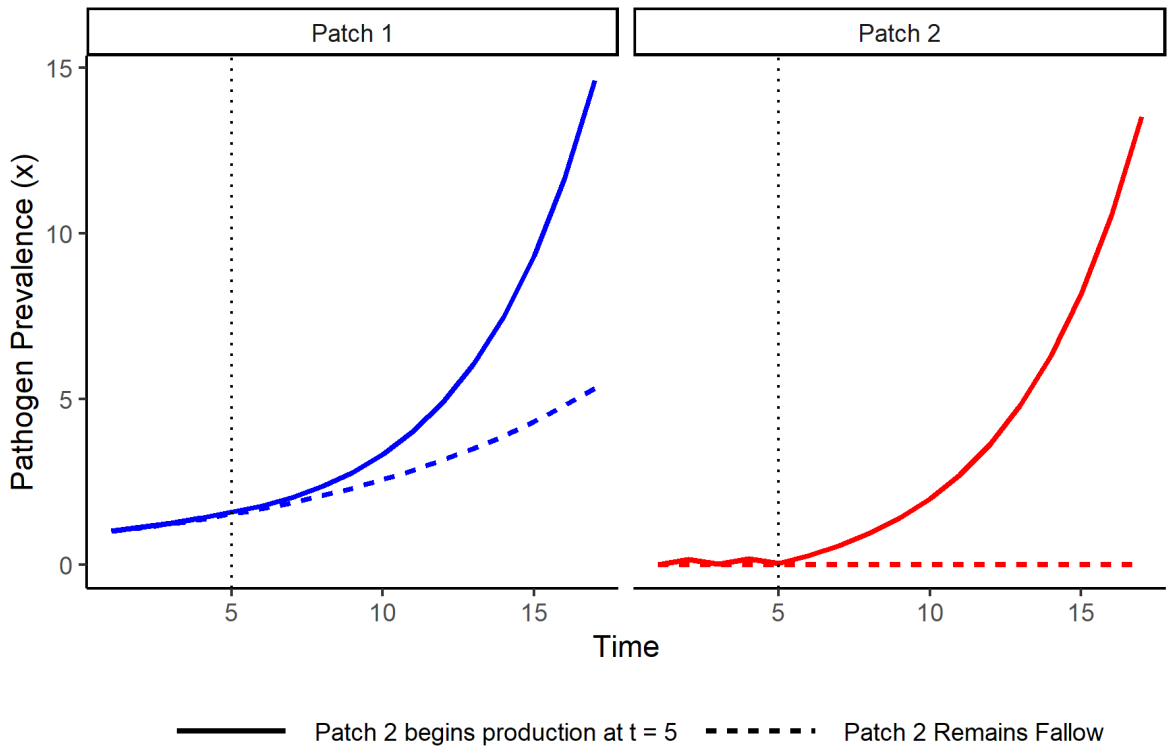

FIGURE 1.3. Pathogen Loads Under 2 Dispersal Scenarios

Note: This figure illustrates the simulated pathogen prevalence in the two patch system under two separate scenarios. The vertical line at $t = 5$ represents the end of the fallow period in patch 2 under the first scenario. Pathogen prevalence is scaled such that the initial pathogen prevalence level in Patch 1 is equal to one.

The simulation results are presented in Figure 1.3. Pathogen levels in Patch 1 are illustrated on the left and pathogen levels in Patch 2 are illustrated on the left. The solid line represents pathogen prevalence in the scenario in which Patch 2 switches from fallow to active while the dashed line represents the scenario in which Patch 2 remains fallow for the duration of the simulation. We have parameterized the model such that $\bar{t} = 5$ and the starting pathogen levels in Patch 2 are near zero.[7]

[7]This is consistent with the existing literature, which suggests that the required fallowing periods are likely effective at reducing the prevalence of *P. salmonis* to negligible levels (Price et al., 2017).

The difference between the solid line and the dashed line in Patch 1 represents the effect of dispersal from Patch 2. We'll refer to this wedge as α, where

$$\alpha_t = (x_{1t}|f_2^1) - (x_{1t}|f_2^2) \tag{1.6}$$

As above, f_2^1 is the entire fallowing schedule in scenario 1, where Patch 2 switches from fallow to active at $\bar{t}$. f_2^2 is the entire fallowing schedule in scenario 2, where Patch 2 remains fallow for all periods in the simulation. As a result of this fallowing schedule, α is a function of time since the end of fallowing in Patch 2. As the figure illustrates, we expect α to grow larger, the longer that Patch 2 remains in operation. This mirrors the intuition from Equation 1.3, but with additional richness. Not only does the time elapsed since the end of fallowing on Patch 2 determine pathogen prevalence on Patch 2, but dispersal ensures that it also enters into the expression for pathogen prevalence on Patch 1.

The effect of dispersal from Patch 2 on Patch 1 is the parameter of interest and the focus of our empirical strategy. The dashed line is the counterfactual outcome as it captures the time path of pathogen prevalence on Patch 1, given that Patch 2 remained fallow. If there is no dispersal between the patches, then there would be no difference between the dashed line and dotted line, and α would collapse to zero. This outcome represents the null hypothesis in our statistical tests. If the difference between the solid and dashed lines is sufficiently large, then we will reject the null, providing evidence that dispersal effects are present. This statistical test is the core of our identification strategy.

1.5. Empirical Strategy and Data

To identify the spillover effects described in the theoretical section, we exploit quasi-random variation in the timing of coordinated fallowing (f in the theoretical model). This variation does not come from the establishment of the policy or from a one-time policy "rollout" across time or space. Instead, it is a recurring feature of the policy design that can be found throughout the study period. As described in the background section, the management policy assigns to each neighborhood a twelve-week period in which all farms within the neighborhood are required to remain fallow. However, production within a neighborhood rarely begins on the first allowable

week after the end of the mandated fallowing. As a result, the "effective" fallowing periods tend to be much longer than the twelve required weeks. The start of any effective fallowing period ($\bar{t}$ in the theoretical model) is exogenous because it reflects the aggregate response of many firms to the fallowing schedule. Each firm will choose when to begin their production cycle. These delays are completely independent of farms in adjacent neighborhoods. In other words, a farm manager cannot control the timing of their exposure to dispersal from adjacent neighborhoods.

FIGURE 1.4. Length of Observed Coordinated Fallow Period

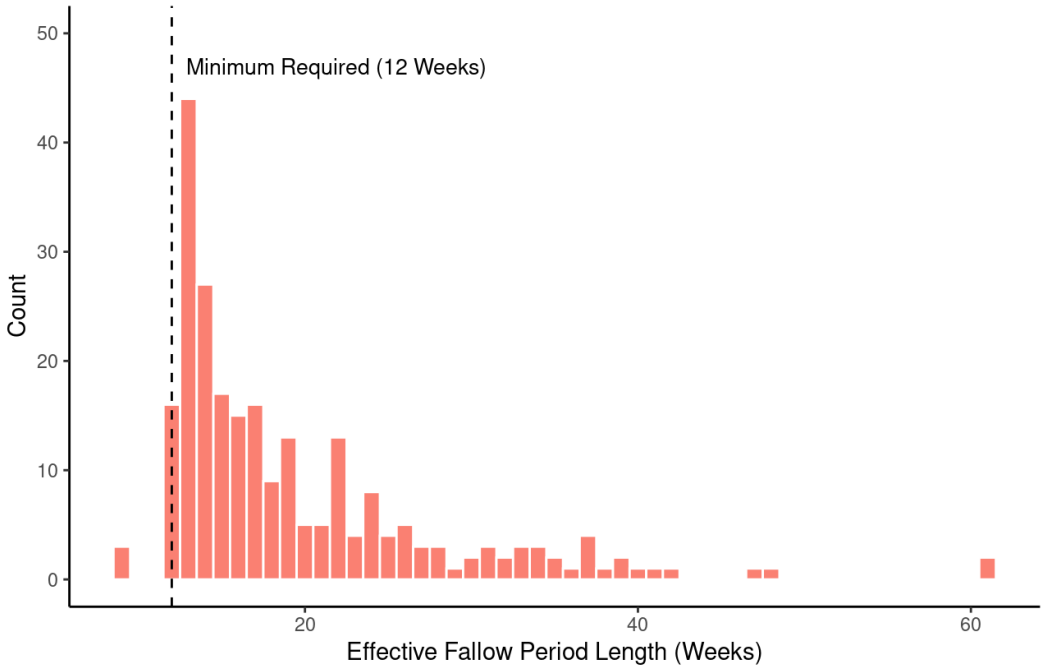

Note: This figure plots the distribution of observed coordinated fallow periods at the neighborhood level. The dashed line represents the minimum fallow length required by law. The observations to the left of the dashed line represent three instances in which the minimum required fallow length was temporarily reduced to two months early in the program's implementation.

The recurring fallowing periods allow us to address the theoretical concern that pathogen levels are simultaneously determined. When a neighborhood switches from a fallow period to a productive period, the large influx of biomass tends to lead to a rapid increase in pathogen prevalence. This is essentially the situation simulated in Figure 1.3. We argue that this increase in pathogen prevalence

(and dispersal) after fallowing amounts to an exogenous shock to farms in adjacent neighborhoods that we can use to measure spillovers from one neighborhood onto another.

We use two complementary econometric approaches. Our main specification is a difference-in-differences (DiD) model with dynamic treatment effects. It measures the effect of a switch from fallow to active in those cycles which are exposed to such a switch, to a control group that never is exposed to the switch. In doing so, it essentially measures α, defined in Equation 1.6. We also estimate an alternative model using generalized method of moments (GMM). Although the GMM model can utilize more data than the DiD model, it requires additional assumptions regarding the spatial-temporal patterns of dispersal. Although we use panel methods, we take an unconventional approach to panel structure. Our data is observed at the production-cycle level, where each cycle occurs at a fixed location (i.e. farm site). Over the course of many years, a single concession will host multiple production cycles. At the end of each cycle, all fish are removed from the water and the concession must remain fallow until the policy-mandated fallow period is complete and the farm chooses to begin a new production cycle. We treat each production cycle as a distinct, cross-sectional unit. Instead of modeling comparisons between production cycles in calendar time, we use production cycle time. This choice is largely practical. As we describe in the next section, our database covers hundreds of production cycles but relatively few are contemporaneous. A modeling strategy designed to compare cycles that occur in the same calendar week would be left with many poorly identified comparisons. Furthermore, all fish undergo important but ultimately unobserved changes during the production cycle that impact their susceptibility to pathogens. We can control for these changes using cycle-week fixed effects, creating a strategy that draws comparisons in cycle time rather than in calendar time.

Our approach also relies on the idea that, on average, exposure to pathogen dispersal increases with spatial proximity. In other words, a farm will be more impacted by dispersal from a neighboring farm that is close by than a farm that is far away. To take advantage of the variation generated by the fallowing policy, we match each farm with its nearest shared neighborhood water border, measured using seaway distance. A detailed description of this matching process is described in Appendix A.1. This concession-neighborhood pair is fixed throughout the analysis and applies

to all production cycles observed on that farm site. It is the primary mechanism through which treatment is assigned to an observational unit in our experimental design.

An observational unit's treatment status cannot change in response to the treatment of another observational unit. This is the stable unit treatment values assumption (SUTVA) (Angrist et al., 1996). A SUTVA violation might arise in this context if spillover effects from a neighborhood impact a farm that is matched with a different neighborhood. This is, in theory, possible because there are neighborhoods with multiple shared borders. However, we don't believe that this is a major problem in our setting. As the summary statistics in Table 1.1 and 1.2 illustrate, farms are, on average, more than 18 sea-way kilometers away from the second-closest shared neighborhood border.[8] This is far longer than the maximum transmission distance suggested in the epidemiology literature.

This empirical approach is unique within the context of marine disease in aquaculture, both in its use of recurring variation from the neighborhood policy and for its focus on causal inference for program evaluation. Much of the focus in previous work has been on using predictive rather than causal models for measuring the spread of pathogens (Kristoffersen et al., 2013; Rees et al., 2014). This obscures the mechanistic relationships between farms as well as the role of policy in mediating dispersal citation. Price et al. (2017) illustrate that fallowing on individual farms can reduce the hazard of SRS in future production cycles but does not consider the role of coordination in those outcomes. Guarracino et al. (2018) argue that selectively implemented coordinated management zones in Norway targeted at sea lice were ineffective at reducing sea lice outbreaks. However, they do not attempt to address simultaneity in lice prevalence nor do they account for the possibility of selection effects in the location of the coordinated zones. In our setting, this is not a problem as all of the salmon farming activities in Regions X and XI are subject to the neighborhood policy. Peacock et al. (2020) take a more structural approach by estimating a population dynamics model to simulate how coordinated application of chemical therapeutics can limit parasite spillovers from farmed salmon populations to wild populations. To some extent, simultaneity is addressed in their setting through observable migrations of wild salmon which pass in proximity to salmon farms.

[8]The second closest neighborhood would be the source of the SUTVA violation in this case.

Although this variation helps to identify their model parameters, they do not directly use causal inference to evaluate management policy.

1.5.1. Data.

1.5.1.1. *Description of Data Sources.* To conduct this analysis, we compile information from many sources. First, we use the Sistema de Información para la Fiscalización de Acuicultura (SIFA) database, the Informe Ambiental para la Acuicultura (INFA) database, and a third database of antimicrobial prescriptions, all of which are maintained by the Servicio Nacional de Pesca y Acuicultura (SERNAPESCA). SIFA and INFA contain observations on the environmental conditions, production outcomes, and sanitary performance of all salmonid farming operations in Chile. Together, they offer the most complete view of the industry available. An important caveat of SIFA and INFA is that they reflect self-reported observations: each firm is ultimately responsible for monitoring the state of their farms and submitting accurate data to SERNAPESCA.

The SIFA database contains weekly observations of cage-level mortality counts, delineated by cause (including from SRS) along with the species in production, and the total number of fish observed in each cage. We aggregate the cage-level observations up to the level of the production center, creating a database of observations and outcomes measured at the farm by week level (including a count of the total number of cages). The INFA database contains information on the water temperature, salinity, and salinity.[9]

The antimicrobial-use database records all of the prescriptions written by licensed veterinarians in the salmon industry. The prescriptions are disease-specific and denote the number of days in which antimicrobial products will be applied. The prescriptions are also connected to specific farm sites, allowing us to link the prescriptions to the SIFA and INFA databases. For each week of a production cycle, we create a binary variable to indicate whether SRS-related treatment is occurring. In addition, we also calculate the proportion of cages that are treated, a measure of treatment of intensity. Not all of the prescriptions could be accurately matched with the SIFA and INFA databases only the cycles which which could be matched.

[9]The INFA database is distinct from the SIFA database because it is used for sea lice monitoring rather than for infectious disease control.

We also rely on a collection of administrative data that is publicly available from the Subsecretaria de Pesca y Acuicultura (SUBPESCA). These include the geographic coordinates of all aquaculture leases, the boundaries of the production neighborhoods, and the schedule of common fallowing periods for the production neighborhoods since the start of coordinated fallowing policies. The geographic boundaries and concession coordinates allow us to calculate seaway distances between each concession and the nearest shared neighborhood border while the fallowing schedule allows us to infer the start and end dates of individual production cycles.

1.5.1.2. *Filtering.* The SIFA and INFA databases contain information for the entire Chilean salmon industry but we choose to limit our sample in a few important ways. A full description of database construction is included in the Appendix but two key features are particularly important for our identification strategy. First, we focus our analysis on Atlantic salmon, even though Chile's industry produces three different species: Atlantic Salmon, Coho Salmon, and Rainbow Trout. Atlantic salmon is the most important both by volume and economic value and are most relevant for policy analysis. Coho salmon are not required to report ambient environmental data on a weekly basis, severely limiting their utility in our model. They are also thought to be much less susceptible to SRS than Atlantic salmon or Rainbow Trout.

Although we focus on outcomes on Atlantic salmon farms, we choose to include information from all three species when calculating the effective date in which neighborhoods switch from fallow to active as well as when calculating total pathogen prevalence levels. We believe this to be the most appropriate assumption given that both coho salmon and rainbow trout are carriers of the disease and can ultimately, contribute to the overall pathogen load in the system.

Second, we limit our analysis to Region X and Region XI where the majority of salmon farming activity occurs. As Figure 1.2 illustrates, few new concessions have been granted in these two regions in the past decade, making the policy context fairly stable during the analysis. The salmon farming industry has been rapidly expanding into Region XII, leading to the creation of many new neighborhoods since 2010. However, SRS infection does not appear to be an important challenge in this region so we will not include observations from this region as part of the analysis.

1.5.1.3. *Summary Statistics.* Table 1.1 provides summary statistics for the filtered data. Each observation comes from a completely observed production cycle. As we'll illustrate clearly in a later

section, the DiD model uses a subset of this data (where the subset is determined by the specifics of the identifying variation), while the GMM model uses the entire dataset.

TABLE 1.1. Summary Statistics

Variable	N	Mean	Std. Dev.	Min	Pctl. 25	Pctl. 75	Max
SRS Mortality Rate (%)	27867	0.037	0.17	0	0	0.018	7.3
Temperature (C)	27867	11	1.4	3.1	10	12	22
Salinity (ppt)	27867	30	3.5	3	30	32	74
SRS Mortality in Adj. Neighborhood (fish)	27867	564	1870	0	0	325	48373
Dist to Nearest Shared Border (km)	591	10.8	8.5	0.4	4.5	15.3	38.9
Dist to 2nd Nearest Shared Border (km)	591	22.3	13.0	1.7	13.1	29.2	84.1
SRS Treatment (=1)	27867						
... 0	21082	76%					
... 1	6785	24%					
SRS Prescription (=1)	27867						
... 0	25574	92%					
... 1	2293	8%					
Adj. Neighborhood Fallow (=1)	27867						
... 0	22893	82%					
... 1	4974	18%					

Note: This table illustrates the number of observations, mean, standard deviation, minimum observed value, 25th and 75th quantiles, and the maximum observed value. The statistics for Distance to Adj. Neighborhood Border are measured at the cycle level and are not used in any model. SRS Treatment is equal to one on weeks in which medicated feed is delivered to fish. SRS Prescription is equal to one on weeks in which a prescription for medicated feed is written. This does not match treatment weeks because a prescription can call for a treatment regime that lasts more than one week. Adj. Neighborhood Fallow equals one when all the farms in the nearest neighborhood are fallow.

These data offer important insight into the system. First, it is important to highlight that the SRS mortality rate is highly variable, creating large standard errors and suggesting non-linearities in pathogen prevalence. It is also worth noting the mean distance of farms from shared neighborhood borders is 11 km. Recall that the epidemiological community has adopted a 10 km transmission radius for SRS transmission. This may give an early indication that, on average, farms are simply too distant for adjacent neighborhoods to be affected by spillover. Nevertheless, there are a large number of farms that are located in close proximity to neighborhood borders. This is a source of heterogeneity that we will explore.

FIGURE 1.5. Average Observed Pathogen Prevalence

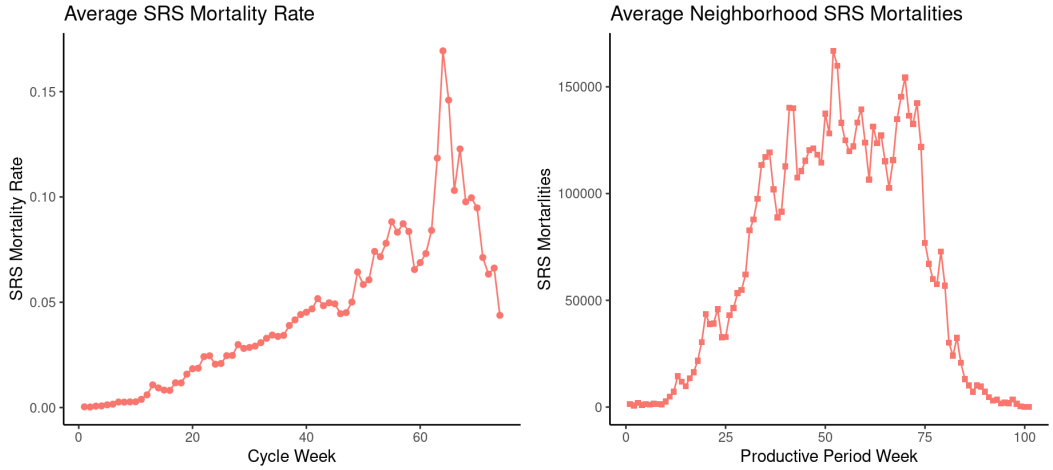

Note: The left figure plots the average SRS mortality rate (in percentage points) for each week of the production cycle. The left panel plots the average number of mortalities attributed to SRS in a neighborhood during each week of a productive period.

In Figure 1.5, we plot the average SRS mortality rate for each week of the production cycle. Pathogen prevalence strongly increases throughout the production cycle, suggesting that our theory accurately reflects average pathogen levels on the farm. When aggregated to the neighborhood, the patterns remain similar with a few important distinctions. Pathogen prevalence increases for much of the productive period. But as farms complete their cycles and begin to harvest, total pathogen prevalence begins to fall. This decline matches the pathogen removal process related to the fallowing described in the theoretical model. These figures also highlight the fact that we can only observe pathogen levels when farms are active. However, the dramatic reduction in host density should dramatically reduce the prevalence of the pathogen, relative to those levels observed during the neighborhood productive period. Nevertheless, the pathogen is endemic and is unlikely to be completely eradicated from the landscape.

1.6. Difference in Differences

Equation 1.6 provides the structural foundation for a differences-in-differences (DiD) model with dynamic treatment effects. We are interested in estimating α but the two scenarios are also mutually exclusive: in any given period, a farm can only be observed in an active state or in a

fallow state. To estimate the equation, we first identify production cycles that are exposed to an adjacent neighborhood that switches from a fallow period to a productive period. This will be the treatment group. In terms of the structural model, treatment assignment is determined by the switch date, $\bar{t}$. From the perspective of any farm, the collective fallowing behavior in the nearest adjacent neighborhood is plausibly exogenous. This means that $\bar{t}$ in Equation 1.6 and treatment assignment is also plausibly exogenous.

We use the following specification for the difference-in-differences model:

$$y_{it} = \lambda_i + \delta_t + \mu t + \sum_{\substack{r \neq 0 \\ \underline{T} \leq r \leq \bar{T}}} \mathbb{1}[D_{it} = r] \cdot \beta_r + \gamma X_{it} + \epsilon_{it} \tag{1.7}$$

where y_{it} is log SRS-related mortality rates during production cycle i in production cycle week t. $D_{it} = 1$ for all periods in which the adjacent neighborhood is in an active period. The parameters of interest are β_r which capture the dynamic effects of the switch from the fallowing period to the productive period on the outcomes of interest, relative to the control cycles. The subscript r refers to the period relative to the start of treatment (this is event-time where treatment starts when $r = 1$). In terms of simulation exercise, r is the period relative to $\bar{t}$ and β_r is a non-parametric approximation of α, whose values we expect that vary over time.

The cycle fixed effects, λ_i, account for time-invariant heterogeneity between production cycles that are relevant for pathogen levels. For example, this would capture heterogeneity due to the characteristics of a particular cohort of fish that are stocked at the farm (e.g. the genetic strain or the quality of sanitary conditions at the hatchery), of the technology on a farm during that cycle (e.g. physical size of the pens, presence of environmental monitors), or of the farm management team (e.g. experience or training). Cycle fixed effects would also account for features that are unique about the particular farm site and relevant for pathogen levels, such as proximity to a freshwater outlet or bathymetry. While these features may not vary between production cycles located at the farm site (recall that we observe many farm sites with multiple cycles within the database), the effect of these features might vary by year. For example, El Niño can drive dramatic fluctuations in rainfall. During a wet year, the impacts of the freshwater outlet might be quite

different on nearby farms than during a dry year. While this might suggest the utility of Year $\times$ Farm controls, we found this practically infeasible. The panel is not balanced in calendar time, creating a large number of very poorly identified Year $\times$ Farm pairs. A farm could, for example, operate in all 52 weeks of one year but only 1 week of the next. A production cycle fixed effect is a tractable, flexible compromise that likely captures much of the key variation.

We also include cycle-week fixed effects to address time-varying features that affect all cycles. Perhaps the most important feature that these controls address are changes in the susceptibility of salmon to pathogens. For example, as noted in a previous section, essentially all salmon are administered a temporary vaccine before they enter the marine stage. This protects the fish during the initial weeks of the production cycle but does not last indefinitely. The inclusion of a time trend, μt is motivated by the simulation exercise and is distinct from cycle-week fixed effects. As the derivation of the structural model suggests, this controls for the average reproduction rate of the pathogen during the cycle. It is worth emphasizing that the model is structured in cycle-week time rather than calendar time. Despite the fact that the database extends over many years, long periods of fallowing create comparatively little overlap across observational units, offering few opportunities for comparison within a calendar-time setting. To address this problem, we use cycle-week time instead. Finally, we also include a vector of cycle-specific, time-varying factors, γX_{it}, including temperature and salinity at the farm site.

