Essays in Institutional Economics

by

Scott Jordan Lustig

Department of Economics
Duke University

Date: _____

Approved:

_____
Charles Becker, Advisor

_____
Timur Kuran

_____
Atila Abdulkadiroglu

_____
Frank Sloan

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2011

Abstract
(Economics)

Essays in Institutional Economics

by

Scott Jordan Lustig

Department of Economics
Duke University

Date: _____

Approved:

_____
Charles Becker, Advisor

_____
Timur Kuran

_____
Atila Abdulkadiroglu

_____
Frank Sloan

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics in the Graduate School of Duke University
2011

Abstract

This dissertation is a collection of three chapters all pertaining to institutional economics. In short, the field of institutional economics is an outgrowth of public economics, in the sense that in many cases the key institutions that frame economic decisionmaking are the product of public policy. However this is not exclusive. Institutional economics' key contribution is the acknowledgement that cultural and social institutions — often developed organically over the course of centuries — can play as significant a role in individuals' economic choices as governmental policy. In the pages that follow, we will address the economic impact of cultural and political institutions in three contexts: Judicial decisionmaking in Islamic courts, the effects of negative health shocks on retirement savings, and the tradeoff between retirement savings and investment in durable goods.

To my parents, Linda and Barry, and my sister, Alissa,

To my grandparents, Ruth and Martin Lustig, and Suzanne and Samuel Friedman,

And to Sevan Vahanian,

for their constant LOVE, ENCOURAGEMENT, and UNDERSTANDING.

Contents

List of Tables

List of Figures

Acknowledgements

A project as lengthy, complicated, and daunting as a doctoral dissertation simply cannot be accomplished without the support of many people. To begin, I would like to thank the members of my dissertation committee for their support and thoughtful advice. Atila Abdulkadiroglu and Frank Sloan provided what is often the most useful form of support — the ability to view a long project with fresh eyes and an outside perspective. Their academic rigor and professional support were instrumental in moving this dissertation from draft to draft, with each iteration undoubtedly improved by their unique insights.

I owe a special gratitude to my co-author and dissertation committee member Timur Kuran. During the wilderness that all doctoral students must traverse when transitioning from a student to a researcher, Timur offered the first glimpse of what lay on the other side. I was initially hired to help Timur analyze the unique dataset that became the basis of our collaboration, and over time our shared intellectual curiosity in the economic development of the Middle East allowed our partnership to grow. I sincerely appreciate the opportunity Timur offered me and hope that our collaboration can continue into the future.

In many ways, I owe my greatest debt of gratitude to my advisor, Charles Becker. He was one of the first professors I met when I started at Duke University. Our shared academic interests and overlapping sense of humor provided a vital lifeline when I first found myself facing a research deadline with little understanding of how to proceed.

In many ways, this dissertation is a testament to Charlie's caring guidance, earnest support, and emphatic insistence that I actually sit down and write! Charlie, I could not have done this without you, and I am incredibly grateful for your guidance and friendship.

Duke University's Department of Economics has been a delightful and generous institution during my five years here. I am especially thankful for their five years of financial support. Furthermore, I would like to thank the Robert K. Steel family for their generosity in funding my research over the summer of 2010.

While a dissertation can be a lonely project, I would like to thank all my friends for standing by me when I was absent, encouraging me when I was working, and celebrating with me at every milestone. I am a lucky man for having friends that stretch as far back as elementary school, freshman year in Tufts' Wren Hall, and fly-out day here at Duke. I know that wherever we all go, we will always be able to support one another with love and friendship.

Finally, I would like to thank my wonderful family. Without their loving support and reassurance that graduate school cannot *possibly* go on forever, the past five years would have been impossible. Nana and Grandpa, as I completed my dissertation and looked out to find where I am going next, I was struck that my inner voice echoed your advice, reassurance, and love. If every person contains within themselves the seeds of their forebears, I know I am well prepared for whatever comes my way, and I am so thankful for your help imparting that spirit in me.

Grandma, I join you in wishing that Grandpa were here for this day. In a way, my research echoes you two very closely. My interest in Ottoman history stems from our family ties, and in some ways economic research is simply the mathematical desire to see something complicated and try to take it apart, not unlike the projects I shared with Grandpa as a child. Thank you for everything.

Sevan, surviving a loved one in graduate school is never easy. A couple that

stays together when separated by space, graduation date, and academic discipline is harder still! As unknowable as they are right now, I believe that our next years together will be better than ever and I cannot wait to start them with you.

Alissa, even as I finish up what is undoubtedly the largest project of my life so far, I am in awe of your good heart, caring nature, and amazing accomplishments. I am so proud of you and lucky to have you as both a role model and a little sister.