Within the discussion of the theoretical model, we acknowledged that farms will respond to disease outbreaks by applying antimicrobial therapeutics. The application of antimicrobials will certainly affect pathogen prevalence (and observed mortality rates), and so we include a binary indicator of treatment during the week as an explanatory variable within X_{it}. However, treatment behavior offers an alternative signal of pathogen prevalence. If farm managers respond to higher observed infection rates by increasing the application rate of antimicrobials, then this indirectly suggests elevated pathogen dispersal rates. It may also be the case that treatment application and disease-related mortalities are simultaneously determined. If, for example, application rates increase after the end of fallowing but observed disease-related mortality rates do not, then this also suggests a dispersal effect that is obscured by the effects of treatment. The difference-in-differences model is sufficiently flexible to offer insight into these effects. To explore this point we estimate

two additional difference-in-differences models where the probability of antimicrobial therapeutics and the proportion of cages treated are treated as dependent variables.

1.6.1. Identifying Assumptions and Robustness Checks. Our design amounts to a comparison between production cycles that are "treated" during different cycle weeks, motivating a staggered timing DiD model.[10] However, treatment is not a standardized unit of exposure but an outcome of a policy that may vary between neighborhoods and over time; we are interested in recovering the average effects. Recent advances in the econometrics literature have illustrated that estimating Equation 1.7 with staggered treatment-timing using ordinary least squares (OLS) will only recover consistent estimates of the average treatment effect when treatment effects are homogeneous (de Chaisemartin and D'Haultfœuille, 2020). Many have proposed estimators that solve this problem under less restrictive assumptions (Callaway and Sant'Anna, 2021; de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021). Our main specification utilizes the estimating procedure outlined by Sun and Abraham (2021) which allows for treatment effect heterogeneity and requires the following identifying assumptions: (1) parallel trends in baseline outcomes (a weaker assumption than that required for OLS) and (2) no anticipatory behavior.

The parallel trends assumption lies at the core of the DiD design. For the Sun and Abraham estimator to be consistent, it must be the case that treated units would have followed the same trajectory as the control units, had they not been treated, conditional on observables. There are no cycles whose nearest adjacent neighborhood remains fallow for all cycle weeks. To create a control group, we use production cycles that are treated later in their production cycles as controls for cycles that are treated earlier. In other words, we compare outcomes between treated cycles and not-yet-treated cycles. Parallel trends will hold if there is no difference between the trajectories of the treated cycles and the not-yet-treated cycles. This is plausible in our setting because the time at which neighborhoods switch from active to fallow is not predictable from the perspective of farms in the adjacent neighborhood. In this case, there should be no difference between the farms that are treated in any given week and those that have not yet been treated, after we condition on cycle week, year, quarter of the year, the application of antimicrobials, salinity, and temperature. The

[10]To be clear, treatment, in this case, has nothing to do with the application of antimicrobial products. It refers to exposure to a switch from fallow to active in the nearest adjacent neighborhood.

exogenous nature of treatment timing also suggests that assumption (2), no anticipatory behavior, is also plausible. No anticipatory behavior implies that farm managers cannot anticipate exposure to the switch from fallow to productive in the adjacent neighborhood and somehow adjust. Given that the timing of exposure is quasi-random, this assumption is reasonable.

To bolster the parallel trends claim, we also have chosen to use only a limited set of cycles that begin operation while their nearest adjacent neighborhood is in a fallow period. This subsample offers the cleanest match to the theoretical model presented in the previous section and reduces the possibility of contamination from the previous cycle. As illustrated in Table 1.2, the subsample is much smaller than the full database. We could utilize a larger sample by including production cycles which begin while the adjacent neighborhood is in production but which continue operating while the adjacent neighborhood completes an entire fallowing period and resumes production. However, we are concerned that exposure to the prior active cycle may lead to a different set of baseline outcomes than those cycles which began production while the adjacent neighborhood fallowed. Nevertheless, we show in the Appendix A.2, that including these cycles does not change the conclusions of the DiD model.

TABLE 1.2. Summary Statistics for DiD Models

Variable	N	Mean	Std. Dev.	Min	Pctl. 25	Pctl. 75	Max
SRS Mortality Rate (%)	4858	0.041	0.19	0	0	0.015	7.3
Temperature (C)	4858	11	1.3	5.3	9.9	12	15
Salinity (ppt)	4858	30	3.8	5	29	32	35
Fallow Exposure (Adj. Neighborhood Fallow =1)	4858	0.21	0.17	0.014	0.067	0.32	0.72
Prescription Written (=1)	3594	0.08	0.27	0	0	0	1
Proportion of Cages Treated	3594	0.045	0.19	0	0	0	1
Dist to Nearest Shared Border (km)	78	10.0	7.9	0.4	3.5	13.6	32.9
Dist to 2nd Nearest Shared Border (km)	78	18.4	8.9	5.6	11.5	24.9	50.3

Note: This table illustrates the number of observations, mean, standard deviation, minimum observed value, 25th and 75th quantiles, and the maximum observed value for the subset of the database used for the three Difference-in-Differences (DiD) models.

1.6.2. DiD Model Estimates. The model estimates from the DiD model are displayed in Figure 1.6. The parameter of interest in the event-study DiD model is the vector of coefficients, β_r. The DiD results do not indicate a worsening of veterinary conditions in the weeks after the neighborhood begins production. The effect of the transition from fallowing to active on SRS mortality rates appears limited, particularly for the first six weeks following the end of fallowing.

While there is some evidence of elevated mortality rates at the end of the observed window, these effects appear both transient and statistically insignificant. Similarly, the probability and intensity of treatment for SRS show essentially no effect after the end of fallowing. Another way to present these results is through a weighted average treatment effect on the treated population (ATT), following Sun and Abraham (2021) and others. These results are presented in Table 1.3 and state that the average effect of exposure is a 0.02 percentage point decrease in the SRS mortality rate, a four percent increase in the probability of a new prescription, and two percentage point increase in the proportion of cages treated. The estimates match the conclusions from the event study: that the effect of exposure is indistinguishable from zero and the point estimates are small.

Nevertheless, this model faces a few key limitations. First, the variation in the width of the confidence intervals reflects the high levels of variation in the SRS mortality rates. This, in combination with the use of a relatively small subsample of the data, reduces the precision of our estimates. In the Appendix, we present the results of a relaxed data filtering process that significantly increases the sample size. We find that the results in Table 1.3 are robust to this alternative database.

A second challenge is that pathogen levels are low during the first weeks of the productive period. This is predicted by the theoretical model and confirmed, on average, in Figure 1.5. At the start of the productive period, fish that are stocked on the farms will be free of SRS and must be reinfected by bacteria that are in the environment. However, low levels of pathogen prevalence will reduce our ability to detect spillover effects. In other words, we may be measuring causal effects which reflect the lowest-intensity spillover effects. In the next section, we will illustrate how a GMM estimator can test for spillovers during moments when they are most detectable.

FIGURE 1.6. Event Study Model Estimates

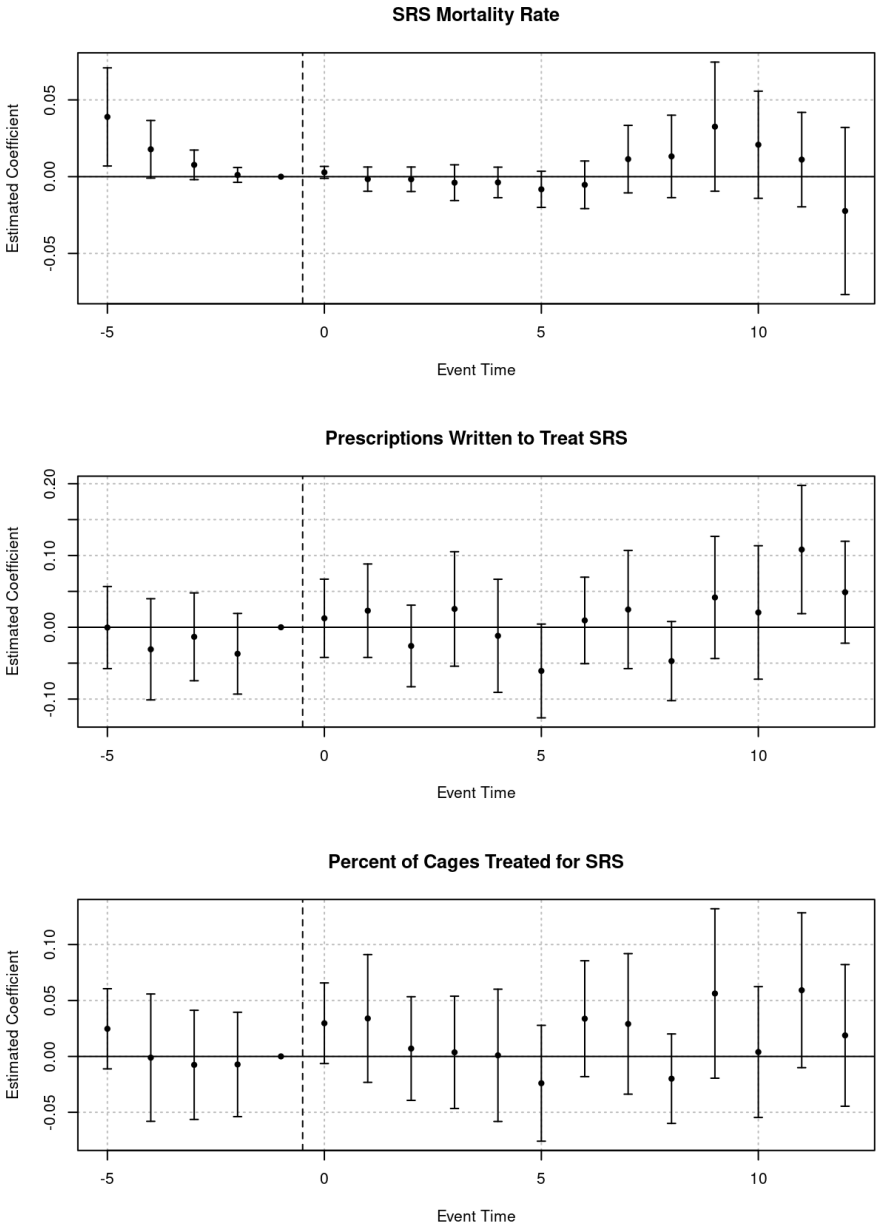

Note: This figure depicts the coefficients of the event study model (β_r in Equation 3) and confidence intervals, calculated at $\alpha = 0.05\%$. The coefficients are calculated relative to the base group at event time -1. The end of fallowing is illustrated by the dashed vertical line at event time -0.5. All models include controls for salinity, temperature, cycle week, and year along with fixed effects for the cycle and cycle-week. Standard errors are clustered at the production-cycle level.

TABLE 1.3. Weighted Average Treatment Effects

Outcome Var.	ATT	Std. Error	t Value	P Value
SRS Mortality Rate	-0.02	0.02	-0.69	0.49
SRS Prescriptions	0.04	0.03	1.22	0.23
Percent of Cages Treated	0.02	0.03	0.65	0.52

Note: This table captures the weighted average treated effect on the treated units, where the weighting scheme is determined by the number of observations in each treatment cohort (Sun and Abraham, 2021).

1.7. Generalized Method of Moments

The DiD strategy is limited by the source of the identifying variation. The (exogenous) switches from fallow periods to active periods are best suited to measuring spillovers during the relatively narrow window after the adjacent neighborhood has just begun production. This is also the moment when pathogen levels in the adjacent neighborhood are lowest (Figure 1.5) and when any spillover will be the hardest to detect. We can increase our chance of detecting spillovers by not limiting ourselves to those observations that fit the identification strategy of the DiD. For this model, we will use all of the data available. We have also already established that the timing of the fallowing periods is exogenous. This same variation can also be exploited in a different model.

Doing so, however, requires additional assumptions over the nature of the spatial spillovers to avoid simultaneity bias. The DiD model solved this by focusing on a point at the end of the fallowing period when the simultaneity problem is likely limited. We will take a different approach for the Generalized Method of Moments (GMM) estimator. In Equation 1.1, recall that both contemporaneous and lagged values of x enter into the expression of pathogen levels on Patch 1. For the GMM estimator, we will assume that $\delta = 0$, eliminating contemporaneous spillovers. We are left with a spatial-temporal autoregressive structure with lagged, spillover effects:

$$x_{1,t+1} = \rho_1 f_{1t} x_{1t} + \rho_1^f (1 - f_{1t}) x_{1t} + \gamma_{21} x_{2t} \tag{1.8}$$

We will use this structure as the basis for the GMM estimator. In the DiD model, pathogen pressure from the adjacent neighborhood does not explicitly enter the estimating equation. Instead, we

effectively use time since the end of fallowing in the nearest adjacent neighborhood as a proxy for pathogen prevalence. In this model, we'll follow 1.8 by explicitly including both a lagged measure of prevalence on farm i as well as a lagged measure of pathogen prevalence in the neighborhood nearest to farm i.

The econometric model for this section is

$$y_{it} = \rho y_{i,t-1} + \beta Y_{i,t-1} + \theta X_{it} + \lambda_i + \delta_t + \epsilon_{it} \tag{1.9}$$

$$Y_{i,t-1} = \sum_{k \in N_i} y_{k,\tau-1} \tag{1.10}$$

where y_{it} is the outcome of interest, and, as in the DiD model, θX_{it}, λ_i, and δ_t represent time-varying covariates, cycle fixed effects, and cycle-week fixed effects respectively. We will not simultaneously estimate pathogen levels on interconnected patches as is common in other spatial settings (LeSage and Pace, 2009). Instead, we are simply interested in understanding how variation in the adjacent neighborhood drives pathogen prevalence during cycle i. To this end, we use the sum of disease-related mortalities observed in the nearest adjacent neighborhood as a measure of pathogen prevalence. This is captured by $Y_{i,t-1}$. Let farm i be paired with its nearest adjacent neighborhood, N_i. $Y_{i,t-1}$ is the sum of pathogen prevalence across all k farms in N_i. Although the model remains indexed in cycle-week time, the measure of pathogen prevalence in the adjacent neighborhood must be linked in calendar time. With some abuse of notation, let τ be the calendar week on farm k which matches with the cycle week on farm i. The coefficient on this Y captures the spillover effect from the neighborhood onto the farm and is the term of interest.

As in the DiD model, it is also important to control for unobserved heterogeneity across observational units using fixed effects, λ_i. However, the inclusion of fixed effects in an autoregressive model will bias the parameter estimates (Nickell, 1981). To account for this, we will use an estimator in the spirit of the Difference Generalized Method of Moments Estimator described in Arellano and Bond (1991), in which time-invariant heterogeneity is eliminated through first differences. However, the additional lagged terms that are introduced through the first-difference operator create endogeneity in both the autoregressive term as well as in the lagged measure of prevalence in the adjacent neighborhood. We will use two sets of instruments to identify these terms.

First, following the logic of Arellano and Bond, we will use deep lags to identify both $y_{i,t-1}$ and $Y_{i,t-1}$. The identifying assumption required is sequential exogeneity in the lagged explanatory variables, such that, for $t \geq 2$,

$$E[\epsilon_{it} y_{i,t-1}] = 0 \tag{1.11}$$

$$E[\epsilon_{it} Y_{i,t-1}] = 0 \tag{1.12}$$

which allow us to define the following moment conditions

$$E[(\epsilon_{it} - \epsilon_{i,t-1}) y_{i,t-k}] = 0 \tag{1.13}$$

$$E[(\epsilon_{it} - \epsilon_{i,t-1}) Y_{i,t-k}] = 0 \tag{1.14}$$

for $t \geq 3, k \geq 2$. In other words, we use the collection of deep lags, $y_{i,t-k}$ and $Y_{i,t-k}$, $t \geq 3, k \geq 2$, as instruments for $y_{i,t-1}$ and $Y_{i,t-1}$. Although these conditions are restrictive, they are not unjustified considering Equation 1.8. Taken at face value, 1.8 describes pathogen levels in x_{t+1} as completely determined by pathogen values in the preceding period, x_t. In other words, prevalence effectively follows a Markovian process, where x_{t-1} only impacts x_{t+1} via its effect on x_t. This is the autoregressive structure required by sequential exogeneity.

We also use time since the end of fallowing in the adjacent neighborhood as an additional instrument for $Y_{i,t-1}$. To understand why this is a valid instrument, recall from Equation 1.3, that pathogen levels in either patch can be expressed as a function of time since fallowing. As in the previous section, time since fallowing remains exogenous from the perspective of a farm in Patch 1, as it is determined by the collective response to a fallowing schedule set by federal regulators. Using a strict interpretation of the models presented in this section, we might conclude that $t - \bar{t}$ contains no additional information that is not already captured by the lagged values of pathogen prevalence. After all, Equations 1.4 and 1.8 illustrate that pathogen prevalence can be equivalently expressed as either a function of lagged levels or as a function of time. However, our claim that lagged values are completely uncontaminated by contemporaneous spillovers from another neighborhood is untestable. This contamination is not an issue for $t - \bar{t}$. However, $(t - \bar{t})$ amounts to an imprecise, average trend. It does not, for example, express any association with pathogen levels

that are contained within the constants, c and k, nor will it account for any other stochastic features affecting pathogen levels (and not expressed in our structural model). For example, time since fallowing will not capture the rate of treatment application in the neighborhood that will ultimately moderate the rate of pathogen growth. Therefore, both instruments have strengths and weaknesses; we will include both in our preferred model specification.

Finally, we will also explore the inclusion of antimicrobial treatment as a control variable. Pathogen levels will be influenced by the application of antimicrobials. However antimicrobial treatments are non-randomly allocated. Farm managers will only treat when SRS mortality levels are high so treatment is subject to selection. [11] If treatment is correlated with mortality in the adjacent neighborhood, $Y_{i,t-1}$, then including treatment without addressing the endogeneity will bias the estimate of β. But, even if they are not correlated, endogenous control may make the model inconsistent. Given the clear relevance of treatment to mortality rates, we will estimate two specifications, with and without antimicrobial treatment, and explore the robustness of our results.

1.7.1. GMM Model Estimates. Before presenting the model results, it is important to check that the data meet the required modeling assumptions. Difference GMM will produce unbiased estimates when the instruments are uncorrelated with the residuals and when differenced residuals are not autocorrelated beyond the second order. We test these assumptions using the Sangan-Hansen Test and the test for autoregression suggested in Arellano and Bond (1991) and present the results in Table 1.4. Given that the p-values for the test statistics associated with the Sangan-Hansen test are all large, there is no evidence of correlation between the residuals and the instruments. As expected in a dynamic panel model, the tests for auto-regressive behavior do indicate that the residuals are autocorrelated of order one but not order two. This offers support that we are using sufficiently lagged instruments for our purposes.

Models (1) and (2) in Table 1.4 contain estimates of Equation 1.9 using Difference GMM. The parameter of interest is the lagged measure of pathogen prevalence in the adjacent model, β in Equation 1.9, and lag asinh(Adj SRS) in Table 1.4. Our preferred model is (2). Interpreted at face value, the model suggests that a one percent increase in SRS mortalities in the adjacent neighborhood causes a 0.0002 percent increase in the farm SRS mortality rate, an estimate that

[11] Recall that, by law, farms are not allowed to apply antibiotics as a prophylactic measure.

is not statistically different than zero. Furthermore, the model estimates confirm that disease prevalence is dominated by an autoregressive process. Coefficients on the lagged dependent variables are large and significant. The model offers a similar conclusion for a firm response to disease incidence.

The null result is also robust to the exclusion of lagged treatment as a regressor, as shown in model (1). As in the DiD model, lag treatment refers to the condition in which medicated feed is actually applied to cages. This contrasts with a prescription that reflects a direct behavioral response from the farm manager but which is unlikely to affect the SRS mortality rate. Treatment should be correlated with the mortality rate, so it is not surprising that the inclusion of treatment impacts the lagged dependent variable. However, it does not appear to have a large impact on the effect of SRS mortalities in the adjacent neighborhood. It is also important to note that the coefficient on lag SRS Treatment is large and positive, even though antimicrobial treatments are meant to reduce SRS mortality rates. Farm managers choose to apply antimicrobials only when mortality rates are high, creating a positive correlation between antimicrobial applications and the SRS mortality rates. As a result, SRS treatment is subject to endogeneity.

The endogeneity of treatment will impact the coefficient on lag asinh(Adj SRS) (β) if the two variables are correlated. A key way that this could occur is if there is a causal relationship between mortalities in the adjacent neighborhood and antimicrobial applications. If information regarding an outbreak is shared between neighborhoods, then a veterinarian might be more likely to prescribe treatment or to prescribe more cages for treatment. In theory, this effect might also appear in the application of treated feed (lag SRS Treatment) but observed prescriptions offer the closest connection to the behavioral mechanism of interest. However, models (3) and (4) suggest that there is no positive relationship between SRS mortalities in the adjacent neighborhood and the probability or intensity of prescriptions to treat SRS.

1.8. Discussion

Marine disease in aquaculture represents a challenging spatial-dynamic production externality. Pathogens are transmissible over long distances, suggesting the need for coordinated control efforts across space. However, determining the appropriate spatial scale for coordination is challenging,

TABLE 1.4. Model Estimates

	Dependent variable:			
	asinh(SRS rate)		Pr(Prescription)	% Cages Prescribed
	(1)	(2)	(3)	(4)
lag asinh(SRS Rate)	0.730***	0.723***	0.115***	0.075***
	(0.027)	(0.025)	(0.030)	(0.020)
lag asinh(Adj SRS)	−0.001	0.0002	−0.009**	−0.003
	(0.002)	(0.001)	(0.003)	(0.002)
lag(SRS Treatment)		0.008***	0.007	0.002
		(0.002)	(0.008)	(0.005)
lag(Prescription)			0.016	
			(0.010)	
lag(% Cages Prescribed)				−0.023***
				(0.006)
asinh(temp)	0.039***	0.020	0.013	−0.014
	(0.013)	(0.013)	(0.042)	(0.026)
asinh(salin)	−0.020*	−0.012	−0.036	0.011
	(0.011)	(0.010)	(0.035)	(0.021)
Observations	27867	27867	27867	27867
Cycle Week FE	X	X	X	X
Sargan-Hansen Test	122.69	221.67	283.08	279.24
Arellano-Bond AR(1)	-7.16***	-7.20***	-17.94***	-17.36***
Arellano-Bond AR(2)	-0.78	-0.77	0.62	0.38

Note: This table presents the results from models using difference generalized method of moments (GMM). SRS Rate refers to mortality rate attributable to *Salmonid Rickettsial Syndrome*, and % of cages treated refers to the proportion of total cages that are treated for SRS. For the covariates, Adj SRS refers to total SRS mortalities observed in the nearest adjacent neighborhood and is the outcome of interest for this study. The operator, asinh(), denotes the inverse hyperbolic sine. Robust standard errors are reported below the coefficient estimates following Windmeijer (2005). The table also presents test statistics for three different tests. The alternative hypothesis of the Sargan Hansen test is that instruments are correlated with the residuals, violate the exclusion restriction, and are therefore invalid. Under the null hypothesis, the test statistic is distributed Chi-Squared with degrees of freedom equal to the difference between the number of instrumemts and the number of regressors. The altnernative hypothesis in the Arellano-Bond test is that differenced residuals exhibit autocorrelation. Here, we test for autocorrelation of degree 1 and degree 2. The test statistic is distributed standard normal under the null hypothesis. *p<0.1; **p<0.05; ***p<0.01.

particularly in a complex, heterogeneous landscape. We investigate this issue within the context of Chile's salmon aquaculture industry. Evidence suggests that Chile's coordinating policy (the neighborhood system) was not explicitly designed to reflect the spatial scale of pathogen dispersal. Given this point and the epidemiological literature suggesting the transmissibility of marine pathogens across long distances, we expected to find evidence of transmission between coordinated groups of farms (neighborhoods). Our analysis offers little evidence that pathogens readily spread between neighborhoods, a finding that is robust to multiple modeling approaches and alternative specifications. The results also suggest little evidence of a behavioral response to variation in pathogen pressure.

1.8.1. Econometric Considerations. In contrast to the common approach in the epidemiological literature, we seek to identify causal relationships between policy-defined groups of farms. Both of the identification strategies that we use are explicitly informed by our model of disease transmission, allowing us to address some of the common econometric challenges for empirical studies in spatial-dynamic settings. Nevertheless, this study faces some key empirical limitations. While the point estimates of the DiD design are all fairly close to zero, tighter confidence intervals would strengthen our claim of a null result. There are at least three explanations for the relative imprecision of the results. First, the data filtering used for the preferred DiD specification significantly reduces the size of the dataset available for estimating the model parameters. We do this to make our parallel trends assumption as credible as possible but show in the Appendix that the DiD results are robust to less restrictive assumptions. Second, transmission is undoubtedly impacted by unobserved features of the landscape. In the paper, we argue that these features are not correlated with the time at which an adjacent neighborhood switches from fallow to active nor are they correlated with the pairing of a farm to its nearest shared border. Nevertheless, controlling for the direction and magnitude of surface currents, for example, would likely improve the precision of our estimates. Bravo et al. (2020) take an important step in this direction by utilizing a hydrodynamic model of Region X to simulate SRS transmission. However, this type of model has yet to be paired with empirical analysis. Third, non-linearities in the spread of disease increase the variance of observed SRS levels. To some extent, this problem is predicted in the theoretical model. But in

the context of a relatively small sample size, it may lead to poor identification of the cycle-week fixed effects at controlling for average pathogen levels on the farm.

From an epidemiological perspective, it is worth emphasizing that, technically, we do not observe pathogen prevalence. Rather, we use disease mortality as a proxy for disease prevalence because prevalence data is not available. This approach matches the epidemiological literature but is not without weaknesses. Most importantly, disease mortality is a lagging indicator of disease presence; in other words, disease presence will always precede disease mortality. However, this lag impacts both the independent and dependent variables. Therefore, on average, it likely has little impact. In fact, the DID design is robust to this conceptual problem because it exploits variation in time since fallowing, rather than observed SRS mortalities. Furthermore, if observed SRS mortalities are a lagging indicator of SRS transmission, then the first-stage relationship between time since the end of the fallow period and observed SRS mortalities is likely an underestimate of SRS presence.

Finally, it is worth reflecting on the role of theoretical structure in our analysis. One of the key empirical barriers to measuring the causal relationships between neighborhoods is simultaneous dispersal. As we showed in the theoretical model, it is plausible that pathogen levels in adjacent neighborhoods are jointly determined. Quasi-random variation in pathogen prevalence generated by the fallowing schedule allows us to bypass this problem using the DiD model with relatively few additional assumptions about the structure of that dispersal process. However, the available scope is limited. By contrast, the GMM model relies explicitly on more restrictive assumptions over the spatial-dynamic processes. These restrictions are not readily testable and lead to a complex estimator with many moment conditions. But it does offer traction on a challenging empirical problem. This trade-off is common in the analysis of coupled human-natural systems (Ferraro et al., 2019).

1.8.2. Policy Implications. This analysis represents an important step toward a cost-benefit evaluation of the neighborhood policy. If there is low dispersal between neighborhoods, then this suggests that the neighborhoods are, on average, appropriately designed to match pathogen transmission. As prior simulation work has illustrated (Werkman et al., 2011), a neighborhood system with low connectivity between neighborhoods will likely experience fewer losses than a neighborhood system with high levels of connectivity between neighborhoods. Thus, our results suggest

that the system is likely experiencing lower disease-related costs than an alternative system with higher levels of connectivity. Given that the policy was not designed to match the features of the landscape relevant for the externality, it is also significant that this second-best policy is likely generating benefits for the industry.

Nevertheless, this result should be interpreted cautiously as we do not consider the costs associated with coordination. Top-down policies like the neighborhood system tend to be more costly than comparable market-based approaches. For example, economic theory suggests that a firm operating multiple farm sites near one another would have a strong incentive to internalize the costs of a production externality. A firm is also well-positioned to coordinate disease control efforts across multiple farm sites in a least-cost manner. However, the coordinating policy does not incentivize spatial clustering; in fact, the data suggest that firms tend to spread out their farms across many different neighborhoods. There are many possible reasons why this might be the case and a full analysis of the spatial behavior of farm utilization is beyond the scope of this paper. But it serves as an important reminder that alternative coordinating mechanisms that are relevant for disease management do exist and there is no evidence to suggest that neighborhood policy is more efficient than an alternative.

Future work might reasonably consider the extent to which other adjustments are relevant for pathogen dispersal. In our data, the average farm site is more than ten kilometers from the nearest adjacent neighborhood, a fairly long distance in terms of the dispersal range of a free-floating pathogen. It could be the case that farms are choosing to utilize farm sites that are located on the interior of production neighborhoods, rather than those on the borders. Nevertheless, most farms experience SRS during their production cycle, regardless of their location, suggesting that this isolation from neighborhood borders does not provide protection. Additional work on the heterogeneity of spillover effects between neighborhoods may provide additional insight and identify regions where the largest spillovers might exist.

Future work might also investigate the extent to which we can detect heterogeneity in spillover effects, particularly as it relates to the distance between a farm and its nearest adjacent neighborhood. It is plausible to believe that spillover effects might be largest for farms that are located in closest proximity to the border. These farms would operate as key links in the farm network,

facilitating pathogen spread between neighborhoods. In other words, the external costs generated by these sites might be disproportionately large. Identifying these sites would be useful for policy-makers looking to reduce the transmission of disease through the targeted removal of farming concessions.

If we cannot detect spillovers between neighborhoods, then it could be possible that each neighborhood is larger than necessary and coordinates too many constituent farms. But this analysis does inform that question. Our results suggest that the spatial scale of the neighborhoods may be sufficiently large to capture the gains from coordination, within the context of SRS. However, there are many other marine pathogens and parasites that are important for the salmon aquaculture industry. This policy likely affects any pathogen that is transmitted horizontally from farm to farm. Our approach lends itself particularly to endemic, rather than exotic pathogens. While many exotic pathogens can induce aggressive responses from the industry to avoid catastrophic losses, endemic pathogens lead to relatively small losses in productivity. In this setting, nearly all producers must contend with endemic pathogens, making an industry-wide neighborhood policy reasonable.

The control of SRS is also relevant for human health because it is responsible for more than 90% of the salmon industry's use of antibiotics. Concern over the development of resistance has led many to call for a reduction in the rate of antibiotic applications in the salmon industry. One way to achieve this reduction is to reduce the transmission of pathogens between farms, one of the goals of the neighborhood program The results of our study indicate that these efforts should focus on transmission between farms within a neighborhood, rather than on transmission between neighborhoods.

1.9. Conclusion

In this paper, we take a first step towards the evaluation of a spatial management policy designed to offer protection from marine disease in Chile's salmon aquaculture industry. We measure transmission between coordinated management units as well as the effects of coordinated fallowing periods. Our strategy departs from the existing literature by focusing on the identification of causal linkages between management units, overcoming important econometric challenges common to systems with spatial-dynamic features.

In contrast to the literature, we find limited evidence of spillovers between spatial management units for both Salmonid Rickettsial Syndrome (SRS) in Chile's aquaculture industry. Our findings suggest that the neighborhood system may be appropriately scaled to the externalities generated by SRS, despite existing evidence of long-distance dispersal. The neighborhood policy is a large and ambitious solution to controlling costly marine diseases. However, it is critically important to match the arrangement of spatial management units to the landscape. Our work offers a first assessment of whether the policy matches a heterogeneous landscape.

This study offers an important first step toward a full evaluation of the costs and benefits of the system. We also lay a foundation for future research agenda that would characterize the costs of the neighborhood policy, incorporate firm behavior, and characterize the heterogeneity of spillovers across the production region. All of this information can be informative for policy-makers looking to improve the effectiveness and efficiency of the spatial management system.

CHAPTER 2

Estimating the Effect of Endogenous Antibiotic Use

2.1. Introduction

Although public policy might be effective for reducing industry costs of disease transmission, individual farmers also have a strong private incentive to control pathogen spread. Accurately characterizing these incentives will be critical for balancing disease control with antibiotic use. This requires understanding and empirically estimating the impact of antibiotic use on farm-level disease dynamics.

However, conventional models in epidemiology tend to abstract away from human behavior. This simplification can make it challenging to draw realistic conclusions about policy because disease risk tends to induce behavioral responses in people or firms that are consequential for disease dynamics. Many have shown the value of explicitly including the linkages between human behavior and the epidemiological or biological dynamics in coupled human-natural systems (Fenichel et al., 2011; Ferraro et al., 2019). But these same linkages can also confound econometric estimates of key model parameters that may be important for credible policy simulation.

In this chapter, we illustrate some of the empirical challenges associated with estimating the parameters of a system that is jointly determined by epidemiological and economic features. Starting with a simulation model, we show that the selective application of antibiotic treatments in response to a disease outbreak will confound estimates of treatment effectiveness that rely on observational data. We then show that a control function (Wooldridge, 2015) can be used to empirically test for the presence of behavioral feedback (i.e. non-random selection) and is a reasonable estimator to recover a policy-relevant parameter of an epidemiological model.

We apply this approach to Chile's salmon aquaculture industry, where producers routinely apply antibiotics to control outbreaks of endemic disease. In order to measure the benefits of antibiotic use and accurately simulate the dynamics of an infection, we must estimate the effectiveness of antibiotic

treatment. Commercially available antimicrobial products are tested in a laboratory setting where veterinarians can control exposure to a pathogen, manipulate the allocation of treatment, and observe the progress of the infection for each individual under study. In the field, antibiotics are not applied to individual sick fish but broadcast to all fish in a pen, regardless of their infection status.[1] Because of this, treatment effectiveness estimates from lab experiments do not map into a farm-level model of disease dynamics, suggesting that these parameters need to be estimated from observational data.

However, farm managers do not randomly apply antibiotics to their farms. They react to farm conditions and selectively apply antimicrobial treatments when it is in their best interest. When an outbreak is spreading rapidly, a farm manager is much more likely to treat than if there is no infection present. Furthermore, it is likely that there are other variables (e.g. fish behavior) that enter this decision-process that are readily observable to the farm manager but that are unobservable to the econometrician. If the farm manager chooses to treat based on these unobservables, then conventional estimates of treatment effectiveness will be biased. In this paper, I show how a control function, a well-known solution to endogenous explanatory variables, can be directly motivated from an epidemiological model, where treatment is endogenous, and be used to estimate the parameters of the model.

After illustrating the efficacy of this approach with a Monte Carlo experiment, we then bring the estimator to Chile's salmon aquaculture industry to estimate the effectiveness of antibiotic treatments to control Salmonid Rickettsial Syndrome (SRS). We take advantage of exogenous variation in the application of antibiotic treatments, generated by environmental conditions to measure the marginal effect of antibiotic use on SRS mortality rates. Although the internal validity of our empirical estimates suffer from weak instrumental variables, we believe that the application offers a strong proof of concept for future work.

This chapter informs two sets of literature. First, estimating the effectiveness of antimicrobial treatments is of exceptional importance to Chile's salmon aquaculture industry. Chile's salmon producers utilize antibiotics at rates that far exceed nearly all other producing countries. As noted in the previous chapters, many public health and policy experts have grown concerned over

[1]Recall that a farm is typically made up of multiple pens.

ecosystem impacts of prolonged exposure to antibiomicrobial use (Quinones et al., 2019) as well as over the development of antibiotic resistance to therapeutics that are important for both fish and human health (Schar et al., 2020). This has driven interest in evaluating the extent to which the benefits of antibiotic applications are worth the cost. To this end, our study is similar to Price et al. (2016). The authors use observational data from Chile's salmon farming industry to measure the the effectiveness of antibiotics for the control of salmon disease. The authors find that antibiotic applications fail to control disease outbreaks in more than 40 percent of applications. However, they do not estimate a parameter that can be readily linked to a farmer's decision to apply antibiotics during production, nor can it be used to simulate counterfactual disease outbreaks under alternative treatment schemes. By contrast, we focus on obtaining a measure of the marginal impact of antibiotic treatment on the weekly, disease-attributed mortality rate. This is an important first step towards measuring the expected marginal benefits of antibiotics that might impact a farmer's decision to treat their fish.

Second, this paper contributes to a growing literature linking epidemiological and economic models (Auld, 2003; Castonguay et al., 2020; Fenichel et al., 2011; Geoffard and Philipson, 1996; Kremer, 1996; Perrings et al., 2014). It is well-known that epidemiological models make strong assumptions over the behavioral response of agents to disease risk. This impacts the extent to which epidemiological models can realistically capture disease dynamics. However, far fewer studies have sought to bring these dynamic models to the data. Agricultural settings provide a unique opportunity to integrate behavioral insights with epidemiological models because the setting is relatively data-rich and the behavioral incentives are relatively simple, in comparison to settings focused on disease transmission in human populations. While we cannot fully observe a disease outbreak, we do observe a strong proxy for disease prevalence (disease mortalities) along with antibiotic applications. This allows us to test the extent to which the hypothesized linkages between economic behavior and the underlying disease dynamics are present in the observable data.

The remainder of this paper is organized in the following manner. Section 2.2 describes an epidemiological model in which we show that treatment behavior is fundamental to measuring treatment effectiveness. Section 2.3 illustrates how treatment effectiveness can be recovered empirically, demonstrating the efficacy of the approach using a Monte Carlo experiment. Section 2.5

describes how the case study and presents results for the effectiveness of antibiotic treatments in Chile's salmon aquaculture industry. Section 2.7 concludes.

2.2. Model

Consider the following model, based on the epidemiological models described in Anderson and May (1979):

$$S_{t+1} = S_t - (\beta_1 + \lambda_t)\frac{S_t I_t}{N_t} - \beta_2 S_t + \beta_3 R_t \tag{2.1}$$

$$I_{t+1} = I_t + (\beta_1 + \lambda_t)\frac{S_t I_t}{N_t} - \beta_4 I_t T_t - \beta_5 I_t \tag{2.2}$$

$$R_{t+1} = R_t - \beta_3 R_t + \beta_4 I_t T_t \tag{2.3}$$

where S_t is the number of susceptible individuals, I_t is the number of infected individuals, R_t is the number of recovered individuals, and N_t is the total number of individuals in period t, such that $N_t = S_t + I_t + R_t$. To simplify notation, define $X_t = [S_t, I_t, R_t]'$ as a vector containing all three states. λ_t is a stochastic term, capturing shocks to the infection rate T_t represents treatment. Treatment impacts the infection dynamics by transitioning fish from an infected state into a recovered state. We assume a non-linear relationship between the impact of treatment and the number of infected individuals, with the rate of recovery increasing with the number infected. This does not necessarily mean that treatment is more effective when there are more infected individuals. It simply implies that more individuals will transition from infected to recovered for a given unit of treatment.[2] We also assume that treatment has no impact on fish that are already in the susceptible or recovered states.

Next, define the mortality rate, m_t to be

$$m_t = \frac{\beta_5 I_t}{N_t} \tag{2.4}$$

[2]Price et al. (2016) argue that the effectiveness of treatment may fall as infection progresses due to low uptake rates of feed based treatment by sick fish. However, I does not measure the extent of an infection in an average fish. It only accounts for the number of infected individuals. For the sake of simplicity, we do not attempt to integrate metabolic feedback mechanisms into the model.

Plugging 2.4 into 2.2 , the expected mortality rate in $t+1$, conditional on X_t, is [3]

$$E[m_{t+1}|X_t, T_t] = \frac{\beta_5}{N_{t+1}}\left(I_t + (\beta_1 + E[\lambda_t|X_t, T_t])\frac{S_t I_t}{N_t} - \beta_4 I_t T_t - \beta_5 I_t\right) \tag{2.5}$$

Suppose that we are interested in the effect of treatment, T_t, on the future mortality rate, m_{t+1}. This is the difference in expected mortality rate, given treatment and no treatment, respectively. This can be succinctly expressed using a potential outcomes framework:

$$\Delta(X) = E[m_{t+1}|X_t, T_t = 1] - E[m_{t+1}|X_t, T_t = 0] \tag{2.6}$$

$$= \frac{\beta_5}{N_{t+1}}\left(\frac{S_t I_t}{N_t}\left(E[\lambda_t|X_t, T_t = 1] - E[\lambda_t|X_t, T_t = 0]\right) - \beta_4 I_t\right) \tag{2.7}$$

It is immediately clear that our assumptions over treatment behavior, T, are consequential. If farms apply treatment in response to farm conditions (e.g. elevated mortality rates), then the observations in which treatment occurs will be systematically different than those in which treatment does not occur. In other words, treatment is endogenous. We illustrate this analytically using two contrasting behavioral assumptions for T_t: (1) treatment that is randomly allocated and (2) treatment behavior that is correlated with both observable and unobservable terms components.

2.2.1. Random Treatment. First, suppose that treatment is applied randomly. This implies that treatment is unrelated to the distribution of the shock to the infection rate, λ:

$$E[\lambda_t|X_t, T_t = 1] = E[\lambda_t|X_t, T_t = 0] = E[\lambda|X_t] \tag{2.8}$$

and Δ collapses to

$$\Delta^R(X) = E[m_{t+1}|X_t, T_t = 1] - E[m_{t+1}|X_t, T_t = 0] \tag{2.9}$$

$$= -\frac{\beta_5\beta_4 I_t}{N_{t+1}} \tag{2.10}$$

$$= -\frac{\beta_4 m_t N_t}{N_{t+1}} \tag{2.11}$$

The non-linear relationship between the impact of treatment and the number of infected appears in Δ. But it is more complicated as it is scaled to reflect the impacts on the mortality rate, rather

[3]It is worth emphasizing that, in this model, $N_{t+1} = N_t - \beta_2 S_t - \beta_5 I_t$ and is, therefore, predetermined. This means that it does not need to remain in the conditional expectation operator.

than on the number of infected individuals. After the fish have been stocked, it is nearly always the case that $N_t \geq N_{t+1}$ due to losses from disease and natural mortality. Assuming $\beta_4 \geq 0$,

$$0 \leq |\Delta^R(X)| \leq \beta_4 m_t \tag{2.12}$$

2.2.2. Selection into Treatment Based on Unobservables. Now suppose that the farm manager selects into treatment based upon observed conditions. Assume that the probability of treatment can be described by the following rule:

$$Pr(T_t = 1|X_t) = \begin{cases} 0 & m_t + \eta_t \leq \zeta \\ f(m_t) + \eta_t & \zeta < m_t + \eta_t < \upsilon \\ 1 & m_t + \eta_t \geq \upsilon \end{cases} \tag{2.13}$$

where ζ and υ are constants while η_t is a random variable. The effect of treatment can now be expressed as the following:

$$\Delta^S(X) = \frac{\beta_5}{N_{t+1}} \left(\frac{S_t I_t}{N_t} (E[\lambda_t | X_t, \eta_t > \zeta - m_t] - E[\lambda_t | X_t, \eta_t \leq \zeta - m_t]) - \beta_4 I_t \right) \tag{2.14}$$

If $E[\eta_t \lambda_t] = 0$, then η is simply another source of stochasticity that will not bias our estimates of the treatment effect. If this is the case, Δ^S collapses to Δ^R. When $E[\eta_t \lambda_t] \neq 0$, selection into treatment will bias the estimated treatment effect. This is clear from our expression of Δ^S, rewritten below in terms of m_t

$$\Delta^S(X) = \underbrace{\frac{m_t S_t}{N_{t+1}} \left(E(\lambda_t | X_t, \eta_t > \zeta - m_t) - E(\lambda_t | X_t, \eta_t \leq \zeta - m_t) \right)}_{\text{Selection Effect}} \underbrace{- \frac{\beta_4 m_t N_t}{N_{t+1}}}_{\substack{\text{Effect when treatment} \\ \text{is random}(\Delta^R)}} \tag{2.15}$$

If treatment is more likely when there there is a positive shock to the infection rate, then $E[\eta_t \lambda_t] > 0$. In this case, the selection effect is positive and the relationship above suggests that $\Delta^R(X) \leq \Delta^S(X)$. In other words, without correcting for selection, empirical estimates of treatment effectiveness will likely be biased towards zero.

2.3. Econometric Model

In order to estimate Δ^S, we need to estimate the structural parameters of the compartment model described in the previous section. However, it is often the case that we cannot observe all of the state variables of interest. For our purposes, assume that we can observe the disease-related mortality rate, m_t, non-disease mortality rate, q_t, the number of fish on the farm, N_t, and treatment, T_t. In terms of the structural parameters, $m_t = \beta_5 I_t/N_t$ and $q_t = \beta_2 S_t/N_t$. It follows that N_{t+1} is predetermined in t because it's components are also observable:

$$N_{t+1} = N_t - \beta_5 I_t - \beta_2 S_t \tag{2.16}$$

$$= N_t(1 - m_t - q_t) \tag{2.17}$$

It follows that:

$$E[m_{t+1}|X_t, T_t] = \frac{\beta_5}{N_{t+1}}\left(I_t + (\beta_1 + E[\lambda_t|X_t, T_t])\frac{S_t I_t}{N_t} - \beta_4 I_t T_t - \beta_5 I_t\right) \tag{2.18}$$

$$= \frac{1}{N_{t+1}}\left(\beta_5 I_t + \frac{\beta_5\beta_1 S_t I_t}{N_t} - \beta_5\beta_4 I_t T_t - \beta_5^2 I_t\right) + \frac{\beta_5 S_t I_t E[\lambda_t|X_t, T_t]}{N_t N_{t+1}} \tag{2.19}$$

$$= \frac{1}{N_{t+1}}\left(m_t N_t + \beta_1 S_t m_t - \beta_4 m_t N_t T_t - \beta_5 m_t N_t\right) + \frac{S_t m_t E[\lambda_t|X_t, T_t]}{N_t(1 - m_t - q_t)} \tag{2.20}$$

$$= \frac{m_t}{1 - m_t - q_t}\left(1 - \beta_5 + \frac{\beta_1 q_t}{\beta_2} - \beta_4 T_t\right) + \frac{q_t m_t E[\lambda_t|X_t, T_t]}{\beta_2(1 - m_t - q_t)} \tag{2.21}$$

$$= (1 - \beta_5)x_t + \frac{\beta_1}{\beta_2}x_t q_t - \beta_4 x_t T_t + \frac{x_t q_t}{\beta_2}E[\lambda_t|X_t, T_t] \tag{2.22}$$

where $x_t = m_t/(1 - m_t - q_t)$.

This formulation suggests the following econometric model:

$$m_{t+1} = \theta_1 x_t + \theta_2 x_t q_t + \theta_3 x_t T_t + \psi_t \tag{2.23}$$

where ψ_t captures the unobserved shock to the infection rate described in the structural model.

$$E[\psi_t|X_t, T_t] = \frac{x_t q_t}{\beta_2}E[\lambda_t|X_t, T_t] \neq 0 \tag{2.24}$$

The treatment effect of interest in terms of the econometric model is:

$$\Delta(X) = E[m_{t+1}|X_t, T_t = 1] - E[m_{t+1}|X_t, T_t = 0] \tag{2.25}$$

$$= \theta_3 x_t + E[\psi_t|X_t, T_t = 1] - E[\psi_t|X_t, T_t = 0] \tag{2.26}$$

This illustrates the econometric challenge, foreshadowed in the previous section. The true effect of treatment is determined by the "structural" parameter, β_4 (and in the equivalent econometric model, θ_3). However, when producers select into treatment, $E[\lambda_t|X_t, T_t] \neq 0$ which implies $E[\psi_t|X_t, T_t = 1] \neq E[\psi_t|X_t, T_t = 0]$. As a result, selection based on unobservable features in the system will bias our estimates of θ_3.

Bias generated from selection into treatment is a well-studied topic within econometrics (Hausman, 19 something). The simplest approach is to instrument for selection and estimate the outcome regression using two-stage least squares (2SLS). However, given the non-linear relationships between x_t, q_t, and λ_t, we will use a control function approach, as described in Wooldridge (2015) and others. As in 2SLS, we first estimate a selection function, using observables and an appropriate instrument, z. The residuals from the selection function are then be used as an estimate of the unobserved error component, λ_t.