Mom and Dad, your support for my decision to go to graduate school lasted longer than my initial resolve, but your love and enthusiasm recharged me and got me through. I know that I owe you so much for your guidance and support. I know how fortunate I am to say that I admire and love you two as individuals, a couple, my parents, and my friends. Thank you most of all. I love you.

1

Introduction

This dissertation is a collection of three chapters pertaining to institutional economics. The field of institutional economics is an outgrowth of public economics, in the sense that in many cases the key institutions that frame economic decisionmaking are the product of public policy. However, this is not exclusive. Institutional economics' key contribution is the acknowledgement that cultural and social institutions — often developed organically over the course of centuries — can play as significant a role in individuals' economic choices as governmental policy.

In the second chapter of this dissertation, which was co-authored with Professor Timur Kuran of Duke University, we examine the institutional framework and decisions of the 17th century Ottoman judicial system to show that Islamic courts of the Ottoman Empire exhibited biases that would have limited the expansion of economic exchange in the Eastern Mediterranean, particularly those between Muslims and non-Muslims. It thus explains why economic modernization in the Middle East required the establishment of secular courts. In quantifying Ottoman judicial biases, the paper also discredits the view that these courts treated Christians and Jews fairly as well as the counter-view that they ruled against non-Muslims

disproportionately. Biases against non-Muslims were in fact institutionalized. By the same token, non-Muslims did relatively well in adjudicated interfaith disputes, because they settled most conflicts out of court in anticipation of judicial biases. Islamic courts also appear to have exhibited biases in favor of state officials. The paper thus refutes the Islamist claim that reinstituting Islamic law (*sharia*) would be economically beneficial.

In the third and fourth chapters, the dissertation focuses on defined-contribution pension systems. While the data employed come from the Chilean pension system, the results developed in this paper can be applied to a wide variety of upper-middle and middle-income countries that employ or plan to employ a defined-contribution public retirement system, such as Mexico, Peru, Kazakhstan, and Azerbaijan. Both chapters focus on the risk of non-compliance, a fundamental yet under-studied risk to defined-contribution programs that can jeopardize the primary public-sector benefit of a defined-contribution program — the risk of unfunded liabilities.

Many countries that adopt a defined-contribution retirement program do so under duress. Financial challenges caused by rising old-age dependency rates — when fewer working-age citizens are able to fund the retirements of pensioners — and often-times economic turmoil or uncontrolled pension generosity threaten the viability of pay-as-you-go pension schemes. Governments overcome this systemic risk by transitioning to a fully-funded defined-contribution pension plan. Ideally, individuals will now save for their own retirement by complying with the mandatory pension contribution requirement. However, when individuals fail to comply with this requirement, there is the risk that when they reach retirement, these people will be unable to afford their retirement and seek the defined-benefit pension that many governments provide to the indigent.

When individuals who failed to comply with the mandatory saving requirement return to a defined-benefit pension program for the poor, the government is once

again faced with the problem of unfunded liabilities. Hence, individuals who do not comply with their savings requirements can become a free-rider and undermine the financial stability that defined-contribution retirement programs seek to establish. While this behavior can have potentially far-reaching consequences, there has been virtually no systematic analysis of what causes an individual to deviate from his savings requirements, or of the seriousness of the non-compliance problem.

While the risk of noncompliance exists in every economy with a defined-contribution retirement program, Chapters 3 and 4 focus on the institutional and regulatory environment in Chile, which is the flagship retirement program of this kind in among middle-income countries. Both chapters focus on the issue of competing institutions. Competing institutions often arise when policymakers require citizens to pursue two governmental goals that draw upon a shared pool of resources and thus force individuals to choose between two mandatory and distinct economic options.

In the third chapter, I explore the interaction between two of Chile's main social protection institutions — its defined-contribution pension program and its public insurance program. Although all adults are required to purchase health insurance and contribute 10 percent of their incomes to their privately administered pension funds by law, the imperfect monitoring characteristic of middle-income countries allows individuals to shirk these requirements. After formulating a model that posits whether health insurance and retirement savings are substitutes or complements, I employ Chile's under-utilized EPS dataset to estimate the effects of insurance status and health shocks on time-to-noncompliance and find that Chile's public health insurance fails to fully insure individuals against health shocks but lengthens individuals' time-to-noncompliance — indicating that health insurance is a complement to retirement savings in good health but a substitute otherwise. Finally, I analyze the policy implications of the results to assess the impact of linking health insurance to pension contributions.