Formally, we'll estimate the probability of treatment under the assumption that η is distributed standard normal and that there is a linear relationship between λ and η. The estimating equation of the first stage is:

$$Pr(T_t = 1) = \Phi(\delta_1 x_t + \delta_2 x_t q_t + \delta_3 z_t) \tag{2.27}$$

Where Φ is the cumulative distribution function of the standard normal distribution. It follows that

$$E[m_{t+1}|X_t, T_t] = \theta_1 x_t + \theta_2 x_t q_t + \theta_3 x_t T_t + E[\psi_t + \epsilon_t|X_t, T_t] \tag{2.28}$$

$$= \theta_1 x_t + \theta_2 x_t q_t + \theta_3 x_t T_t + \frac{x_t q_t}{\beta_2} E[\lambda_t|X_t, T_t] \tag{2.29}$$

$$\equiv \theta_1 x_t + \theta_2 x_t q_t + \theta_3 x_t T_t + \frac{\gamma x_t q_t}{\beta_2} r(X_t, T_t, z_t) \tag{2.30}$$

where γ captures the linear relationship between λ and η, and r is the control function:[4]

$$r(x_t, z_t, T_t) = T_t \left(\frac{\phi(\delta_1 x_t + \delta_2 x_t q_t + \delta_3 z_t)}{\Phi(\delta_1 x_t + \delta_2 x_t q_t + \delta_3 z_t)} \right) - (1 - T_t) \left(\frac{\phi(\delta_1 x_t + \delta_2 x_t q_t + \delta_3 z_t)}{1 - \Phi(\delta_1 x_t + \delta_2 x_t q_t + \delta_3 z_t)} \right) \tag{2.31}$$

The estimating equation for the second stage uses an estimate of r obtained from the first stage:

$$m_{t+1} = \theta_1 x_t + \theta_2 x_t q_t + \theta_3 x_t T_t + \theta_4 x_t q_t \hat{r_t} + \epsilon_t \tag{2.32}$$

where $\hat{r}$ is the predicted value of r. This will lead us to unbiased estimates of the treatment effect:

$$\Delta(X) = E[m_{t+1}|X_t, T_t = 1, z_t] - E[m_{t+1}|X_t, T_t = 0, z_t] \tag{2.33}$$

$$= \theta_3 x_t \tag{2.34}$$

Clearly, these results depend upon the parametric assumptions over the distribution of the unobservables. We can explore the tradeoffs associated with this choice later in the simulation.

2.4. Monte Carlo Experiments

Monte Carlo experiments illustrate the efficacy of the proposed estimator. To simulate the compartmental model, we'll make the parametric assumption that λ, η, and z are jointly distributed multivariate normal:

$$\begin{bmatrix} \lambda \\ \eta \\ z \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}) \tag{2.35}$$

where $\mathbf{\Sigma}$ is the variance covariance matrix, with elements denoted σ_{ij}. Given that z is designed as an instrumental variable, we let $\sigma_{\lambda z} = 0$ and $\sigma_{\eta z} > 0$. To illustrate the utility of the estimator, we'll compare three different relationships between λ, η, and T.

Table 2.1 contains the parameters for two different data generating processes, where the differences are driven solely by the covariance between the two sources of unobserved stochasticity, λ and η. The first scenario simulates random allocation of treatment, where treatment behavior is

[4]Under the assumption that λ and η are distributed multivariate normal, it can be shown that $\gamma = \sigma_{\lambda\eta}/\sigma_\eta$, given that $E[\lambda_t|X_t, \eta_t > \zeta - m_t] = \mu_\lambda + \frac{\sigma_{\lambda\eta}}{\sigma_\eta} \frac{\phi(\alpha)}{1-\Phi(\alpha)}$

TABLE 2.1. Simulation Parameters

Parameter	Random Treatment	Selection into Treatment
β_1	7.50×10^{-1}	7.50×10^{-1}
β_2	1.00×10^{-2}	1.00×10^{-2}
β_3	5.00×10^{-2}	5.00×10^{-2}
β_4	4.00×10^{-1}	4.00×10^{-1}
β_5	5.00×10^{-2}	5.00×10^{-2}
p	5.00×10^{-1}	
ν	1.50×10^{-2}	1.50×10^{-2}
ζ	1.00×10^{-3}	1.00×10^{-3}
σ_λ	5.00×10^{-2}	5.00×10^{-2}
σ_η	7.50×10^{-3}	7.50×10^{-3}
σ_z	3.75×10^{-1}	3.75×10^{-1}
$\rho_{\lambda\eta}$	0.00	8.00×10^{-1}
$\rho_{\eta z}$	0.00	1.00×10^{-1}
$\rho_{\lambda z}$	0.00	0.00
T	75.00	75.00

completely independent of any observable or unobservable features in the model. We set the covariance between λ and η equal to zero and the probability of treatment equal to a constant, $p \in [0, 1]$. In the second scenario, we simulate treatment determined both by the observed mortality rate and a shock that is correlated with the infection dynamics. As such, the the covariance between λ and η is greater than zero.

To give a sense of the dynamics of the system, the average outcomes are displayed in Figure 2.1. It important to note that the two scenarios illustrate qualitatively different dynamics. This is solely driven by our (behavioral) assumptions over the probability of treatment. Under selection, the probability of treatment varies over time, consistent with the behavioral model. By contrast, the probability of treatment is, on average, constant across the production cycle. Yet the probability of treatment matters. When the probability of treatment is equal to 0.5, the ending state of the system is nearly identical to the ending state under selection.

Next, we will conduct the Monte Carlo experiment to illustrate the distribution of the estimates for the effectiveness of treatment, $\hat{\beta}_4$. For each treatment scenario (treatment randomly allocated and treatment governed by selection), we will estimate equation 2.23 using ordinary least squares and a control function. The result of these experiments is illustrated in Figure 2.3. When treatment is randomly allocated, we see that both the control function and ordinary least squares (OLS)

FIGURE 2.1. Base Model Simulation

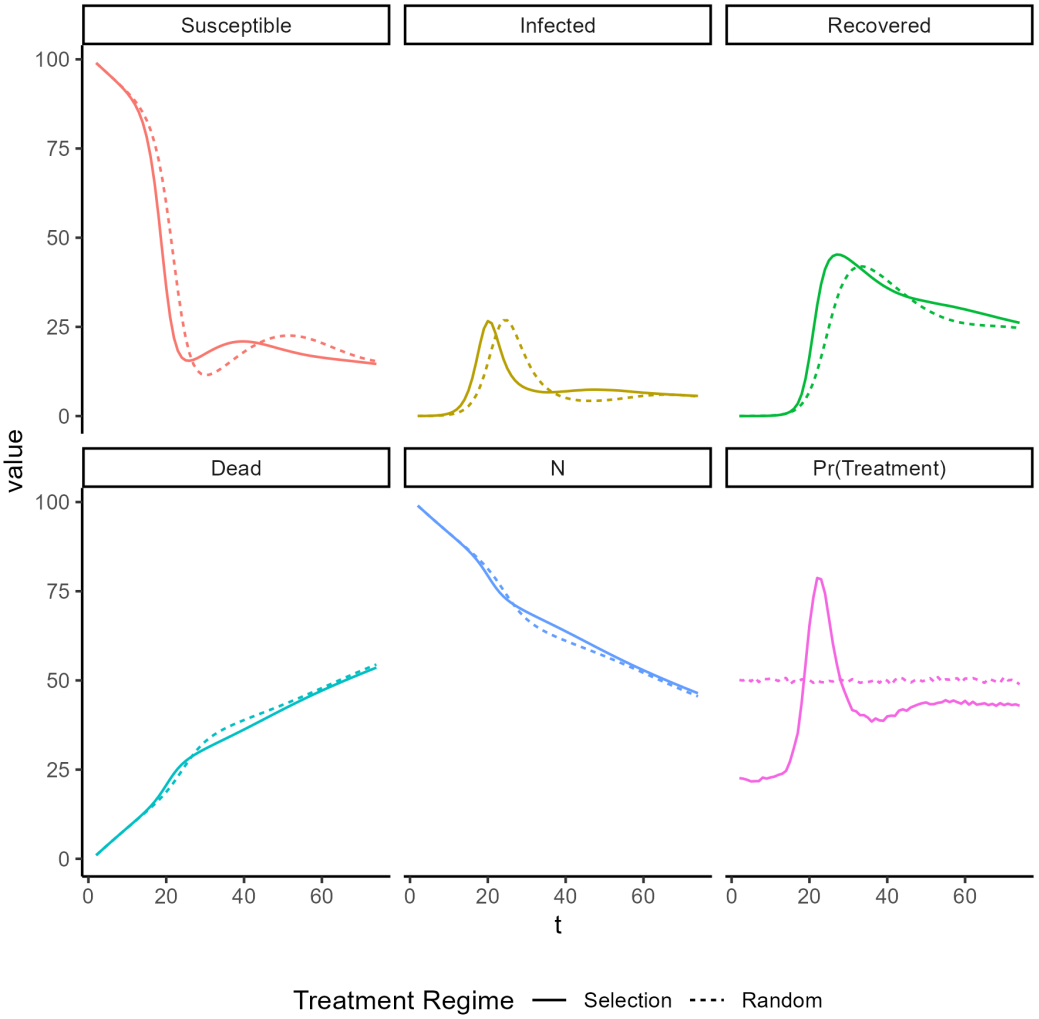

Note: This figure illustrates the mean outcomes across 10,000 simulations. The solid line represents the trajectories when treatment is random and the dashed line represents the trajectories when treatment is a function of the mortality rate.

recover unbiased estimates of θ_3 (which correspond to the structural parameter β_4). When selection is driven by a combination of observed and unobserved factors, OLS no longer returns an unbiased estimate. Instead, the estimate is biased towards zero, consistent with the presence of an unobserved feature that is positively correlated with the mortality rate. However, the control function does appear to generate an unbiased estimate but with less precision than OLS under random treatment assignment.

FIGURE 2.2. Distribution of $\hat{\theta_3}$

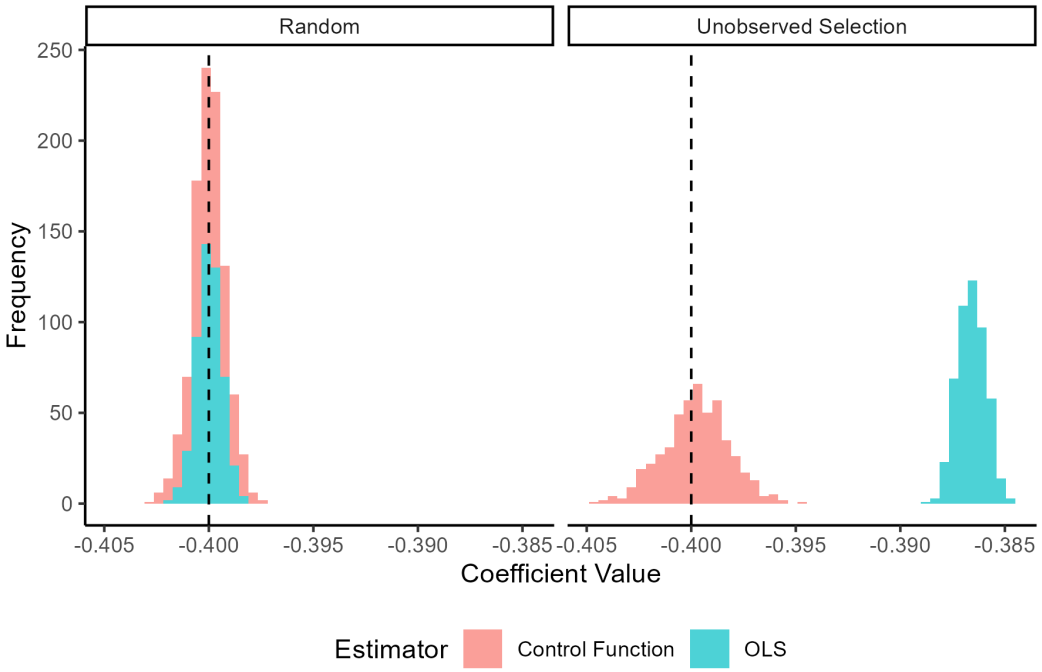

Note: This figure illustrates the distribution of the estimates for $\hat{\theta_3}$. The Monte Carlo experiment consisted of 500 draws, 127 simulated production cycles per draw each of length T. The dotted line represents the true value of the parameter.

In addition to estimating the effectiveness of treatment, we also report the distribution of the F-statistic associated with the significance of θ_4 in the second stage. If we reject the test statistic, then we have evidence that producers are non-randomly selecting into treatment. As we have illustrated, this generates a biased estimate of θ_3. In Figure 2.3, we see that this test identifies the presence of selection with relatively good accuracy. When selection is randomly allocated, the test statistic incorrectly rejected the null hypothesis 19% of the time. When producers are actively allowed to select into treatment, due to a combination of unobserved and observed characteristics, the significance test correctly rejected the null hypothesis in every instance.

2.5. Empirical Application to Antibiotic Use in Salmon Aquaculture

In this section, I describe how this estimator might be deployed to estimate the effectiveness of antibiotic use with observational data. As noted in the previous chapter, disease is one of the most important production risks facing Chilean salmon producers. Although the salmon are restricted

FIGURE 2.3. Statistical Test for Selection into Treatment

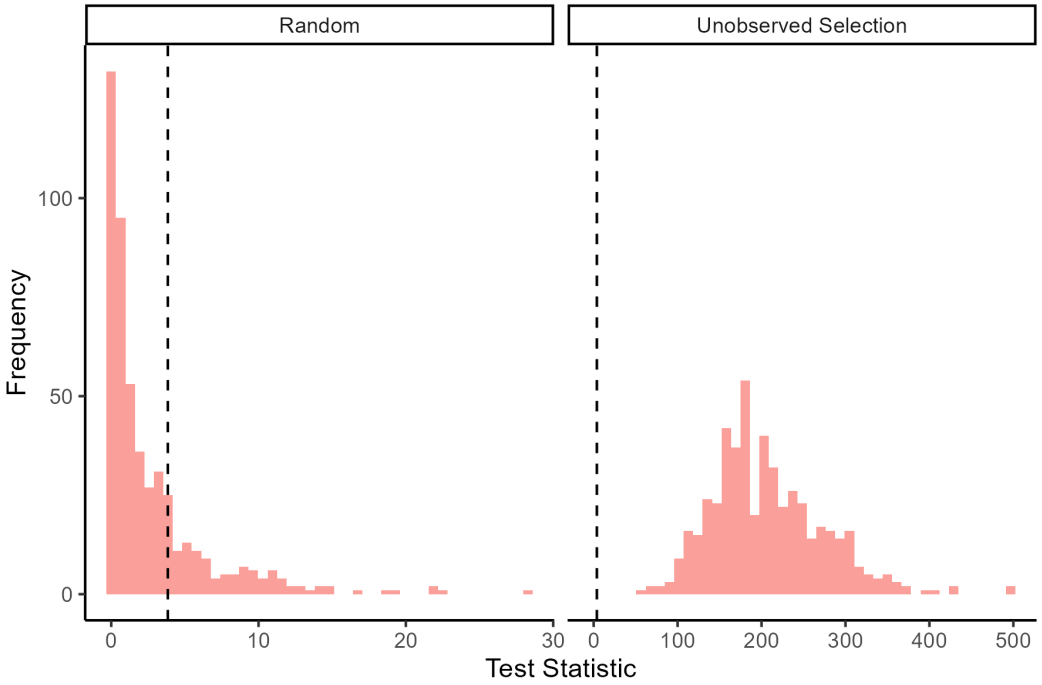

Note: This figure illustrates the distribution F-Statistic evaluating the significance of $\hat{\theta_4}$. The dotted line represents the critical value, assuming $\alpha = 0.05$. When the test statistic exceeds the critical value, it is statistically significant at a $1 - \alpha$ level of confidence.

to the pen, water and organic matter can move freely into and out of the pen, creating an "open" production system (Hennessy and Marsh, 2021). Farm managers can feed their fish, they cannot prevent exposure to pathogens or control the ambient environment conditions. In Chile, the most important pathogen is *P. salmonis*, the endemic, causative agent for Salmonid Rickettsial Syndrome (SRS). SRS infections account for the vast majority of all disease-related production losses in Chile's industry (SERNAPESCA, 2022). While vaccines for SRS are commercially available, they only provide protection during the first two to three months of the marine phase (Maisey et al., 2017). After this period, farm managers address SRS infections through antibiotics, typically administered via medicated feed. SRS accounts for 90% of all antibiotics used in Chile's aquaculture industry (SERNAPESCA, 2022) and has become central to growing concern over antimicrobial resistance.

The Chilean federal government regulates the use of antibiotics, and forbids their application before the arrival of the infection. In contrast to many other livestock systems, producers cannot

use antibiotics to prevent the arrival of pathogens nor can they be prescribed for non-medical applications such as growth-enhancement. Antibiotics must be prescribed by a veterinarian following a clinical diagnosis. The prescription and associated diagnosis must then be reported to the federal government.

However, the large scale of the farms prevents veterinarians from individually managing fish on the farm. Veterinarians do not track the health status of individual fish nor can they selectively apply antibiotics to infected individuals. When disease is detected on a farm, veterinarians and farm managers must make a judgement of when to treat, based upon their assessment of disease spread. SRS can be readily detected using polymerase chain reaction (PCR) tests (Rozas and Enríquez, 2014). But, given the scale of the farms and the stress associated with handling, it is not practical for veterinarians to test the entire farm and count the number of infected fish. Instead, they roughly estimate infection dynamics based upon a combination of PCR tests, mortality rates, and observable clinical signs. Infected fish tend to be lethargic, exhibit erratic swimming behavior, and, in some cases, develop haemorraghic ulcers (Almendras and Funealba, 1997). Once veterinarians choose to apply treatment, medicated feed is delivered to both infected and uninfected individuals alike.

Despite their widespread use in the industry, recent evidence suggests that antibiotic treatments are unreliable. One retrospective study of 118 farms concluded that antibiotics failed to control outbreaks in 47% of attempted treatments (Price et al., 2016). Some believe that the high treatment failure rates indicate growing resistance to antibiotic treatments (Cartes et al., 2017). Others suggest that in-feed antibiotic treatments will have low success rates if the infection reduces fish metabolism (Price et al., 2016). Whatever the cause, treatment failure is a major issue impacting the industry.

2.5.1. Identifying an Instrument. Based upon the evidence from the veterinary literature, farm managers likely select into treatment in response to a wide range of factors. Some of these features, such as the mortality rate, may be observable to the econometrician, while others, such as the behavior of fish in the pen, may may only be observable to the farm manager. As illustrated in the previous section, this will make it challenging to recover an unbiased estimate of the effectiveness of treatment. We see additional evidence of this in Panel A of Figure 2.4, which plots the

binned average probability of treatment. We see that the probability of treatment increases as the SRS mortality rate increases, suggesting that producers are actively selecting into treatment when disease incidence is high.

FIGURE 2.4. Probability of Antibiotic Treatment

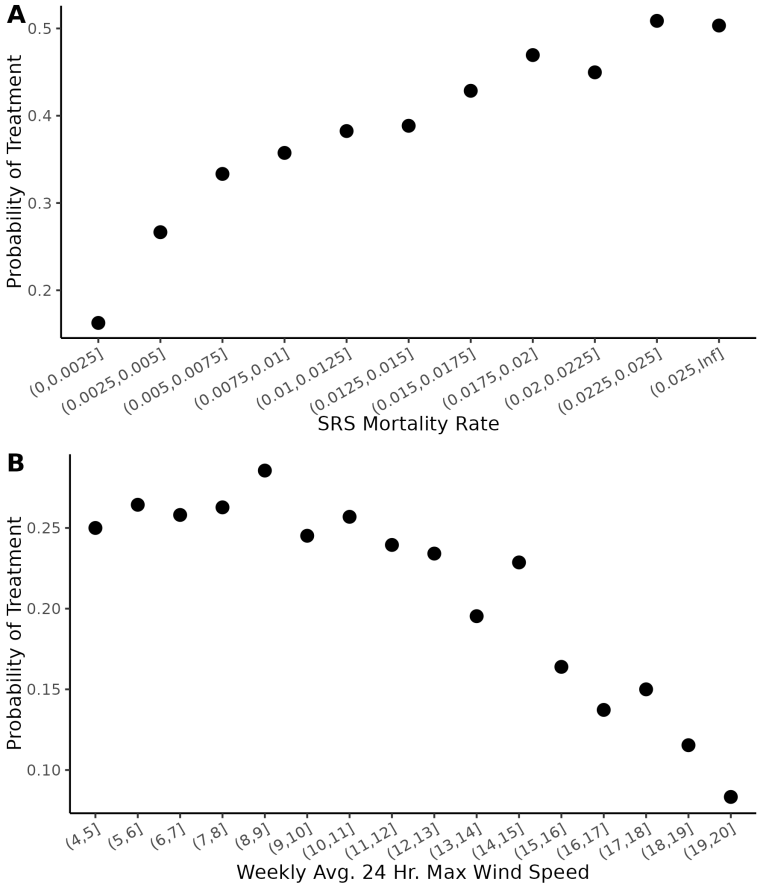

Note: Panel A depicts the average probability of antibiotic treatment by SRS mortality rate experienced on the farm. Panel B depicts the average probability of treatment as a function of the weekly average 24 hour maximum wind velocity.

In order to identify the effect of treatment, we need an instrument that is correlated with treatment behavior but which is plausibly uncorrelated with unobserved selection process. We propose using wind speed at four ports in southern Chile to predict variation in treatment that is plausibly exogenous from infection dynamics. As described in the previous section, farms do not tend to store medicated feed locally; it must be delivered when needed. These deliveries are vulnerable to weather-related shocks. To construct an instrument for wind speed we use information

collected from weather stations maintained by Servicios Climáticos of the Dirección Meteorológica de Chile. The weather stations are located in or near four ports that serve the salmon industry in Regions X and XI: Puerto Montt, Dalcahue, Puerto Aysen, and Melinka.

When wind speeds are very high, navigation becomes hazardous. During storms, there are times in which farm managers would like to treat their fish but cannot, simply because they are not able to deliver the medicated feed to the farm sites. To predict these delays, I calculate a 24 hour rolling average wind speed at each farm site. For each farm site-week, I choose the largest observed rolling average value. A large value indicates that the port has experienced windy, and possibly hazardous, conditions during the prior 24 hours, reducing the likelihood of medicated feed delivery. Given that we cannot observe any information about the ports which service individual farm sites in any given week, I estimate a distance-weighted average wind speed at each farm, where distant ports are weighted less than closer ports. Panel B of Figure 2.4 illustrates that the probability of treatment falls with higher average maximum wind speed experienced.

The key assumption underpinning this instrument is that wind speed does not impact future mortalities, except through its contemporaneous effect on the application of treatment. This is plausible for many reasons. First, net pens extend as much as 30 meters below the water, offering space in the pen that is largely unaffected by surface conditions. Second, while farm sites may not stockpile medicated feed, they do often maintain semi-permanent structures as well as technology to remotely monitor the farm site (e.g. remotely controlled cameras located within each pen). This means that poor weather conditions will not prevent farm managers from observing and responding to changes in the farm's health status. Furthermore, companies have, in many cases, chosen farm sites in protected bays that offer natural protection from adverse weather conditions and prevailing winds. Third, the instrument captures the weather conditions at the ports, rather than at the farm site. Many farm sites are located many kilometers away from any port; while we might expect weather conditions to be correlated across space, the observed conditions at the port may be quite different than those at the farm site. Even if the weather conditions did impact mortality rates in some other way, weather conditions at the ports may not be a strong predictor of weather conditions at the farm sites.

2.5.2. Data. We use a combination of data sources in this section. We utilize hourly observations of wind velocity from weather stations maintained by the Servicios Climáticos of the Dirección Meteorológica de Chile at four ports that serve the salmon industry in Regions X and XI, including Puerto Montt, Dalcahue, Puerto Aysen, and Melinka. These hourly data are aggregated to the weekly level, using procedure described above. We then match the weather data to the weekly production and antibiotic-use data for Atlantic salmon deployed in Chapter 1. We only use the production cycles which can be matched to non-missing weather data for all observed cycle weeks. This significantly restricts the sample from the available database to specific cycles in 2015-2019.

2.5.3. Empirical Results. The model estimates are reported in Table 2.2. The key term of interest, x $\times$ treat, is not statistically significant, except in the control function with a probit first stage (CF(c)). However, the coefficient on the un-interacted term, x, are consistently significant and close to one, regardless of the model. Although this is not explicitly an autoregressive model, x is a function of the lagged dependent variable. The statistical significant on this parameter may reflect a persistent autoregressive structure that explains much of the variation in mortality rates. It is interesting to note that the inclusion of additional controls appears to have little impact on the model estimates, especially since the additional controls include a cycle fixed effect. Given that the OLS estimates are not impacted by the inclusion of these terms, the model terms are not correlated with these additional controls This implies that there is not significant bias generated by the interaction between a lagged dependent variable a fixed effect.

The model results suggest that the estimates of treatment effectiveness using ordinary least squares (OLS) are smaller and closer to zero than those estimated using a control function. This matches the predictions from the analytical section of the analysis and would be consistent with selection into treatment on unobservables that are correlated with mortality rates. However, the control function does not offer strong evidence for selection. None of the F-statistics for the endogeneity tests are statistically significant. This holds for regardless of the functional form for the first stage and across inclusion of different controls.

The results of the first stage suggest a key weakness in the proposed instrumental variable. When additional controls are added into the model, the test statistic on the F-test drops to nearly zero and the coefficient on x $\times$ treat moves moves closer to zero. This is consistent with an

TABLE 2.2. Model Estimates

	OLS(a)	OLS(b)	CF(a)	CF(b)	CF(c)
x	0.849***	0.861**	0.943***	1.017***	1.274***
	(0.179)	(0.276)	(0.197)	(0.195)	(0.198)
x × q	-33.812**	-34.313	-24.624	-25.842	-34.316*
	(10.652)	(25.485)	(17.491)	(17.145)	(17.365)
x × treat	-0.107	-0.104	-0.155	-0.200	-0.562*
	(0.243)	(0.336)	(0.253)	(0.248)	(0.253)
x × q × $\hat{r}$			54.665	56.541	35.123+
			(37.231)	(37.589)	(20.444)
Num.Obs.	10361	10361	10361	10361	10361
Controls	X		X		
First Stage			Linear	Linear	Probit
First Stage (F)			0.00	18.94	27.15
EEV Test (F)			2.16	2.26	2.95

Note: This table reports the estimates from five different models, where each column represents a different model. The first two columns report ordinary least squares estimates and the last three columns report control function estimates. Bootstrapped standard errors are reported in parentheses below the parameter estimates. Models with controls include cycle, month, year, salinity, and temperature as additional regressors. The first stage row reports the functional form of the first stage of the control function. First Stage (F) reports the F-Statistic from a test of instrument strength. EEV Test (F) reports the F-statistic testing the endogeneity of treatment.

instrument that is correlated with the additional covariates added into the model. In this case, it is likely that that the measures of maximum wind speed exposure vary by season. Therefore, the inclusion of weather, salinity, and monthly controls may explain much of the variation in the instrument.

Despite the econometric challenges listed above, it is interesting to note that the magnitude of the estimates for treatment effectiveness using the control function are much larger than those estimated using OLS. Taken at face value, the control function estimate of β_4 (denoted x × treat in Table 2.2) is at least 50% larger than the OLS estimate. This difference is certainly economically significant, and suggests the possibility of significant bias.

2.6. Discussion

The purpose of this chapter is to illustrate that a policy relevant parameter with epidemiological meaning can be recovered from observational data. We illustrate the efficacy of a control-function based estimator using a Monte Carlo simulation. We bring the estimator to data from Chile's

salmon aquaculture industry and offer a first step towards how it might be deployed. Our approach leverages a model that captures both the epidemiological dynamics and the human response to disease risk. In doing so, we were able to point to a specific mechanism driving selection into treatment and predict its likely impact on the epidemiological dynamics as well as its econometric implications.

In the Monte Carlo simulations, the probability of treatment under selection is determined by a fairly simple model. In general, it assumes that the probability of treatment increases with the observed mortality rate. On the surface, this is a somewhat myopic view of disease control because it does not allow for decision-makers to pursue a strategy that maximizes the benefits of treatment over the entire production cycle. This impacts the states in which treatment occurs. In this model, the deterministic components of the selection function ensure that treatment tends to occur when mortality rates are high. In a dynamic behavioral setting, however, treatment might occur at very low levels of mortality in order to avoid the future, non-linear increases in costs. Similarly, a farm manager might avoid treatment during certain periods of high mortality (e.g. at the end of the production cycle) simply because the dynamic benefits of treatment cannot be captured.

This behavior has econometric implications. The analytical model suggests that selection bias will increase as the correlation between the unobserved shock to the infection rate and treatment probability strengthens. If the decision-maker is allowed to incorporate the dynamic benefits of treatment, then the correlation between contemporaneous, unobserved shocks to the infection dynamics and the probability of treatment might be weaker than under a myopic decision process. However, the current estimator and endogeneity test cannot distinguish between static and dynamic decision-making.

Although the underlying economic behavior may be the cause for our statistically insignificant result, it is more likely the result of a weak instrument. The instrument proposed in this chapter, a measure of average maximum wind speed in a given week proved poor, despite negative correlation between wind speed and treatment probability. This critically undermines the internal validity of the empirical estimates. As noted in the main body of the paper, the strength of the instrument appeared to decline dramatically once additional exogenous regressors were added into the model. It seems plausible that seasonal variation in disease incidence may be correlated with seasonal

patters in wind speed. As a result, the instrument may not be adding much additional variation into the estimator. However, these estimates might be improved with additional environmental variables. We justified the instrument via a particular mechanism: that wind speed at nearby ports reduced the probability of medicated feed deliveries due hazardous navigating conditions. But alternative measures might be more informative (for example, a similar argument could be made regarding wave height).

As a final point, there are at least two additional weaknesses in the epidemiological model. The additive shock generates a convenient, linear, functional form and offers a reasonable approximation of a stochastic epidemiological system. However, this structure is stable only for a relatively narrow set of parameters (Diekmann et al., 2010) and limits the extent to which we can explore the extremes of the parameter space. Second, although we explore the impacts of fixed effects in the empirical section of the paper and hypothesize on the interaction between the fixed effects and the (approximately) autoregressive structure in the estimating model, we do not incorporate this into the Monte Carlo experiment. Both of these features add important robustness to this approach.

2.7. Conclusion

In this chapter, we used an epidemiological model of a disease with endogenous treatment to illustrate how the effectiveness of antibiotic treatments might be recovered using observational data. Treatment effectiveness is a critically important parameter for understanding the private incentives for antibiotic use and for simulating counterfactual polices that might impact disease dynamics. We show that estimates of treatment effectiveness, including those informed a conventional epidemiological model are likely to underestimate the impact of treatment on disease dynamics. We offer an alternative approach using a control function to solve this problem and recover a parameter with a structural interpretation. Although our empirical application was limited by a weak instrumental variable, we believe that the approach offers a strong road map for future work.

Treatment estimates that are biased towards zero create a problem for policy simulation. If selection bias forces the estimate of treatment effectiveness towards zero, then the model will also underestimate the marginal benefits of treatment. If treatment is applied in a myopic fashion, the implications of the bias should be straightforward to predict and will simply depend upon the shape

of the instantaneous cost and benefit functions. But if disease management decisions are dynamic (i.e. they consider both current and future costs and benefits), then the impacts of the bias may have unexpected impacts on simulated, optimal, treatment behavior. We return to this point in Chapter 3.

CHAPTER 3

The Cost-Effectiveness of Antibiotic Reduction Policies

3.1. Introduction

Drug-resistant infections are becoming increasingly common in both humans and animals, prompting serious concern from the public health and veterinary communities (Laxminarayan et al., 2013; Van Boeckel et al., 2015, 2017). A growing body of evidence links these infections to the consistent use of antimicrobial products. But reducing the use of antimicrobials is challenging. Even as the evidence of resistance grows, antimicrobial products remain indispensable for fighting infection. In places with inadequate access to antibiotics, hundreds of thousands of people die each year from preventable conditions such as a pneumonia and sepsis (Laxminarayan et al., 2016). A similar story exists in agriculture, where pests and pathogens threaten the productivity of the global food system. When producers fail to control communicable diseases, they create an externality for their neighbors. In the absence of alternative disease-fighting technologies, antibiotics not only reduce private losses from disease but they also help to avoid transmission by reducing the number of infected individuals in the system.

Regardless of the application, policymakers interested in reducing the use of antibiotics must reconcile the long-term benefits of avoided resistance with the immediate losses from disease. In practice, however, realistic appraisals of resistance dynamics are rare, making it challenging to identify the efficient level of antibiotic-use. This has prompted policymakers and stakeholders to suggest second-best targets for antibiotic reduction that are intended to promote meaningful improvements.

In this paper, we explore how the design of antibiotic-use restrictions are likely to impact the productivity of a farm that is exposed to a communicable disease. Drawing on the growing literature around economic epidemiology, we model the farm as a forward-looking, profit maximizing entity where the biomass at harvest is determined by a compartmental model of disease dynamics. Farm

managers control the disease by applying chemical therapeutics that allow individuals to recover from the infection. We assume that pathogens are transmissable both between individuals on the farm and between farms in a region. We solve the model as a finite-horizon, deterministic optimal control problem and simulate a selection of counterfactual antibiotic-use restrictions that are parameterized to achieve an exogenous policy target.

We use Chile's salmon aquaculture industry as a motivating example. Chile is the world's second largest producer of salmon but uses antibiotics at a rate that far exceeds most other salmon producing countries. The vast majority of antibiotics are applied to treat *Salmonid Rickettsial Syndrome* (SRS), an endemic disease caused by a bacterial infection. Both policymakers and industry participants have publicly expressed interest in reducing antibiotic use, and have recently begun to make improvements towards this goal. The Chilean experience is particularly interesting because the country has also implemented a spatially and temporally explicit program to mitigate the spread of disease between farms. As part of our analysis, we illustrate how this pre-existing program might interact with antibiotic-use restrictions by creating a predictable trajectory of infection pressure experienced by the farms in a region.

The role of public policy in the control of pathogens and diseases is well appreciated within agricultural economics (Hennessy and Marsh, 2021). This paper draws inspiration from the bioeconomic literature in which economic models are combined with models of dynamic biological processes. Our approach bears similarity to both Bicknell et al. (1999) and to Horan and Fenichel (2007). Both evaluate optimal management strategies for disease that is transmissable between wildlife and livestock by combining an epidemiological model of disease with an economic (i.e. public policy) objective. Similar to our setting, they also focus on a pathogen that persists in the ecosystem and for which eradication may not be practical. However, our focus is on policies to reduce antibiotic use, rather than on strategies to minimize the costs of disease. In doing so, we highlight the tension inherent in many policies seeking to balance veterinary and public health. Furthermore, we analyze incentives facing individual decision-makers in the face of targeted policy, rather than modelling the social planner's solution for the entire system.

This paper makes three contributions to the literature. First, it extends our understanding of the linkages between private decision-making and disease dynamics. Many have illustrated that

this feedback is consequential for public policy, both within the context of human disease (Auld, 2003; Dangerfield et al., 2022; Francis, 1997; Gersovitz and Hammer, 2005; Kremer, 1996) as well as in agricultural and livestock production (Atallah et al., 2017; Bicknell et al., 1999; Fuller et al., 2017; Horan and Fenichel, 2007).[1] We explore this same dynamic within the context of second-best policy. In our case, we consider how an existing policy designed to control the spread of marine pathogens and a set of simulated policies targeted at reducing antimicrobial use.

Second, it explores the capital-theoretic nature of farmed fish though a unique lens. If fish are valuable assets that grow over time, then the farmer's problem is to maximize the net present value of harvest (Bjørndal, 1988). Mirroring similar advancements in the forestry literature, the foundational theoretical models for aquaculture have been extended to include additional dimensions of the production process (Guttormsen, 2008; Heaps, 1995; Mistiaen and Strand, 1998; Yu and Leung, 2006). However, despite the fact that infectious disease remains one of the most important production risks facing the global aquaculture industry (Asche et al., 2009; Shinn et al., 2018), few in the literature have rigorously incorporated infectious disease into conventional models. Abolofia (2014) suggests that a problem of this type could be solved and proposes a similar optimization problem in the context of a marine parasite; but the author does not pursue a numerical solution. As a result, much of our knowledge about the management of marine disease on fish farms comes from ecology and veterinary medicine (Frazer et al., 2012; Krkošek et al., 2010; Peacock et al., 2016; Revie et al., 2005), where it is common to assert restrictive or myopic assumptions over producer behavior. We will this gap by combining the compartmental epidemiological model with a dynamic model of firm decision-making.

Finally, this work contributes to the nascent literature on the economics of antibiotic use. Despite increasing awareness of the costs of antibiotic resistance in both humans and livestock, the design and implementation of antibiotic-use policies remains elusive (Laxminarayan et al., 2013; Van Boeckel et al., 2017). Traditionally, research on antibiotic use has focused on resistance, where a drug's effectiveness can be modelled as a non-renewable resource (Laxminarayan and Brown, 2001). We take a different approach, focusing instead on antibiotics as a production input targeted at fighting disease. We address resistance only indirectly, through the simulation of policies which

[1]Recent experience with COVID-19, in particular, has generated a surge in interest in this idea for public health applications. McAdams (2021) provides a preliminary review of this literature.

target a specific reduction in antimicrobial use. In doing so, we illustrate that the non-linearities in the disease dynamics are ultimately consequential for the behavioral response to both disease risk and to antibiotic use policy.

Our model illustrates that the private optimal utilization of antibiotics is a function of the net shadow value of recovered fish. However, this shadow value is a function of ambient disease pressure. In some cases, when ambient disease pressure is high (increasing the rate of re-infection after treatment), the shadow value of antibiotic treatment falls. This implies that a reduction in ambient disease pressure may not immediately imply a reduction in antibiotic use. We also show that a cap on the volume of antibiotic use offers the most cost-effective approach to reaching the target levels of abatement. This aligns with long-standing theoretical and empirical conclusions from the environmental economics literature (Baumol and Oates, 1988). Furthermore, our simulations illustrate the net present value of the fish farm falls as a result of the restriction, but at a slower rate than the reduction in antibiotic use. This is because the firm can respond to restriction across the entire production cycle, and because the restriction alters dynamics of pathogen dispersal between farms. We also illustrate that a cap on the number of weeks of antibiotic application is more cost-effective than a restriction on the maximum weight at which antibiotics can be applied. This is the result of dynamic incentives from disease spread on the farm as well as the altered trajectory of pathogen pressure between farms. However, the cap on the number of application weeks creates a far larger externality than the restriction on the maximum weight at treatment.

3.2. Model of Optimal Treatment

In this section we develop a model of a salmon farm over a single production cycle. The objective of the farm manager is to maximize profits, π, generated by the sale of uninfected biomass remaining at the end of the production cycle. We assume that the fish represent a single age class, are stocked instantaneously at $t = 0$, and do not reproduce during the production cycle.[2] We also assume that the farm is exposed to an infectious disease that is present in the ambient environment. The model contains two main components: an epidemiological model of disease and an economic model of treatment decisions.

[2]In other words, the number of fish on a farm strictly decreases during the cycle.

3.2.1. Epidemiological Submodel. The epidemiology of the disease is captured using a conventional compartmental model that is adapted to our setting. Let S represent the number of individuals that are susceptible to infection, I represent the number of individuals that are infected by the disease, and R represent the number of individuals that are recovered. Fish experience mortality in the S and I states at a rate of β_2 and β_5 per week, respectively, where $\beta_5 > \beta_2$. These individuals are permanently removed from the system. Ambient pathogen levels in the water, F, represent the average pathogen concentration in the water, measured as cells/m^3. Pathogen levels naturally dissipate at a rate of β_7 per week but increase due to shedding from the farm at a rate of β_9 cells/($m^3 \times$ individual $\times$ week) and from exogenous factors at a rate of β_8 per week.

FIGURE 3.1. Compartmental Model of Farm-Level Disease Transmission

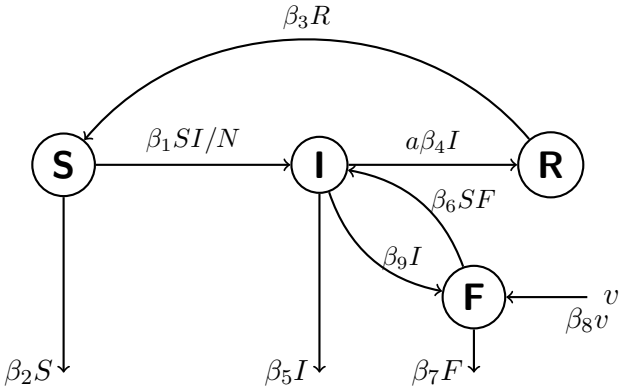

Note: This figures illustrates the dynamics captured in the epidemiological submodel. The circles represent state variables and the arrows represent flows between states. Arrows which leave a node but are not connected to another represent mortality rates. Arrows do not start at a node but end at a node represent an exogenous factor. The parameters that determine the rate of flow are listed beside each arrow. a is a control variable that is determined endogenously in an economic submodel.

Fish transition from susceptible to infected as a result of a contact with infected individuals at a rate of β_1 per week as well as outside sources at a rate of β_6 per week. Once infected, fish can only enter the recovered state once they have been treated, $0 \leq a \leq 1$. This occurs at a rate β_4/(treatment unit $\times$ week). Fish in the recovered stage transition become susceptible again at a rate of β_3 per week.

F represents an external reservoir of waterborne pathogens that continuously spread to the farm from the outside. This extension to the conventional $SIRS$ model has been used to model other water-born disease, such as cholera (Igoe et al., 2023; Tien and Earn, 2010). As the state dynamics illustrate, the ambient disease levels in F are jointly determined by infections on the farm, I, and by exogenous factors, v. The ambient pathogen levels are also subject to a natural decay rate, β_7. Modelling the disease in this way creates three main pathways through which disease is spread. First, disease can spread directly through contact between susceptible and infected fish. Second, disease can spread from infected fish to the local environment and back to susceptible fish on the farm. This allows for the persistence of the pathogen within the local environment, even if treatment occurs. Third, disease can spread from an exogenous source v into the local environment and subsequently to susceptible fish susceptible fish. When $v = 0$, the model is closed and the epidemiological dynamics would be completely determined by the conditions and actions on the farm. When $v > 0$, the system is partially determined by exogenous features of the system. We explore the implications of this in a later section. This is a rich representation of farm-level disease dynamics and is appropriate for our interest in simulating farm-level policies.[3]

3.2.2. Economic Submodel. The farm manager's problem is to identify the series of treatment decisions throughout the production cycle that will maximize net profits. Revenues are generated through the sale the fish that are remaining at the end of the period. The number of fish in any period, $N_t = S_t + I_t + R_t$, is determined by the epidemiological submodel. In each period, the farm manager chooses the level of antibiotic treatment, a_t, where $0 \leq a_t \leq 1$, where treatment moves infected individuals into a recovered state. The value of each fish is also determined by the average weight of each fish, W. We assume that W evolves according to a logistic growth function, with W_0 and W_T exogenous, and that there is no feedback between weight and the other states in the model. In doing so, the time path for weight is nonlinear but is, essentially, exogenous to the rest of the system. Furthermore, setting W_0 and W_T effectively fixes the length of the production cycle T (Appendix). However, the ending number of fish in each state, S_T, I_T, R_T, and the ambient

[3]This is not the only possible specification for this problem. One could, for example, model the set of pathogens across a collection of adjacent salmon farms as a metapopulation (Rowthorn et al., 2009). This undoubtedly would offer deeper insights regarding the impacts of control actions on one farm on the outcomes of another but may require a simpler representation of disease dynamics on the farm.

pathogen concentration, F_T, at the end of the production are freely determined by the optimal control problem.

This can be summarized using the following optimization problem:

$$\max_{a_t} \ \pi = e^{-\delta T} p W_T (S_T + R_T) - \int_0^T e^{-\delta t} C(a_t, N_t, W_t) dt \tag{3.1}$$

$$s.t. \quad \dot{S} = -\beta_1 \frac{S_t I_t}{N_t} - \beta_2 S_t + \beta_3 R_t - \beta_6 S_t F_t \tag{3.2}$$

$$\dot{I} = \beta_1 \frac{S_t I_t}{N_t} - \beta_4 I_t a_t - \beta_5 I_t + \beta_6 S_t F_t \tag{3.3}$$

$$\dot{R} = -\beta_3 R_t + \beta_4 I_t a_t \tag{3.4}$$

$$\dot{F} = -\beta_7 F + \beta_8 v_t + \beta_9 I_t \tag{3.5}$$

$$W_t = \frac{k W_0}{W_0 + (k - W_0) e^{-rt}} \tag{3.6}$$

$$0 \leq a_t \leq 1 \tag{3.7}$$

$$0 \leq S_t, \ 0 \leq I_t, \ 0 \leq R_t, \ 0 \leq F_t \tag{3.8}$$

$$S_0, I_0, R_0, W_0, W_T \text{ known} \tag{3.9}$$

As noted above, W follows a logistic growth function with solution in equation 3.6, where r is the intrinsic growth rate for fish weight and k is the maximum fish weight. c represents the marginal cost of treatment per unit of biomass and p is the harvest value of a unit of biomass. T and W_0, are left free, while W_T is fixed. In other words, we allow the farm manager to choose the optimal cycle length and starting weight of the fish but fix the weight of the fish at harvest.[4]

The first term in the objective function represents the discounted value of biomass at harvest. Mathematically, this acts as a salvage value for the program. For now, we will assume that all live fish, regardless of their infection status, have the same value at harvest. The second term captures accumulated treatment costs over the production cycle. For now, we leave the cost function undefined but assume that $\frac{\partial C}{\partial a}, \frac{\partial^2 C}{\partial a^2} > 0$ and that $\frac{\partial^2 C}{\partial a \partial N}, \frac{\partial^2 C}{\partial a \partial W} > 0$. In other words, we assume that costs are non-linear in the control and that control costs increase in biomass. The shape of the cost function is consequential. First, the non-linearity in the control variable avoids a numerically

[4]This would be consistent with a farm that is operating under a forward contract.

challenging linear optimal control problem. Second, by assuming that treatment costs increase in biomass, rather than the number of infected individuals, we capture the fact that farm managers apply treatment to a proportion of the farm, rather than to individual fish. This proportion is captured by a and is constrained between zero and one. This means that costs rise with biomass but not, necessarily, rise with the extent of the infection.

3.2.3. Parameterizing the Trajectory of v . The model above emphasizes the private benefits of treatment but does not comprehensively captured the external benefits of treatment generated from avoided transmission. Treatment determines the rate at which fish recover, reducing the number of infected individuals, I, and the ambient pathogen load, F. However, we assume that v is exogenous and take the following steps to operationalize this approach:

(1) Solve the optimization problem, assuming $F_0 = \bar{F} > 0$ and $v_t = \left(\frac{\beta_7}{\beta_8}\right) F_0$ for all t. This initializes the problem such that $\dot{F} = 0$.

(2) Obtain the optimal time path for I^*

(3) Solve the optimization problem again, assuming $F_0 = \bar{F}$ and $v_t = I_t^*$.

We are interested in the impacts of a antibiotic reduction policy (i.e. a mandatory reduction in disease-control efforts), given that a coordinated disease control strategy is already in place. As noted in previous chapters, the existing neighborhood strategy requires farms to follow coordinated production and fallow periods. The epidemiological model suggests that this will create a non-random, time-varying pattern of ambient disease prevalence, to which all farms in the area will be exposed. Data observed from the industry largely corroborates this point. The variable v captures the trajectory of this externality and the optimal disease control strategy will account for the full trajectory of v. But it does not reflect a Nash equilibrium.[5] This approach simplifies the spatial-dynamic features of the problem in order to focus on the within-farm disease dynamics, in the presence of a specific but realistic trajectory of ambient disease.

3.2.4. Ambient Disease and Eradication. The possibility of disease eradication from the farm and/or the ambient environment bears significant practical importance. If lasting eradication

[5]To be explicit, this is neither a open-loop Nash equilibrium or a feedback Nash equilbrium. It is not an open loop equilibrium because the trajectory for I that is used as an input for v is determined by a slightly different optimization problem and, therefore, does not reflect an equilibrium. This also does not reflect a feedback equilibrium because the actions of one farm do not impact the strategy of another.

were possible, then the series of actions to achieve this corner solution might be part of an optimal solution. But the epidemiological dynamics of this model largely preclude eradication.

The extent to which eradication is functionally possible depends upon the relationship between I and F in the epidemiological model. Eradication on the farm occurs when the number of infected fish is equal to zero ($I = 0$) while $N > 0$. These conditions hold at the start of the production cycle because we assume that none of the fish entering the pen are infected ($I_0 = 0$). I will remain at zero so long as $F = 0$. The state dynamics equation for F illustrates two channels through which F can become positive. First, the initial conditions can stipulate $F_0 > 0$. Assuming for the moment that there is no spillback from the farm into the ambient environment ($\beta_9 = 0$), then the dynamics of F are determined by disease decay, $\beta_7 F$ and the exogenous disease source $\beta_8 v$. The relationship between these two terms will determine whether F is increasing or decreasing ($\dot{F}$ positive or negative). But exponential decay ensures that $F > 0$, for any decay rate, so long as $v \geq 0$. If spillback from the farm into F is possible ($\beta_9 > 0$), then these dynamics will reinforce the effect of v by accelerating the increase in the ambient pathogen levels. The second possibility is that $F_0 = 0$ but that $v_t > 0$ at some $t < T$. This would capture the scenario in which there are no pathogens either on the farm or in the ambient environment without an exogenous shock. As before, F will remain positive once infected, so long as $v \geq 0$. Similarly, if $\beta_6, \beta_9 > 0$, then the farm will become infected and remain infected throughout the production cycle.

When $F, I > 0$ and $\beta_6 > 0$, then it is not possible to return to $I = 0$ and achieve eradication. This is largely driven by the way treatment a, enters into the epidemiological dynamics. First, the rate at which fish recover from treatment is a function of the number of infected individuals, $\beta_t I_t a_t$. This means that treatment cures a fixed proportion of fish, a decay process that may force I towards zero but will never achieve $I = 0$ in a finite horizon. Second, we do not interact a with F or S. Treatment does not prevent new infections, it only increases the recovery rate.[6] As a result of these two features, infection from F is unpreventable and permanent under our specification. By extension, eradication is essentially impossible. Therefore, we will not seriously consider this course of action as part of an optimal control strategy.