In the fourth chapter, I explore the interaction between Chile's pension program and homeownership. Chilean society has a very high rate of homeownership, a direct effect of Santiago's consistent efforts to subsidize mortgages and encourage the development of Chile's housing stock. However, due to Chile's mortgage regulations, individuals must save up a substantial down payment when purchase a home. This down payment, coupled with individuals' varying access to credit, can often be impossible to accumulate while complying with the government-mandated pension savings program. Thus individuals suffer a large incentive to fail to comply with the mandatory pension contribution requirement. In this chapter, I find that individuals who are unable to utilize government housing subsidies and who are looking to purchase a house in the next three years are significantly more likely to fail to contribute to their pension program — most likely in order to divert their pension savings towards purchasing a new home. Ultimately, I find that this tradeoff between investment in capital-intensive durable goods and compliance with Chile's pension requirement is most extreme at the middle of the income spectrum, as poorer individuals have access to government subsidies and wealthier individuals can afford to save for a home while sequestering 10 percent of their income in illiquid pension savings.

The main lesson for policymakers in Chapters 3 and 4 is that forward-looking and otherwise law-abiding individuals may fail to contribute to their pensions due to competing institutional opportunities. Due to the illiquidity of pension savings, short- and medium-term expenditures can occasionally divert a large enough share of household income that individuals are forced to choose between contemporaneous consumption and long-term savings. We must note that these competing expenditures often fulfill a social purpose: The very same policymakers who encourage long-term savings similarly require all individuals to purchase health insurance and incentivize homeownership. When there is a limited capacity to punish those who

fail to comply with the mandatory contribution requirement, constrained households occasionally will choose other investment and insurance vehicles over contributing to their pensions.

In all the chapters that follow, the unifying theme is one of social institutions causing unintended consequences. As an example that complements our work, Kuran (2011a) charts how commercial institutions and Islamic jurisprudence developed over hundreds of years. While initially mutually supportive in the early-modern era, Islamic jurisprudence failed to keep up with the development of alternative methods of commercial and economic organization during the Ottoman Empire and eventually led to regional economic stagnation. Whereas Islamic jurists who shaped the legal institutions of the Ottoman Empire failed to keep up with the development of modern commercial institutions, a similar framework of competing institutions aptly describes the conflict between encouraging individuals to save for retirement and spend their income on health insurance and durable goods.

In every instance, those who manage the development of these institutions — be they religious scholars, jurists, businessmen, or governmental policymakers — are forward-looking individuals tasked with maintaining social stability through the provision of justice, reduction of risk, and encouragement of responsible economic planning. Yet, in every case, the *interactions* of these myriad institutions create unforeseen consequences that frequently undermine the very objectives that the individual institutions were crafted to support. It is with this unfortunate paradox in mind that we can begin to learn the generalizable lessons that stem from the unique institutions examined in the following chapters.

2

Structural Inefficiencies of Islamic Courts: Ottoman Justice and Its Implications for Modern Economic Life

2.1 Introduction

Underdeveloped countries are often advised to improve their judicial systems in order to strengthen contract enforcement and increase gains from exchange. Specifically, they are urged to institute laws enforced independently from the executive and legislative branches of government and impartially across society, without regard to such personal traits as sex, ethnicity, and religion.[1] The advice typically takes for granted that laws will be secular, in the sense that they will not draw legitimacy from any particular religion.

Yet, in predominantly Muslim countries, there is a demand that legal reforms should draw on Islamic legal traditions stretching back more than a millennium. Sometimes it takes the form of a call for reinstituting Islamic law, or the sharia. Certain elements of Islamic law have been studied for their economic implications.

[1] For a few examples, see Heckman et al. (2004) and Dam (2006).

They include the prohibition of interest, lack of legal personhood, the Quranic rules of inheritance, and penalties for apostasy.[2] What has not been analyzed, at least not rigorously, is the Islamic system of commercial adjudication. Did it satisfy the objectives of judicial impartiality and independence? This paper is to seeks answers with reference to the Ottoman justice system in the seventeenth century.

There are four reasons for this choice. The first has to do with data availability. Although Islamic law was long the law of the land over a huge land mass, nowhere were Islamic court proceedings recorded as thoroughly, and the records preserved as fully, as in the Ottoman Empire. We have constructed a data base of Ottoman commercial trials spanning the entire seventeenth century. By far the largest such data base, it consists of cases adjudicated in Istanbul, capital of the Ottoman Empire and commercial hub of the Eastern Mediterranean. Second, the seventeenth century constitutes the latest period when Ottoman legal practices were essentially free of Western influences.[3] Third, at the time Istanbul was a highly cosmopolitan city with huge Christian and Jewish minorities and also an expanding foreign presence. Given the importance of cross-communal contacts in today's increasingly integrated global economy, it is of interest to examine the operation of Islamic courts in a setting involving inter-faith business ties.