[6]Changing the structure of the model to prevent infection would not make eradication inevitable (or economically viable). If $F > 0$, then eradication would only be possible under certain specific assumptions over how a enters the the state dynamics for S and I.

3.2.5. Private Optimal Control Strategy. The optimal private solution to this problem can be described using the following current-value Lagrangian expression:[7]

$$\begin{aligned}\mathcal{L} = &- C(a_t, N_t, W_t) + \lambda_{1t}\dot{S} + \lambda_{2t}\dot{I} + \lambda_{3t}\dot{R} + \lambda_{4t}\dot{F} \\ &+ \underbrace{\theta_{1t}S_t + \theta_{2t}I_t + \theta_{3t}R_t + \theta_{4t}F_t}_{\text{State Constr. } -S,-I,-R,-F\leq 0} + \underbrace{\mu_{1t}(1-a_t) + \mu_{2t}a_t}_{\text{Control Constr. } a_t\leq 1,\, -a_t\leq 0}\end{aligned} \tag{3.10}$$

where λ, θ, and μ are Lagrange multipliers to address the state dynamics, non-negativity constraints, and the constraints on control, respectively. The Karush-Kuhn-Tucker conditions must hold such that, $\forall t \in [0, T]$:

$$\frac{\partial \mathcal{L}}{\partial a_t} = -\frac{\partial C}{\partial a_t} + \beta_4 I_t(\lambda_{3t} - \lambda_{2t}) - \mu_{1t} + \mu_{2t} \leq 0 \tag{3.11}$$

$$0 \leq a_t \leq 1, \qquad a_t \frac{\partial \mathcal{L}}{\partial a_t} = 0, \qquad (1 - a_t)\frac{\partial \mathcal{L}}{\partial a_t} = 0 \tag{3.12}$$

$$\mu_{it} \geq 0 \qquad \mu_{it}\frac{\partial \mathcal{L}}{\partial \mu_{it}} = 0, \qquad i = 1, 2 \tag{3.13}$$

$$\frac{\partial \mathcal{L}}{\partial \theta_{jt}} \geq 0 \qquad \theta_{jt} \geq 0 \qquad \theta_{jt}\frac{\partial \mathcal{L}}{\partial \theta_{jt}} = 0, \qquad j = 1, 2, 3, 4 \tag{3.14}$$

$$\begin{aligned}\dot{\lambda_1} = &- C(a_t, N_t, W_t) + \lambda_{1t}\left[\delta - \beta_1\left(\frac{I_t}{N_t} - \frac{S_t I_t}{N_t^2}\right) - \beta_2 - \beta_6 F_t\right] \\ &+ \lambda_{2t}\left[\beta_1\left(\frac{I_t}{N_t} - \frac{S_t I_t}{N_t^2}\right) + \beta_6 F_t\right] + \theta_1\end{aligned} \tag{3.15}$$

$$\begin{aligned}\dot{\lambda_2} = &- C(a_t, N_t, W_t) + \lambda_{1t}\left[-\beta_1\left(\frac{S_t}{N_t} - \frac{S_t I_t}{N_t^2}\right)\right] \\ &+ \lambda_{2t}\left[\delta + \beta_1\left(\frac{S_t}{N_t} - \frac{S_t I_t}{N_t^2}\right) - \beta_4 a_t - \beta_5\right] + \lambda_{3t}\beta_4 a_t + \lambda_{4t}\beta_9 + \theta_2\end{aligned} \tag{3.16}$$

[7] See Chiang (1992) for an extensive discussion regarding the relationship between the canonical Hamiltonian setup and the Lagrangean approach used here.

$$\dot{\lambda_3} = -C(a_t, N_t, W_t) + \lambda_{1t}\left[\beta_1 \frac{S_t I_t}{N_t^2} + \beta_3\right] + \lambda_2\left[-\beta_1 \frac{S_t I_t}{N_t^2}\right] + \lambda_{3t}(\delta - \beta_3) + \theta_{3t} \tag{3.17}$$

$$\dot{\lambda_4} = -\lambda_{1t}\beta_6 F_t + \lambda_{2t}\beta_6 F_t + \lambda_{4t}(\delta - \beta_7) + \theta_{4t} \tag{3.18}$$

$$0 = \left[\lambda_{1T} - e^{-\delta T} p W_T\right] S_T \tag{3.19}$$

$$0 = \lambda_{2T} I_T \tag{3.20}$$

$$0 = \left[\lambda_{3T} - e^{-\delta T} p W_T\right] R_T \tag{3.21}$$

$$0 = \lambda_{4T} F_T \tag{3.22}$$

By explicitly including non-negativity constraints, S, I, R, and F can include zero in their optimal time path. This allows for a realistic depiction of disease in this context. However, if all of the state variables are positive, then $\theta_i = 0$. Similarly, if $0 < a < 1$ and the integral constraint is satisfied, then $\mu_i = 0$.

Equations 3.11, 3.12, and 3.13 illustrate the optimal conditions for treatment. Equation 3.11 depends upon the shadow value of I and R, making an analytical solution impractical. To gain additional insight, consider an interior solution such that $0 < a_t < 1$. It follows immediately that $\mu_1 = \mu_2 = 0$, and $\frac{\partial \mathcal{L}}{\partial a_t} = 0$. Rearranging terms,

$$\frac{\partial C}{\partial a_t} = \beta_4 I_t(\lambda_{3t} - \lambda_{2t}) \tag{3.23}$$

For an interior solution, the marginal costs of treatment must be equal to the marginal benefits of the recovered fish. Treatment moves fish from an infected state to a recovered state at a rate of $\beta_4 I$. Although the farmers do not capture an instantaneous payoff from moving fish from an infected to a recovered state, there are dynamic benefits generated by the increase in recovered fish and the decrease in infected fish. Assuming that the shadow value of a recovered fish exceeds the shadow value of an infected fish, $\beta_4 I_t(\lambda_{3t} - \lambda_{2t}) > 0$ for $I > 0$. When the marginal costs exceed the shadow value of treated fish, $\frac{\partial C}{\partial a_t} > \beta_4 I_t(\lambda_{3t} - \lambda_{2t})$, treatment is no longer optimal. Given that treatment is restricted to non-negative values, $a = 0$ and $\mu_2 > 0$. Similarly, when the benefits of treatment are sufficiently large, the optimal value of a may exceed one. Given the constraint on a, the solution will be restricted to $a = 0$ and $\mu_1 > 0$. These results suggest that it may be optimal for

no treatment occurs for long periods of time ($a = 0$) as well as the possibility of maximal treatment ($a = 1$) to occur.

Equations 3.14 are the complementary slackness conditions associated with the non-negativity constraints on the state variables. Equations 3.15 - 3.18 describe the dynamics of the shadow values for each state variables. As the complex functional forms illustrate, they are dependent upon both economic and epidemiological components of the model.

When $t = T$, the fish are harvested but we assume that only the susceptible and recovered fish have value. The transversality conditions in 3.19-3.21 capture these dynamics. The shadow value of susceptible and recovered fish is equal to the present value of an individual fish: $e^{-\delta T} p W_T$. By contrast, infected fish have no value, so the shadow value of the remaining infected fish is equal to zero. This foreshadows an important incentive for producers: if the shadow value of infected fish approaches zero at the end of the production cycle, then the optimal strategy will be to move as many infected fish as possible into the recovered stage. This is clearly illustrated in 3.23, as this implies that the difference between the λ_3 and λ_2 will increase as t approaches T. Finally, given the finite-time horizon of the problem, the shadow value of additional infected cells in F is also equal zero and is captured by 3.22.

3.3. Policy Instruments to Abate Antibiotic Use

Given that the external costs of antibiotic applications are unknown, we cannot identify the efficient level of antibiotic application. However, we can evaluate the extent to which alternative policy designs represent cost-effective approaches to reducing antibiotics. To this end, suppose that the status quo volume (in kilograms of active product) of antibiotic utilized by a profit maximizing farm is

$$V^* = \int_0^T a_t^* W_t (S_t + I_t + R_t) dt \tag{3.24}$$

and policymakers set a maximum antibiotic volume standard of $\bar{V} = \phi V^*$, where $\phi \leq 1$. In theory, ϕ could be set to optimally balance the intertemporal trade off between the long-term development of antibiotic resistance and short-run losses due to disease. But, in practice, these dynamics are often unknown. We will assume that ϕ is an exogenous target set by a policy maker; given this

target, we will identify the cost-effectiveness of a suite of policy instruments that might be used to achieve the targeted level of abatement. Each policy described below will add a new set of restrictions to the optimization problem described in 3.1-3.9.

Cap on Total Antimicrobial Product. Perhaps the most direct policy is to simply place a cap on the total antimicrobial product applied during the production cycle at $\bar{V}$. Mathematically, this is operationalized by setting the following integral constraint:

$$\int_0^T a_t W_t (S_t + I_t + R_t) dt \leq \bar{V} \tag{3.25}$$

where $\bar{V}$ is defined, as above. This means that the firm is free to achieve the desired level of abatement by adjusting both the timing and the intensity of antibiotic applications. Conventional economic theory suggests that this will be the most cost-effective approach to achieving the antibiotic use target (Baumol and Oates, 1988).

Cap on Total Antimicrobial Periods. It may be challenging and costly to accurately measure the total volume of antimicrobial product utilized on a farm site. For example, different antibiotic products may have different doses and applied at different rates. This makes it challenging to calculate the total antibiotic load in the system. To this end, we consider a second, less precise measure of antibiotic application: the total farm weeks of treatment. Recall that, in our model, a represents the proportion of cages treated during a period: $a = 1$ represents 1 farm-week of treatment. Therefore, policy-makers could cap antibiotic treatments at A farm-weeks such that

$$\int_0^T a_t dt \leq \bar{A} \tag{3.26}$$

For purposes of comparison, we select a value of A such that the total volume of antibiotics meets the desired target.

Maximum Treatment Weight. As a final policy instrument, we will also simulate the impacts of a restriction on treatment applications at the end of the production cycle. More specifically, we will consider the impact of preventing treatment τ weeks from harvest. Mathematically this enters the farm's optimization problem as an integral constraint:

$$0 = \int_{T-\tau}^{T} a_t dt \tag{3.27}$$

Given that the growth rate of the fish is completely exogenous in this model, this is equivalent to identifying a maximum average weight at which treatment can be applied. Such a restriction appears intuitive. As fish get larger, the dose required to cure a sick fish increases, driving non-linear increases in the volume of antibiotics per treatment week, simply because the fish are getting larger. This restriction would avoid treatments when fish require the largest dose. This policy would also be relatively easy to enforce. Firms are required to report information on the size of their fish on a regular basis during the production cycle along with their use of antibiotics, making it very simple for a regulator to identify non-compliant firms.

It is worth noting that this policy is similar to existing restrictions on antibiotic applications that intended to avoid the sale of fish that are contaminated with antibiotic residues. Producers cannot sell fish which have been treated with antibiotics in the previous month, creating a *de facto* restriction on maximum weight at treatment. Although this restriction is intended to protect consumer safety, it is likely to have implications for the application of antibiotics.

3.4. Numerical Methods and Calibration

The optimization problems outlined in the previous section are complex.We approximate the optimal treatment regime using pseudospectral collocation (e.g. Garg et al. (2010)), a numerical method that is robust to corner solutions, as well as state and control constraints. It is a well-known approach that has proven efficient and effective in similar contexts (Fuller et al., 2017; Sanchirico and Springborn, 2011). Collocation approximates a continuous function by fitting a flexible polynomial at a discrete set of points (nodes). At each node, the algorithm minimizes the difference between the true function and the polynomial approximation (the residual). The parameters of the polynomial are calculated via a series of individual non-linear optimization problems solved at each node (Judd, 1998). We use as many as 90 Gauss-Legendre collocation points, where the finite number of nodes reflects the tradeoff between computational burden and numerical precision.

TABLE 3.1. Economic, Biological, and Epidemiological Parameters

Parameter	Units	Value
p	\$ kg^{-1}	6.33
c_1	\$ treatment^{-1}	2.50×10^{-1}
c_2	\$ treatment^{-1} fish^{-1} kg^{-1}	2.50×10^{-2}
c_3	\$ treatment^{-1} fish^{-1} kg^{-1}	1.00×10^{-2}
W_T	kg	5.50
W_0	kg	2.00×10^{-1}
r	week^{-1}	8.74×10^{-2}
T	weeks	66.00
k	kg	6.00
β_1	week^{-1}	2.00×10^{-1}
β_2	week^{-1}	4.00×10^{-4}
β_3	week^{-1}	8.75×10^{-2}
β_4	treatment^{-1} week^{-1}	6.00×10^{-1}
β_5	week^{-1}	1.24×10^{-2}
β_6	ml^3 cell^{-1} week^{-1}	1.00×10^{-3}
β_7	week^{-1}	1.00×10^{-2}
β_8	week^{-1}	1.00×10^{-3}
β_9	cells ml^{-3} week^{-1} fish^{-1}	1.00×10^{-3}
F_0	cells ml^{-3}	1.00×10^{-1}
δ		5.00×10^{-4}
ϕ		5.00×10^{-1}

Table 3.1 summarizes the economic and epidemiological parameters used in this study. Although the model is motivated by ongoing challenges associated with *Salmonid Rickettsial Syndrome* (SRS), the epidemiological dynamics of this disease have proven challenging to quantify. Recent research has pushed towards a precise measurement of *P. salmonis* shedding rates using studies conducted in the laboratory (Long and Jones, 2021). But it remains unclear whether these estimates are really transferable to farm conditions. With this in mind, we calibrate the model using convenient approximations. Although these estimates are rough, we believe that they are sufficient to illustrate the comparative performance of the policies. Additional calibration is necessary before the quantitative solutions can be used for planning purposes.

We use the following quadratic cost function:

$$C(a_t, W_t, N_t) = c_1 a_t + c_2 a_t W_t N_t + c_3 (a_t W_t N_t)^2 \tag{3.28}$$

with parameters chosen to illustrate convexity with respect to biomass. The biological parameters governing fish weight are directly from industry data. We use the median starting fish weight and median harvest weight of Atlantic salmon harvested in Chile between 2013 and 2020 as estimates for W_0 and W_T. Additional information regarding the calibration of the fish weight function can be found in Appendix B.1.

The targeted reduction in antibiotic applications, ϕ, is set to match with public commitments made by the salmon industry in recent years (White, 2024). In light of this, we explore the implications of abating antibiotic use (volume) by 50%, relative to the base model. When simulating the policy that caps the volume of antibiotic use, $\bar{V} = \phi V^*$ enters explicitly as a constraint in the optimization function. For policies which restrict the number of applications or establish a maximum fish weight at which treatments can be applied, we iteratively search for the value of $\bar{A}$ and τ that would achieve the abatement target. The specific parameters used to calibrate the policy simulations are reported in Appendix B.2.

3.5. Simulation Results

3.5.1. Base Scenario. We will begin by illustrating the optimal treatment strategy in the absence of an antibiotic abatement policy. When costs are quadratic, treatment at some level is nearly always profitable. Nevertheless, there are two distinct pulses of treatment: one near the beginning of the production cycle and the other towards the end of the production cycle. These two pulses are driven by different incentives. At the beginning of the production cycle, the marginal costs of treatment are relatively low because the total farm biomass is small. Furthermore, there are large dynamic benefits associated with reducing the number of infected individuals at the beginning of the production cycle, as this avoids non-linear increases in infection levels in subsequent weeks. In other words, the net shadow benefits of shifting individuals from an infected state into a recovered state ($\lambda_3 - \lambda_2$) is large because this avoids infections later in the cycle. The second pulse near the end of the production cycle is likely optimal because infected individuals have no value at harvest (by assumption). Thus, it becomes optimal to shift as many individuals as possible from an infected state into a recovered state.

FIGURE 3.2. Base Model Simulation

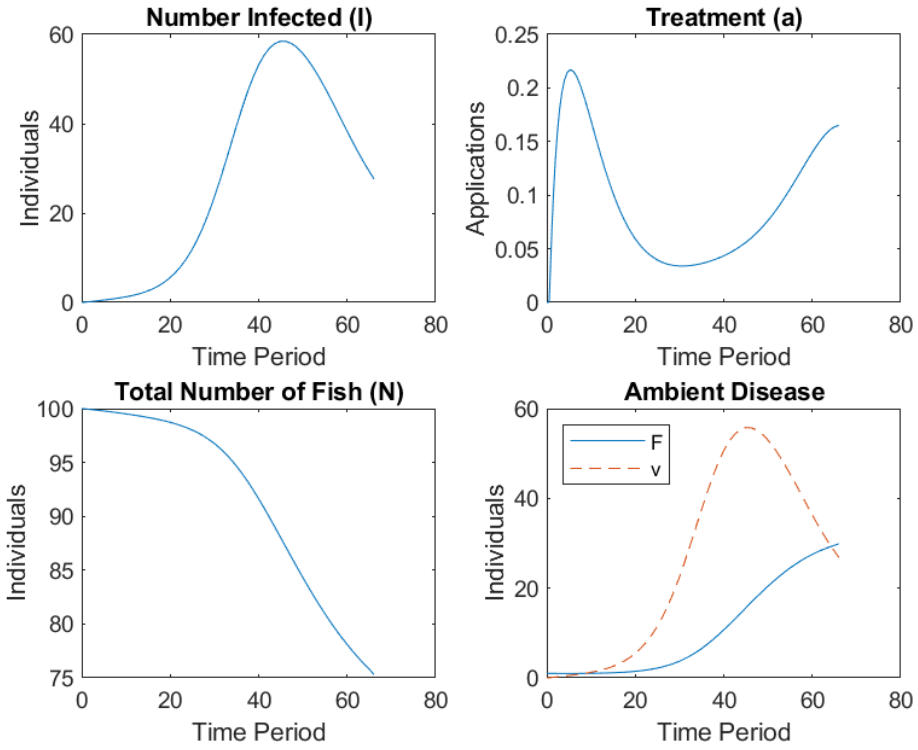

Note: The figure presents the simulation results for the base model. The first panel depicts the number of infected individuals (I), the second panel depicts treatment intensity (a, the third panel depicts the total number of live fish remaining on the farm at each time period, where $N_t = S_t + I_t + R_t$. The fourth panel depicts the ambient disease level, where v is defined as in section 3.2.3

The base model model simulation also illustrates how our parameterization of v relates to the level of ambient disease pressure, F. As described in a previous section, v varies over time to reflect a particular trajectory of exogenous disease pressure resulting from the coordination of adjacent farm sites. As we can see in Figure 3.2, v generally increases over time, mirroring the level of infected individuals. The second pulse of treatment lowers both the number of infected individuals on the farm, but it also is responsible for the decline in v at the end of the cycle. However, despite the fact that I and v fall at the end of the cycle, F continues to grow (though at a slower rate).

Sensitivity analysis on the base model reveals an unexpected set of relationships between infection pressure, F, treatment intensity, a, and treatment effectiveness, β_4. The results of the

sensitivity analysis are reported in Appendix B.3, but we will briefly discuss them here in order to provide additional context for optimal treatment strategy in the presence of endogenous spillovers. Specifically, it is reasonable to hypothesize that reducing outside infection pressure and increasing the effectiveness of antibiotics will reduce the need for antibiotics, all else equal. However, our model suggests a different set of behaviors. First, as treatment effectiveness increases (β_4), the optimal number of applications and the volume of antibiotic use increases. In other words, if the benefit of increased treatment effectiveness exceeds the additional cost of application, then improved antibiotics might drive higher, rather than lower rates of application. Second, a reduction in the infection pressure may not lead to a reduction in antibiotic use. Antibiotic use is highest when starting levels of F are the smallest. If F were allowed to go to zero, then eradication of the disease would be possible. When F_0 approaches zero, the optimal control strategy appears to approach an eradication-style strategy. Although true eradication is not possible, high levels of antibiotic applications to drive F to very low levels may be extremely profitable.

TABLE 3.2. Simulation Results

	Base	Cap on Volume	Cap on No. Apps.	Max Treat Weight	Cap on Volume + Max Treat Weight
π	1.00	0.88	0.88	0.38	0.70
No. of Applications	1.00	0.30	0.30	0.34	0.30
Volume of Antibiotics	1.00	0.50	0.50	0.50	0.50
Cycle Disease Mortality Rate	1.00	1.58	1.58	1.27	1.46
Cycle Cumulative F	1.00	1.66	1.64	1.41	1.58

This table illustrates the performance of each policy relative to the base model. Each reported value is a normalized metric that can be interpreted as a percent.

3.5.2. Policy Simulations. Table 3.2 presents the simulation results relative to the outcomes in the base model. We focus on five key metrics to understand the performance of the policies: the value of the functional, π, the volume of antibiotic applications, the number of antibiotics applications, the disease mortality rate for the entire production cycle, and the accumulated pathogen pressure experienced during the entire production cycle. The outcomes for each of these measures in the base model is normlized to one. Based on the simulations, a cap on the volume of antibiotics is the most cost-effective approach to achieving the targeted level of abatement. This is not surprising and aligns with conventional economic theory. Interestingly, the restriction on the

number of applications achieves a nearly identical outcome, despite being a less precise policy tool. The restriction on the timing of antibiotic treatments is the least cost-effective instrument. This is likely because it dramatically limits the extent to which a farm can distribute antibiotics across the production cycle.

It is also important to emphasize the magnitudes of the policy impacts. The cap on volume and the cap on the number of applications reduced the volume of antibiotics by 50% but only reduced π by 12%. The same reduction in antibiotics lead to a 62% reduction in π. These differences persist despite the fact the maximum treatment weight restriction experienced lower mortality rates and lower cumulative disease pressure. Although these results clearly depend upon the parameterization, they underscore the importance of flexibility in antibiotic abatement.

Figure 3.3 illustrates the optimal treatment schedules under the base scenario and alternative policy simulations. In each case, mandatory antibiotic abatement shifts treatment towards the end of the production cycle. This holds, regardless of the policy instrument. Both of the policies which include treatment weight also lead to larger maximum treatment intensity than policies which restrict the number or volume of applications.

The results presented here reflect the cost-effective approach for individuals, private decision-makers, rather than the cost-effective solution for the entire industry. Table 3.2 suggests that while a cap on the volume of antibiotics or on the number of applications may be a cost-effective approach for an individual farm to reduce antibiotics, they may impose relatively high costs on neighbors via disease transmission.

FIGURE 3.3. Treatment Schedule Under Alternative Abatement Policies

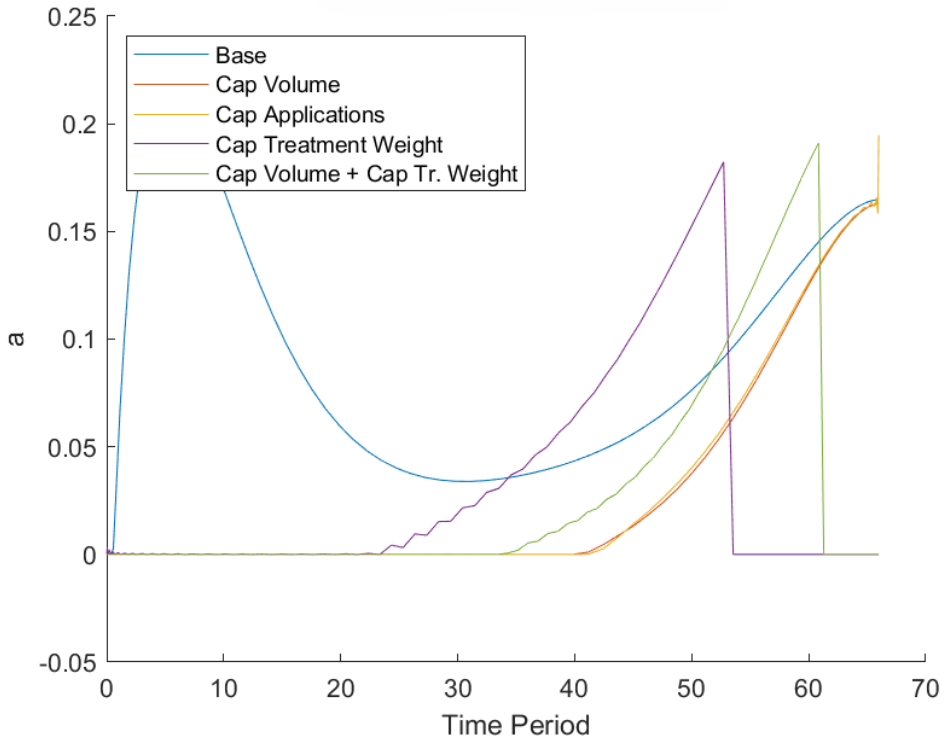

Note: This figure illustrates the treatment schedule for the base parameterization and for each of antibiotic abatement policies described in the main body of the paper.

3.6. Discussion

This paper explores cost-effective approaches to reducing the use of antibiotics in the presence of a communicable agricultural disease. Using Chile's salmon aquaculture industry as a motivating example, we illustrate that a mandatory reduction in the use of antibiotics is likely to increase the disease mortality rate, increase costs, and increase the disease-related externality, regardless of the instrument. However, a cap on the volume of antibiotic use is the most cost-effective approach to reducing the use of antibiotics. All of the antibiotic reduction instruments lead to an increase in the the mortality rate, an increase in costs, and an increase in the ambient disease level (our measure of the disease-related production externality). Interestingly, a restriction on the number of

applications generates a nearly identical solution but both are more cost-effective than a constraint on maximum treatment weight.

The simulations illustrate the importance of dynamics in the production process. Each policy generates a response that is non-uniform across the production cycle. Instead of inducing a vertical or horizontal shift in the treatment schedule, the restrictions redistribute antibiotic applications towards the end of the production cycle. In general, this reflects the fact that the marginal benefits of treatment are a function of the time-varying shadow values of susceptible, infected, and recovered fish. However, there are two mechanisms that determine the shadow value. First, treatment determines the number of fish available for harvest. Under the restrictions, the largest net benefits from antibiotics come at the end of the cycle, where treatment converts infected fish (assumed to have no value at harvest) to recovered fish (which can be sold at the market price). As the end of the cycle approaches, it is clear that the value of an infected fish falls in comparison to recovered or susceptible fish.

The shadow values are also determined by the ambient disease level. Infected fish shed bacteria into the environment, driving an increase in F. When the link between the farm and the ambient disease level is strong (i.e. for large values of β_6 and β_9), farms have an additional incentive to use antibiotics and maintain low levels of I and F. This spillover-spillback linkage creates an additional dynamic private incentive for using antibiotics. However, as the end of the production cycle approaches, the value of avoiding high levels of ambient infection falls. If harvest is imminent, then there is no value to avoiding future infection via spillback from F. In other words, there are no more dynamic benefits associated with treatment at the end of the cycle. Given that the restrictions on antibiotic use pushed treatment to the end of the cycle, this suggests the benefits from increased harvest are more valuable than the dynamic benefits of avoided future infection.

The combination of policy restrictions captured by treatment cap and harvest timing is perhaps the most realistic. In order to avoid contamination, farms are restricted from selling fish that have been treated with antibiotics at any time in the previous month. This creates an intermediate solution that bears similarity to both the treatment cap and to the restriction on harvest timing. However, it is less cost-effective at achieving the targeted reduction than the cap on antibiotic volume in the absence of the timing restriction. This illustrates the intuitive point that seemingly

irrelevant polices (such as those intended to protect seafood consumers) may have important implications for effectiveness of the abatement policies. In this case, it not only makes abatement more expensive by shifting when treatments are applied but, in doing so, it actually reduces the overall ambient disease burden in the system.

One of the key limitations of this study is that it does not integrate non-antibiotic disease control measures such as culling or vaccination into the optimization program. We make this choice for the sake of simplicity but it is well known that these measures can impact disease spread. Nevertheless, our results are informative for understanding the value of alternative technologies. If a constraint on antibiotic use tends to push applications towards the end of the cycle, then the value of an alternative technology will depend upon the time at which they take effect. For example, SRS vaccines are widely used by the Chilean salmon industry but are known to offer temporary and imperfect protection. An improved vaccine would reduce the benefits of (or need for) antibiotics during the weeks in which the vaccine is active. But the vaccine would not necessarily change the marginal benefits of antibiotic consumption later in the production cycle.

In its current formulation, this model focuses on identifying least cost approaches to reducing antibiotic use for the individual farm, but these results do not necessarily extend to the entire industry. While the parameterization of v is convenient, it does not reflect the possibility of strategic feedback between adjacent farms. The current model abstracts from this in order to obtain a realistic, time-varying trajectory for v that is responsive to alternative instruments; in some sense, v is a sort of reduced-form measure of the activity of other farms. But, As noted in an earlier section, our solution does not reflect either an open loop or feedback Nash equilibrium (although we expect that the current results are likely similar to the open loop solution). A feedback Nash equilibrium would allow farms to stategically respond to the expected actions of their neighbors. The impact of incorporating these strategic elements are challenging to predict or contextualize, given that we do not attempt to solve for an industry optimal treatment strategy.

3.7. Conclusion

When antibiotics are important for controlling infectious disease, restricting their use is costly. Using Chile's salmon aquaculture industry as a motivating example, we use calibrated optimization

model to show that reducing the use of antibiotics increase the private costs of production and the externality associated with disease transmission. Furthermore, we illustrate the disease dynamics both within and between farms creates a complex web of incentives. In certain situations, a reduction in the ambient disease level may increase, rather than decrease antibiotic use. Furthermore, the static and dynamic benefits of antibiotic use change over the course of the production cycle.

We illustrate that a cap on the volume of antibiotics is the most cost-effective approach for private decision-makers to achieve a targeted abatement level. This holds despite the fact that this strategy is also associated with the highest mortality rate and the largest disease-related externality of any of the policy simulations. A promising avenue for future work would be to more realistically include the strategic interactions between individuals farms in order to better evaluate the extent to coordinated production cycles influence the need for antibiotics.

Concluding Remarks

This dissertation explores private incentives and public policy for disease control, focusing specifically on Chile's salmon aquaculture industry. Transmissible disease creates a formidable and costly production externality that has driven Chilean producers to utilize antibiotics at a rate that far exceeds other countries. In response to sustainability concerns, the industry has committed to reducing to reducing their use of antibiotics, with many firms working in collaboration with organizations such as Monterey Bay Aquarium Seafood Watch.

We inform these efforts by exploring two main approaches for controlling disease and limiting antibiotic use. First, we consider policy designed to limit transmission between farms. We illustrate empirically that an existing spatial management program improves production conditions for salmon farmers and is appropriately matched the spatial scale of the production externality. This is surprising, given that the evidence that the spatial configuration of the program was not specifically optimized for disease management. Second, we careful consider the private incentives for disease control and, specifically, for antibiotic use. Although antibiotics are critical for disease control in this system, the risk of resistance poses an important sustainability challenge to Chile's industry. Using a parameterized simulation exercise, we illustrate the tradeoffs between efforts to reduce antibiotic use and disease control. Intuitively, we find that restricting the use of antibiotics will increase disease-related losses along with the scale of the disease-externality. Nevertheless, the cost-effectiveness of the policies varies widely, depending upon the chosen policy instrument.

One of the key overarching themes of this dissertation is that structural knowledge of the study system is extremely valuable for empirical work. Theory not only helps to support logically consistent interpretations, but it also useful to motivate a reduced-form empirical strategy. In the first chapter, we used a structural model to motive our empirical strategy and justify alternative and increasingly restrictive assumptions required for alternative econometric estimators. In the second chapter, we use an epidemiological model with endogenous treatment to identify a specific but likely

pervasive form of bias that will undermine estimates of a treatment effectiveness, a critical structural parameter. Although we did not present a direct pathway for an empirical application in chapter three, it offers a provocative set of predictions regarding optimal behavioral responses to ambient disease pressure and treatment effectiveness, even before we pursue the policy simulations. These models along with the simulated, dynamic impacts of alternative antibiotic abatement scenarios, will be informative for future efforts at reducing the Chilean industry's reliance on antibiotics.

This dissertation also provides a foundation for a rich research agenda on pest and disease management as well as the economics of dynamic systems, in general. For example, We have not addressed strategic behavior in either our description of the private incentives facing fish farmers or in the design of the coordinating policy. To some extent, this omission is justified through our focus on economic decisions that are made over relatively short time-horizon, during it which strategic behavior may be limited. However, firms (which own and operate multiple farms) may exhibit a wide range of behaviors across many different margins that are fixed within a production cycle but can vary on a longer time horizon (e.g. utilization of lease sites, choice of species, etc.). These margins for adjustment are consequential and may provide a more holistic view of the impacts of disease management policy.

APPENDIX A

Supplementary Material for Chapter 1

A.1. Database Description

In this section, we describe the process by which the data were cleaned and processed.

A.1.1. Geographic Distances. To implement the empirical strategy, we need to match each production center with its nearest shared neighborhood border. To do this, we measured the seaway distance between each production site and each shared neighborhood border segment that extended over water. We eliminated the segments which offered a division between a water body that was completely enclosed by a landmass and the border line but which did not contain any production sites.

Seaway distances were calculated using the following procedure: First, we converted the polygons representing the production sites into single points by calculating the centroid of the polygon. Next, we converted the publicly available vector representation of the Chilean coastline into a series of raster files. The resolution of each raster file was chosen such that each pixel was no larger than 40 m^2 or roughly 0.0004 degrees. After a detailed examination of the production region, we determined that this resolution would be sufficient to allow a route-finding algorithm to traverse relevant water passages. We then estimated the shortest border distances by allowing the route-finding algorithm to search in eight directions at each step.

In some cases, the centroid of the production sites did not overlap with a raster pixel representing water. In such cases, we removed the portion of the production site that overlapped with land and re-calculated the centroid. If this procedure did not eliminate the problem, we dropped the production site from the analysis. This led to the loss of four total concessions.

A.1.2. Data Filtering. The SIFA and INFA databases contain information for the entire Chilean salmon industry but we choose to limit our sample in a few important ways. A full description of database construction is included in the appendix but two key features are particularly

important for our identification strategy. First, we focus our analysis on Atlantic salmon, even though Chile's industry produces three different species: Atlantic Salmon, Coho Salmon, and Rainbow Trout. Atlantic salmon is the most important both by volume and economic value and are most relevant for policy analysis. Coho salmon are not required to report ambient environmental data on a weekly basis, severely limiting their utility in our model. They are also thought to be much less susceptible to SRS than Atlantic salmon or Rainbow Trout.

Although we focus on outcomes on Atlantic salmon farms, we choose to include information from all three species when calculating the effective date in which neighborhoods switch from fallow to active as well as when calculating total pathogen prevalence levels. We believe this to be the most appropriate assumption given that both coho salmon and rainbow trout are carriers of the disease and can ultimately, contribute to the overall pathogen load in the system. We test the robustness of our results to this assumption in the appendix.

Second, we limit our analysis to Region X and Region XI where the majority of salmon farming activity occurs. As Figure 1.2 illustrates, few new concessions have been granted in these two regions in the past decade, making the policy context fairly stable during the analysis. The salmon farming industry has been rapidly expanding into Region XII, leading to the creation of many new neighborhoods since 2010. However, SRS infection does not appear to be an important challenge in this region so we will not include observations from this region as part of the analysis.

A.2. The DiD Estimator

FIGURE A.1. Event Study Model Estimates with Relaxed Data Filtering

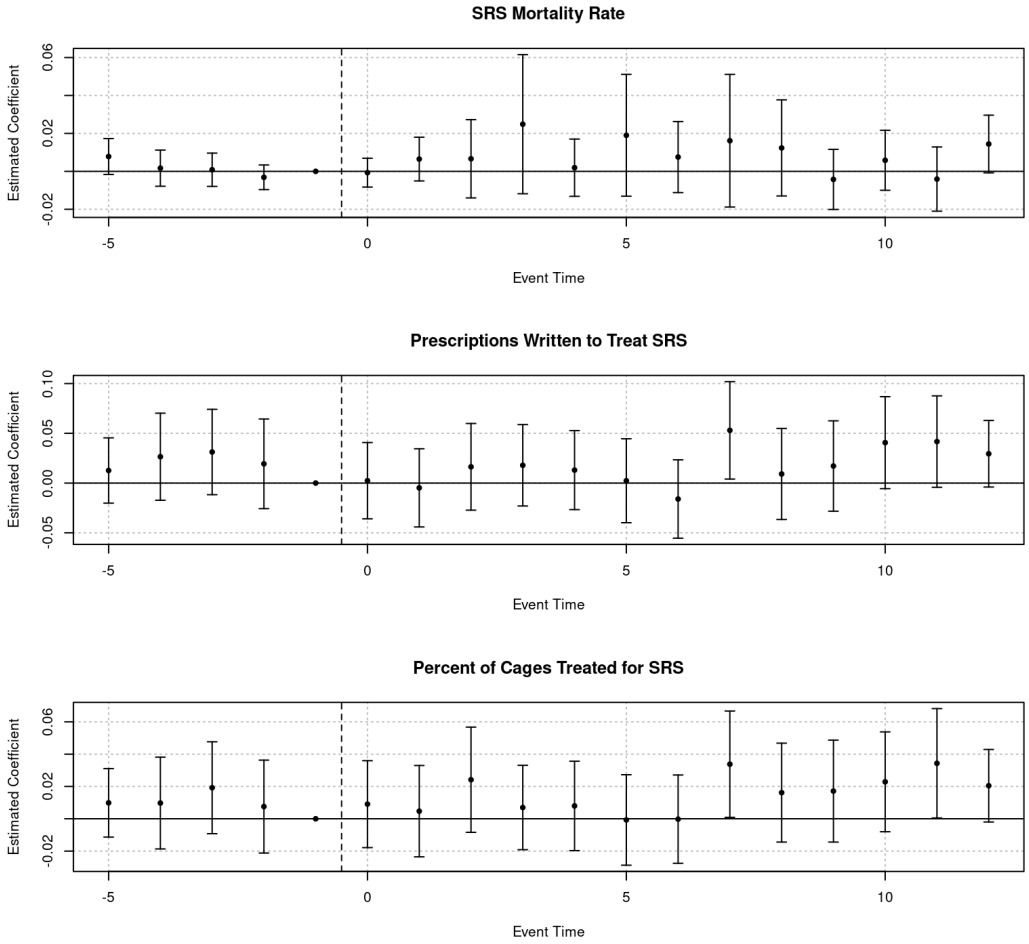

Note: This figure depicts the coefficients of the event study model (β_r) and confidence intervals, calculated at $\alpha = 0.05\%$. The coefficients are calculated relative to the base group at event time -1. The end of fallowing is illustrated by the dashed vertical line at event time -0.5. All models include controls for salinity, temperature, cycle week, and year along with fixed effects for the cycle and cycle week. Standard errors are clustered at the production-cycle level. The estimating procedure is identical to that in the main paper however, we include a slightly less strict definition of the treatment group. We include production cycles whose nearest adjacent neighborhood started in production but then switched into a fallow period and then switched back into production. This significantly increases the sample size but does not change our results.

APPENDIX B

Supplementary Material for Chapter 3

B.1. Calculating Parameters of Weight Function

We assume that weight follows a logistic growth function (Thyholdt, 2014), where fish weight is governed by the following differential equation:

$$\dot{W} = rW_t\left(1 - \frac{W_t}{k}\right) \tag{B.1}$$

The solution to the logistic growth function is well known and has been expressed in many different forms. We use the following:

$$W_t = \frac{kW_0}{W_0 + (k - W_0)e^{-rt}} \tag{B.2}$$

We assume that the initial stocking weight, W_0, the harvest weight, W_T, the carrying capacity, k, and the length of the production cycle are all exogenous. In doing so, we also determine the growth rate, r:

$$r = \frac{\log(W_T(k - W_0)) - \log(W_0(k - W_T))}{T} \tag{B.3}$$

To give a sense of the growth function under the model parameters, W_t is plotted below:

FIGURE B.1. Treatment Schedule Under Alternative Abatement Policies

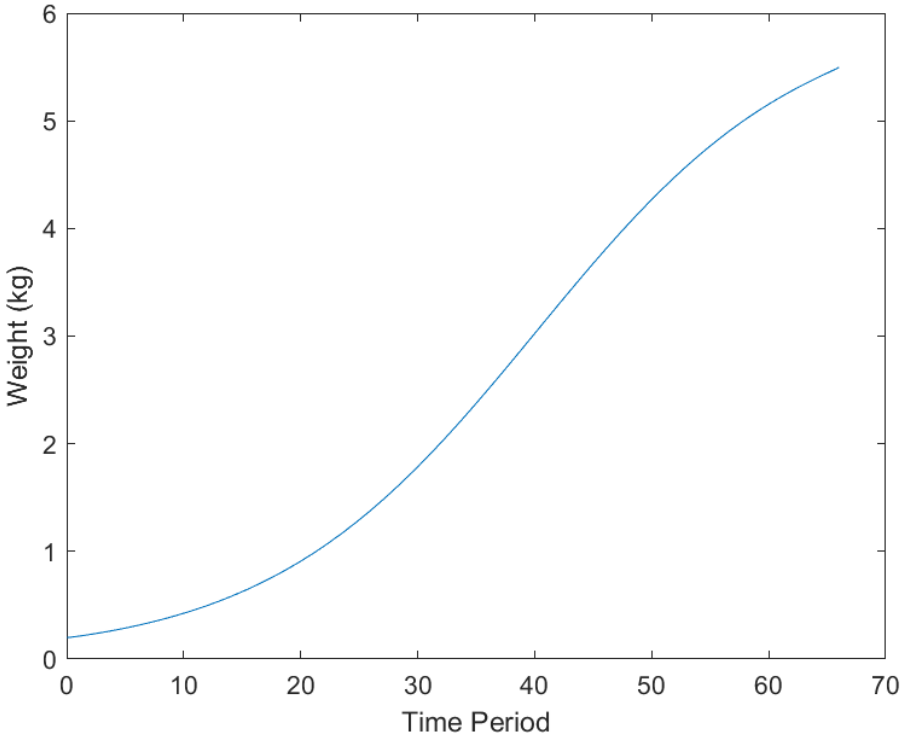

Note: This figure illustrates the growth function under the baseline parameters.

TABLE B.1. Fish Weight Model Parameters

	Parameter
W_0	0.2
W_T	5.5
k	6
T	66
r	0.087351

B.2. Policy Parameters

TABLE B.2. Parameters for Abatement Policy Simulation

	Base	Cap on Volume	Cap on No. Apps.	Max Treat. Weight	Cap on Volume + Max Treat Weight
ϕ	1	0.5	0.5	0.5	0.5
$\bar{V}$		629.4081			629.4081
$\bar{A}$			1.8489		
τ				12.4375	4

This table includes the parameter values use to simulate the antibiotic abatement policies. Each row represents a different parameter and are defined within the text. Each column represents a scenario. The base model includes no policy restrictions.

B.3. Sensitivity of Base Model

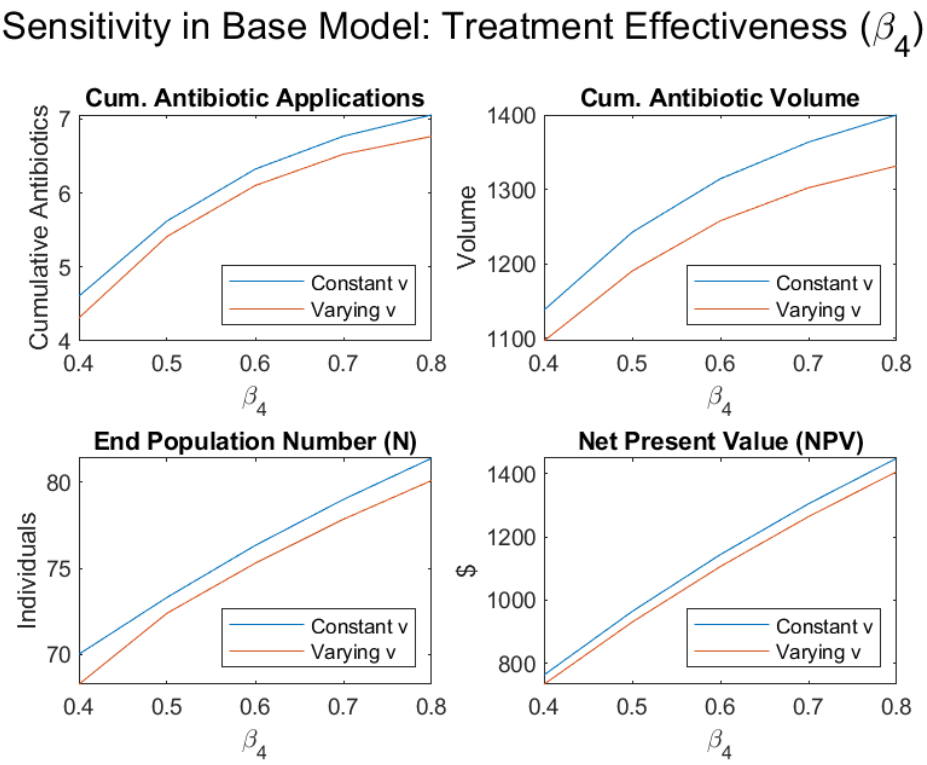

Note: This figure illustrates the sensitivity of the base simulation to different values of treatment effectiveness (β_4 in the epidemiological model) In each panel, there are two scenarios: constant v assumes that exogenous disease pressure is not time-varying. Varying v assumes that exogenous disease pressure does vary over time, following the path of infected individuals.

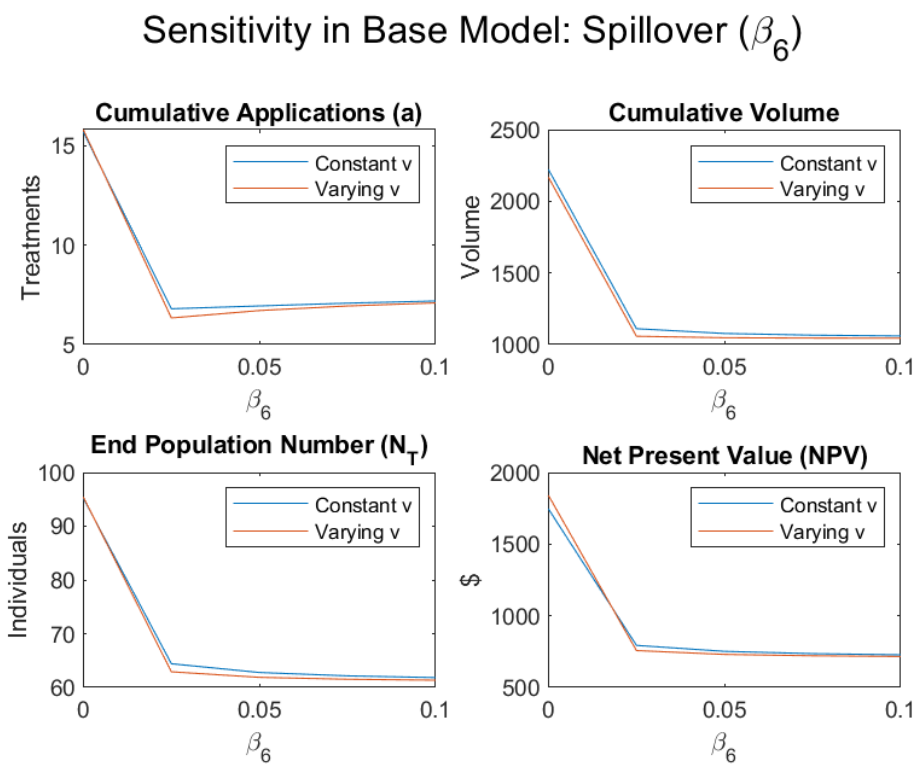

Note: This figure illustrates the sensitivity of the base simulation to different values of spillover intensity (β_6 in the epidemiological model) In each panel, there are two scenarios: constant v assumes that exogenous disease pressure is not time-varying. Varying v assumes that exogenous disease pressure does vary over time, following the path of infected individuals.

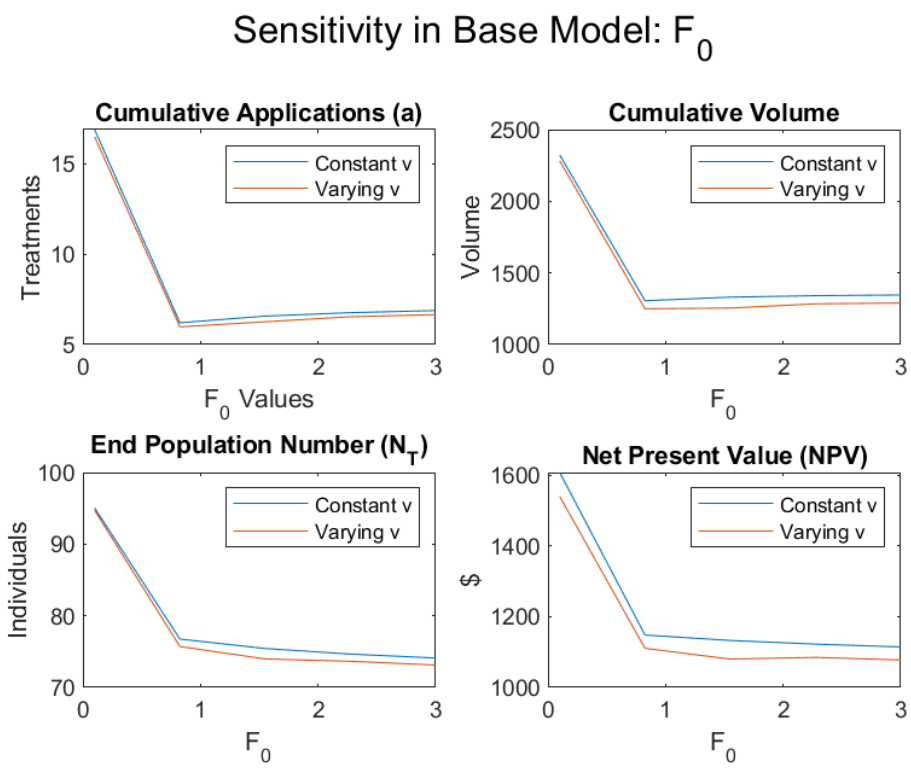

Note: This figure illustrates the sensitivity of the base simulation to different starting values of ambient infection (F_0 in the epidemiological model) In each panel, there are two scenarios: constant v assumes that exogenous disease pressure is not time-varying. Varying v assumes that exogenous disease pressure does vary over time, following the path of infected individuals.

Bibliography

Abolofia, J. N. (2014). *The Bioeconomics of a Common Property Pest: Parasitic Sea Lice and Farmed Salmonids*. [Doctoral dissertation, University of California, Davis]. ProQuest Dissertations and Theses.

Aguilar-Manjarrez, J., Soto, D., and Brummett, R. E. (2017). *Aquaculture zoning, site selection and area management under the ecosystem approach to aquaculture: a handbook*. Food and Agriculture Organization of the United Nations ; The World Bank, Rome. OCLC: 994368479.

Albers, H. J., Fischer, C., and Sanchirico, J. N. (2010). Invasive species management in a spatially heterogeneous world: Effects of uniform policies. *Resource and Energy Economics*, 32(4):483–499.

Alix-Garcia, J. M., Shapiro, E. N., and Sims, K. R. E. (2012). Forest Conservation and Slippage: Evidence from Mexico's National Payments for Ecosystem Services Program. *Land Economics*, 88(4):613–638. Publisher: University of Wisconsin Press Section: Articles.

Almendras, F. E. and Funealba, I. (1997). Salmonid rickettsial septicemia caused by piscirickettsia salmonis:a review. *Diseases of Aquatic Organisms*, 29:137–144.

Alvial, A. (2017). Chile Case: The Spatial Planning of Marine Cage Farming (Salmon). In *Aquaculture zoning, site selection and area management under the ecosystem approach to aquaculture. Full document*, number ACS113536, pages 170–197. World Bank Group, Washington D.C.

Anderson, R. M. and May, R. M. (1979). Population biology of infectious diseases: Part I. *Nature*, 280(5721):361–367.

Ando, A., Camm, J., Polasky, S., and Solow, A. (1998). Species Distributions, Land Values, and Efficient Conservation. *Science*, 279(5359):2126–2128. Publisher: American Association for the Advancement of Science.

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of Causal Effects Using Instrumental Variables. *Journal of the American Statistical Association*, 91(434):444–455.

Arellano, M. and Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *The Review of Economics Studies*, 58(2):277–297.

Asche, F., Hansen, H., Tveterås, R., and Tveterås, S. (2009). The salmon disease crisis in Chile. *Marine Resource Economics*, 24(4):405–411.

Atallah, S. S., Gómez, M. I., and Conrad, J. M. (2017). Specification of spatial-dynamic externalities and implications for strategic behavior in disease control. *Land Economics*, 93(2):209–229.

Auld, M. C. (2003). Choices, beliefs, and infectious disease dynamics. *Journal of Health Economics*, 22(3):361–377.

Avendaño-Herrera, R., Mancilla, M., and Miranda, C. D. (2023). Use of antimicrobials in chilean salmon farming: Facts, myths and perspectives. *Reviews in Aquaculture*, 15(1):89–111.

Baumol, W. J. and Oates, W. E. (1988). *The theory of environmental policy*. Cambridge University Press, 2 edition.

Bhat, M. G. and Huffaker, R. G. (2007). Management of a transboundary wildlife population: a self-enforcing cooperative agreement with renegotiation and variable transfer payments. *Journal of Environmental Economics and Management*, 53(1):54–67.

Bicknell, K. B., Wilen, J. E., and Howitt, R. E. (1999). Public Policy and Private Incentives for Livestock Disease Control. *Australian Journal of Agricultural and Resource Economics*, 43(4):501–521.

Bjørndal, T. (1988). Optimal Harvesting of Farmed Fish. *Marine Resource Economics*, 5(2):139–159. Publisher: The University of Chicago Press.

Bravo, F., Sidhu, J., Bernal, P., Bustamante, R. H., Condie, S., Gorton, R., Herzfeld, M., Jiménez, D., Mardones, F. O., Rizwi, F., and Steven, A. D. L. (2020). Hydrodynamic connectivity, water temperature, and salinity are major drivers of piscirickettsiosis prevalence and transmission among salmonid farms in Chile. *Aquaculture Environment Interactions*, 12:263–279.

Bravo, S. and Campos, M. (1989). Coho salmon syndrome in Chile. *FHS/AFS Newsletter*, 17(3).

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Cartes, C., Isla, A., Lagos, F., Castro, D., Munoz, M., Yanez, A., Haussmann, D., , and Figueroa, J. (2017). Search and analysis of genes involved in antibiotic resistance in chilean strains of

Piscirickettsia salmonis. *Journal of Fish Diseases*, 40(8).

Castonguay, F. M., Sokolow, S. H., De Leo, G. A., and Sanchirico, J. N. (2020). Cost-effectiveness of combining drug and environmental treatments for environmentally transmitted diseases. *Proceedings of the Royal Society B: Biological Sciences*, 287(1933):20200966.

Chiang, A. C. (1992). *Elements of Dynamic Optimization*. Waveland Press, Inc.

Coase, R. H. (1960). The problem of social cost. *The Journal of Law and Economics*, 3:1–44.

Dangerfield, C., Fenichel, E. P., Finnoff, D., Hanley, N., Hargreaves Heap, S., Shogren, J. F., and Toxvaerd, F. (2022). Challenges of integrating economics into epidemiological analysis of and policy responses to emerging infectious diseases. *Epidemics*, 39:100585.

de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, 110(9):2964–2996.

Diekmann, O., Othmer, H. G., Planqué, R., and Bootsma, M. C. J. (2010). The discrete-time kermack–mckendrick model: A versatile and computationally attractive framework for modeling epidemics. *Proceedings of the National Academy of Sciences*, 118(39).

Dresdner, J., Chávez, C., Quiroga, M., Jiménez, D., Artacho, P., and Tello, A. (2019). Impact of *Caligus* treatments on unit costs of heterogeneous salmon farms in Chile. *Aquaculture Economics & Management*, 23(1):1–27.

Epanchin-Niell, R. S. and Wilen, J. E. (2015). Individual and cooperative management of invasive species in human-mediated landscapes. *American Journal of Agricultural Economics*, 97(1):180–198.

Fenichel, E. P., Castillo-Chavez, C., Ceddia, M. G., Chowell, G., Parra, P. A. G., Hickling, G. J., Holloway, G., Horan, R., Morin, B., Perrings, C., Springborn, M., Velazquez, L., and Villalobos, C. (2011). Adaptive human behavior in epidemiological models. *Proceedings of the National Academy of Sciences*, 108(15):6306–6311.

Ferraro, P. J., Sanchirico, J. N., and Smith, M. D. (2019). Causal inference in coupled human and natural systems. *Proceedings of the National Academy of Sciences*, 116(12):5311–5318.

Fischer, C., Guttormsen, A. G., and Smith, M. D. (2017). Disease risk and market structure in salmon aquaculture. *Water Economics and Policy*, 03(02):1650015.

Francis, P. J. (1997). Dynamic epidemiology and the market for vaccinations. *Journal of Public Economics*, 63(3):383–406.

Frazer, L. N., Morton, A., and Krkošek, M. (2012). Critical thresholds in sea lice epidemics: evidence, sensitivity and subcritical estimation. *Proceedings of the Royal Society B: Biological Sciences*, 279(1735):1950–1958.

Fuller, K. B., Sanchirico, J. N., and Alston, J. M. (2017). The spatial-dynamic benefits from cooperative disease control in a perennial crop. *Journal of Agricultural and Resource Economics*, 42(2):127–145.

Garg, D., Patterson, M., Hager, W. W., Rao, A. V., Benson, D. A., and Huntington, G. T. (2010). A unified framework for the numerical solution of optimal control problems using pseudospectral methods. *Automatica*, 46(11):1843–1851.

Geoffard, P.-Y. and Philipson, T. (1996). Rational Epidemics and Their Public Control. *International Economic Review*, 37(3):603–624. Publisher: [Economics Department of the University of Pennsylvania, Wiley, Institute of Social and Economic Research, Osaka University].

Gersovitz, M. and Hammer, J. S. (2005). Tax/subsidy policies toward vector-borne infectious diseases. *Journal of Public Economics*, 89(4):647–674.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Guarracino, M., Qviller, L., and Lillehaug, A. (2018). Evaluation of aquaculture management zones as a control measure for salmon lice in Norway. *Diseases of Aquatic Organisms*, 130(1):1–9.

Guttormsen, A. G. (2008). Faustmann in the Sea: Optimal Rotation in Aquaculture. *Marine Resource Economics*, 23(4):401–410. Publisher: [MRE Foundation, Inc., University of Chicago Press].

Hazen, E. L., Scales, K. L., Maxwell, S. M., Briscoe, D. K., Welch, H., Bograd, S. J., Bailey, H., Benson, S. R., Eguchi, T., Dewar, H., Kohin, S., Costa, D. P., Crowder, L. B., and Lewison, R. L. (2018). A dynamic ocean management tool to reduce bycatch and support sustainable fisheries. *Science Advances*, 4(5):eaar3001.

Heaps, T. (1995). Density Dependent Growth and the Culling of Farmed Fish. *Marine Resource Economics*, 10(3):285–298. Publisher: The University of Chicago Press.

Hennessy, D. A. (2007). Biosecurity and Spread of an Infectious Animal Disease. *American Journal of Agricultural Economics*, 89(5):1226–1231.

Hennessy, D. A. and Marsh, T. L. (2021). Economics of animal health and livestock disease. In *Handbook of Agricultural Economics*, volume 5, pages 4233–4330. Elsevier.

Horan, R. D. and Fenichel, E. P. (2007). Economics and Ecology of Managing Emerging Infectious Animal Diseases. *American Journal of Agricultural Economics*, 89(5):1232–1238. Publisher: [Agricultural & Applied Economics Association, Oxford University Press].

Horan, R. D. and Wolf, C. A. (2005). The Economics of Managing Infectious Wildlife Disease. *American Journal of Agricultural Economics*, 87(3):537–551.

Igoe, M., Casagrandi, R., Gatto, M., Hoover, C. M., Mari, L., Ngonghala, C. N., Remais, J. V., Sanchirico, J. N., Sokolow, S. H., Lenhart, S., and de Leo, G. (2023). Reframing Optimal Control Problems for Infectious Disease Management in Low-Income Countries. *Bulletin of Mathematical Biology*, 85(4):31.

Iizuka, M. (2016). Transformation of Institutions: Crisis and Change in Institutions for Chilean Salmon Industry. In Hosono, A., Iizuka, M., and Katz, J., editors, *Chile's Salmon Industry*, pages 137–174. Springer Japan, Tokyo.

Ives, A. R. and Settle, W. H. (1997). Metapopulation dynamic and pest control in agricultural systems. *The American Naturalist*, 149(2):220–246.

Judd, K. L. (1998). *Numerical Methods in Economics*. The MIT Press.

Kremer, M. (1996). Integrating Behavioral Choice into Epidemiological Models of AIDS. *The Quarterly Journal of Economics*, 111(2):549–573. Publisher: Oxford University Press.

Kristoffersen, A. B., Rees, E. E., Stryhn, H., Ibarra, R., Campistó, J. L., Revie, C. W., and St-Hilaire, S. (2013). Understanding sources of sea lice for salmon farms in Chile. *Preventive Veterinary Medicine*, 111(1-2):165–175.

Krkošek, M., Bateman, A., Proboszcz, S., and Orr, C. (2010). Dynamics of outbreak and control of salmon lice on two salmon farms in the Broughton Archipelago, British Columbia. *Aquaculture Environment Interactions*, 1(2):137–146.

Lannan, C. N. and Fryer, J. L. (1994). Extracellular survival of *Piscirickettsia salmonis*. *Journal of Fish Diseases*, 17(5):545–548.

Laxminarayan, R. and Brown, G. M. (2001). Economics of Antibiotic Resistance: A Theory of Optimal Use. *Journal of Environmental Economics and Management*, 42(2):183–206.

Laxminarayan, R., Duse, A., Wattal, C., Zaidi, A. K. M., Wertheim, H. F. L., Sumpradit, N., Vlieghe, E., Hara, G. L., Gould, I. M., Goossens, H., Greko, C., So, A. D., Bigdeli, M., Tomson, G., Woodhouse, W., Ombaka, E., Peralta, A. Q., Qamar, F. N., Mir, F., Kariuki, S., Bhutta, Z. A., Coates, A., Bergstrom, R., Wright, G. D., Brown, E. D., and Cars, O. (2013). Antibiotic resistance—the need for global solutions. *The Lancet Infectious Diseases*, 13(12):1057–1098. Publisher: Elsevier.

Laxminarayan, R., Matsoso, P., Pant, S., Brower, C., Røttingen, J.-A., Klugman, K., and Davies, S. (2016). Access to effective antimicrobials: a worldwide challenge. *The Lancet*, 387(10014):168–175.

LeSage, J. and Pace, R. K. (2009). *Introduction to Spatial Econometrics*. Chapman and Hall/CRC, 0 edition.

Lipscomb, M. and Mobarak, A. M. (2017). Decentralization and Pollution Spillovers: Evidence from the Re-drawing of County Borders in Brazil. *The Review of Economic Studies*, 84(1 (298)):464–502. Publisher: [Oxford University Press, The Review of Economic Studies, Ltd.].

Long, A. and Jones, S. R. M. (2021). Piscirickettsia salmonis shedding and tissue burden, and hematological responses during cohabitation infections in chum Oncorhynchus keta, pink O. gorbuscha and Atlantic salmon Salmo salar. *PLOS ONE*, 16(3):e0248098. Publisher: Public Library of Science.

Maisey, K., Montero, R., and Christodoulides, M. (2017). Vaccines for piscirickettsiosis (salmonid rickettsial septicaemia, SRS): the Chile perspective. *Expert Review of Vaccines*, 16(3):215–228.

McAdams, D. (2021). The Blossoming of Economic Epidemiology. *Annual Review of Economics*, 13(1):539–570. _eprint: https://doi.org/10.1146/annurev-economics-082120-122900.

Mistiaen, J. A. and Strand, I. (1998). Optimal Feeding and Harvest Time for Fish with Weight-Dependent Prices. *Marine Resource Economics*, 13(4):231–246. Publisher: [MRE Foundation, Inc., University of Chicago Press].

Murray, A. G. and Peeler, E. J. (2005). A framework for understanding the potential for emerging diseases in aquaculture. *Preventive Veterinary Medicine*, 67(2-3):223–235.

Murray, A. G. and Salama, N. K. (2016). A simple model of the role of area management in the control of sea lice. *Ecological Modelling*, 337:39–47.

Nickell, S. (1981). Biases in Dynamic Models with Fixed Effects. *Econometrica*, 49(6):1417.

Oates, W. E. (1972). *Fiscal federalism*. The Harbrace series in business and economics. Harcourt, Brace, Jovanovich, New York, NY.

Oglend, A. and Soini, V.-H. (2020). Implications of entry restrictions to address externalities in aquaculture: the case of salmon aquaculture. *Environmental and Resource Economics*, 77(4):673–694.

Olmstead, A. L. and Rhode, P. W. (2004). An Impossible Undertaking: The Eradication of Bovine Tuberculosis in the United States. *The Journal of Economic History*, 64(3):734–772.

Olson, T. K. and Criddle, K. R. (2008). Industrial evolution: a case study of chilean salmon aquaculture. *Aquaculture Economics & Management*, 12(2):89–106.

Ostrom, E. (2015). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1 edition.

Peacock, S. J., Bateman, A. W., Krkošek, M., and Lewis, M. A. (2016). The dynamics of coupled populations subject to control. *Theoretical Ecology*, 9(3):365–380.

Peacock, S. J., Krkošek, M., Bateman, A. W., and Lewis, M. A. (2020). Estimation of spatiotemporal transmission dynamics and analysis of management scenarios for sea lice of farmed and wild salmon. *Canadian Journal of Fisheries and Aquatic Sciences*, 77(1):55–68.

Perrings, C., Castillo-Chavez, C., Chowell, G., Daszak, P., Fenichel, E. P., Finnoff, D., Horan, R. D., Kilpatrick, A. M., Kinzig, A. P., Kuminoff, N. V., Levin, S., Morin, B., Smith, K. F., and Springborn, M. (2014). Merging Economics and Epidemiology to Improve the Prediction and Management of Infectious Disease. *EcoHealth*, 11(4):464–475.

Pfeiffer, L. and Lin, C. Y. C. (2012). Groundwater pumping and spatial externalities in agriculture. *Journal of Environmental Economics and Management*, 64(1):16–30.

Price, D., Ibarra, R., Sánchez, J., and St-Hilaire, S. (2017). A retrospective assessment of the effect of fallowing on piscirickettsiosis in Chile. *Aquaculture*, 473:400–406.

Price, D., Stryhn, H., Sanchez, J., Ibarra, R., Tello, A., and St-Hilaire, S. (2016). Retrospective analysis of antibiotic treatments against piscirickettsiosis in farmed Atlantic salmon *Salmo salar*

in Chile. *Diseases of Aquatic Organisms*, 118(3):227–235.

Rees, E. E., Ibarra, R., Medina, M., Sanchez, J., Jakob, E., Vanderstichel, R., and St-Hilaire, S. (2014). Transmission of Piscirickettsia salmonis among salt water salmonid farms in Chile. *Aquaculture*, 428-429:189–194.

Reimer, M. N. and Haynie, A. C. (2018). Mechanisms matter for evaluating the economic impacts of marine reserves. *Journal of Environmental Economics and Management*, 88:427–446.

Revie, C. W., Robbins, C., Gettinby, G., Kelly, L., and Treasurer, J. W. (2005). A mathematical model of the growth of sea lice, Lepeophtheirus salmonis, populations on farmed Atlantic salmon, Salmo salar L., in Scotland and its use in the assessment of treatment strategies. *Journal of Fish Diseases*, 28(10):603–613. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1365-2761.2005.00665.x.

Reynolds, M. D., Sullivan, B. L., Hallstein, E., Matsumoto, S., Kelling, S., Merrifield, M., Fink, D., Johnston, A., Hochachka, W. M., Bruns, N. E., Reiter, M. E., Veloz, S., Hickey, C., Elliott, N., Martin, L., Fitzpatrick, J. W., Spraycar, P., Golet, G. H., McColl, C., Low, C., and Morrison, S. A. (2017). Dynamic conservation for migratory species. *Science Advances*, 3(8):e1700707.

Robalino, J., Pfaff, A., and Villalobos, L. (2017). Heterogeneous Local Spillovers from Protected Areas in Costa Rica. *Journal of the Association of Environmental and Resource Economists*, 4(3):795–820.

Rouhi Rad, M., Manning, D. T., Suter, J. F., and Goemans, C. (2021). Policy Leakage or Policy Benefit? Spatial Spillovers from Conservation Policies in Common Property Resources. *Journal of the Association of Environmental and Resource Economists*, 8(5):923–953.

Rowthorn, R. E., Laxminarayan, R., and Gilligan, C. A. (2009). Optimal control of epidemics in metapopulations. *Journal of The Royal Society Interface*, 6(41):1135–1144.

Rozas, M. and Enríquez, R. (2014). Piscirickettsiosis and Piscirickettsia salmonis in fish: a review. *Journal of Fish Diseases*, 37(3):163–188. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jfd.12211.

Sanchirico, J. N. and Springborn, M. (2011). How to get there from here: Ecological and economic dynamics of ecosystem service. *Environmental and Resource Economics*, 48:243–267.

Sanchirico, J. N. and Wilen, J. E. (1999). Bioeconomics of Spatial Exploitation in a Patchy Environment. *Journal of Environmental Economics and Management*, 37(2):129–150.

Sanchirico, J. N. and Wilen, J. E. (2005). Optimal spatial management of renewable resources: matching policy scope to ecosystem scale. *Journal of Environmental Economics and Management*, 50(1):23–46.

SERNAPESCA (2022). Informe sobre uso de antimicrobianos en la salmonicultura nacional. Technical report, Subdireccion de Acuicultura, Departmento de Salud Animal, Valparaiso.

Shinn, A. P., Pratoomyot, J., Griffiths, D., Trong, T. Q., Vu, T. N., Jiravanichpaisal, P., and Briggs, M. (2018). Asian shrimp production and the economic costs of disease. *Asian Fisheries Science*, 31:29–58.

Singerman, A. and Useche, P. (2019). The Role of Strategic Uncertainty in Area-wide Pest Management Decisions of Florida Citrus Growers. *American Journal of Agricultural Economics*, 101(4):991–1011.

Smith, M. D., Sanchirico, J. N., and Wilen, J. E. (2009). The economics of spatial-dynamic processes: Applications to renewable resources. *Journal of Environmental Economics and Management*, 57(1):104–121.

Smith, P. and Mardones, F. O. (2020). Piscirickettsiosis (*Piscirickettsia salmonis*). In Woo, P. T. K., Leong, J.-A., and Buchmann, K., editors, *Climate Change and Infectious Fish Diseases*. CAB International.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Taranger, G. L., Karlsen, Ø., Bannister, R. J., Glover, K. A., Husa, V., Karlsbakk, E., Kvamme, B. O., Boxaspen, K. K., Bjørn, P. A., Finstad, B., Madhun, A. S., Morton, H. C., and Svåsand, T. (2015). Risk assessment of the environmental impact of Norwegian Atlantic salmon farming. *ICES Journal of Marine Science*, 72(3):997–1021.

Tecklin, D. (2016). Sensing the Limits of Fixed Marine Property Rights in Changing Coastal Ecosystems: Salmon Aquaculture Concessions, Crises, and Governance Challenges in Southern Chile. *Journal of International Wildlife Law & Policy*, 19(4):284–300.

Thornton, P. K. (2010). Livestock production: recent trends, future prospects. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1554):2853–2867.

Thyholdt, S. B. (2014). The Importance of Temperature in Farmed Salmon Growth: Regional Growth Functions for Norwegian Farmed Salmon. *Aquaculture Economics & Management*, 18(2):189–204. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/13657305.2014.903310.

Tien, J. H. and Earn, D. J. D. (2010). Multiple Transmission Pathways and Disease Dynamics in a Waterborne Pathogen Model. *Bulletin of Mathematical Biology*, 72(6):1506–1533.

Van Boeckel, T. P., Brower, C., Gilbert, M., Grenfell, B. T., Levin, S. A., Robinson, T. P., Teillant, A., and Laxminarayan, R. (2015). Global trends in antimicrobial use in food animals. *Proceedings of the National Academy of Sciences*, 112(18):5649–5654.

Van Boeckel, T. P., Glennon, E. E., Chen, D., Gilbert, M., Robinson, T. P., Grenfell, B. T., Levin, S. A., Bonhoeffer, S., and Laxminarayan, R. (2017). Reducing antimicrobial use in food animals. *Science*, 357(6358):1350–1352.

Werkman, M., Green, D., Murray, A., and Turnbull, J. (2011). The effectiveness of fallowing strategies in disease control in salmon aquaculture assessed with an SIS model. *Preventive Veterinary Medicine*, 98(1):64–73.

White, C. (2024). At CSARP+ launch, Chilean salmon industry leaders acknowledge failure to develop SRS vaccine has stymied antibiotic reduction effort. *Seafood Source News*.

Wilen, J. E. (2007). Economics of Spatial-Dynamic Processes. *American Journal of Agricultural Economics*, 89(5):1134–1144.

Windmeijer, F. (2005). A finite sample correction for the variance of linear efficient two-step GMM estimators. *Journal of Econometrics*, 126(1):25–51.

Wooldridge, J. M. (2015). Control function methods in applied econometrics. *Journal of Human Resources*, 50(2):420–445.

Yu, R. and Leung, P. (2006). Optimal pest management: A reproductive pollutant perspective. *International Journal of Pest Management*, 52(3):155–166. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/09670870600774273.