Finally, given that Istanbul was the seat of imperial power, its courts would have operated most closely according to the ideals of justice and impartiality that Ottoman sultans claimed to pursue. Like other Muslim dynasties, Ottoman sultans promised to punish state officials who mistreated their subjects; they also promised that all subjects, regardless of faith, would get fair hearings in imperial courts. Keeping Istanbul's judges essentially honest promoted political stability. Unlike judges in

[2] On apostasy, see Saeed and Saeed (2004); on the interest ban, El-Gamal (2006); and on inheritance and legal personhood, Kuran (2011a).

[3] By the eighteenth century non-Muslim Ottoman merchants were beginning to do business under European legal systems (Kuran, 2004; Masters, 2001).

places remote from the capital, those posted in Istanbul also received compensation regularly, which counteracted the temptation to engage in corruption. Thus, the court records of seventeenth-century Istanbul allow us to study Islamic law in a setting where it would be expected to perform as closely as possible to Islamic ideals.

Critics of the Islamic system of justice, from contemporaneous observers of the Ottoman courts to modern legal scholars, have held that as a matter of practice Islamic justice has been unpredictable, biased in favor of state officials, and rigged against non-Muslims. The pro-state biases of the Islamic courts stem, say critics, from their subordination to the sultan. Indeed, Ottoman judges served as personal representatives of the sultan, who had the duty to deliver justice. As for the courts' pro-Muslim biases, they were rooted partly in the ingroup biases of Muslim judges. Also relevant, however, were procedures that treated Muslim testimony as inherently more credible than non-Muslim testimony. Insofar as they existed, the biases in question could not be countered through formal judicial review. The rulings of a court could be reversed only through a personal appeal to the sultan, which for most litigants was not a realistic possibility.

To one degree or another, pre-modern courts openly discriminated against outsiders everywhere. In the absence of equal-rights norms that are central to modern judiciaries, they favored local interests without apology. The Islamic courts of the Ottoman Empire provide no exception. In barring non-Muslims from testifying as a witness against Muslims, they followed what was once a universal pattern. This procedural discrimination lends credibility, of course, to highly critical western accounts of these courts (Porter, 1771; Masters, 2001). Legal reformers of the nineteenth century not only agreed with these criticisms but they considered the biases in question as harmful to Middle Eastern economic development. Revealingly, the modern commercial courts that they established had a secular character.

For all this history of criticism, distinguished Ottoman historians who are familiar

with the historical records report that, although these courts were prone to corruption, they were not noticeably biased against local Christians and Jews, or European foreigners, or any other group (Ekinci, 2004). Lacking evidence of bias, they infer that Ottoman judges treated all groups fairly. This is puzzling. If the evidence-generating procedures of the Islamic courts were stacked in favor of Muslims, how could their judgments have been unbiased? Conversely, if the courts were unbiased against non-Muslims, why did European observers and Ottoman reformers find them blatantly unfair? It could be that the European claims reflect hostility to Islam. Yet, certain critics of Islamic courts heaped praise on other Ottoman institutions, which raises the question of why they were negative in this particular context.

Our unique data set provides an opportunity to reconcile the seemingly contradictory accounts of Ottoman justice. We start with a description of the Ottoman judicial system. Theoretical and empirical insights from the law and economics literature follow. Subsequent sections of the paper address, in turn, the various biases that afflicted Ottoman trials. We conclude with implications for modern attempts to revive Islamic legal institutions. The inefficiencies of the Islamic courts of seventeenth-century Istanbul were not aberrations, we suggest. Rather, they stemmed from structural features of the Islamic legal system that modern promoters of the sharia have not ruled as irrelevant to modern life.

2.2 Legal Marketplace in Seventeenth-Century Istanbul

In the seventeenth century, the legal system of the Ottoman Empire was based on Islamic law. Although it had its own particularities, it closely resembled the legal systems of previous and contemporaneous Muslim-governed states. Accounts of Mamluk courts in fourteenth- century Cairo and of Abbasid courts in tenth-century Baghdad are strikingly similar to those of seventeenth-century Istanbul. In terms of organization, procedures, and principles of justice, the Islamic courts of the Ottoman

Empire did not depart significantly from other Islamic courts.

Every Ottoman court was headed by a judge (*kadi*) who performed, in addition to several executive functions, two distinct judicial functions. On the one hand, he registered, and in so doing authenticated, contracts, settlements, and transactions. A registered contract could be consulted should it become necessary to forestall or resolve a dispute. On the other hand, the judge conducted trials to resolve disputes brought before him.[4] A dispute could involve a criminal matter or what we would now characterize as a civil matter. In either case, the judge would hear the plaintiff, give the defendant a chance to respond, if necessary conduct an investigation of his own, and pronounce a verdict. Occasionally he would postpone a verdict to allow a litigant to bring evidence. A verdict might involve an order to fulfill a contractual term or pay damages. The records contain disputes involving debt, divorce and custody, estate settlement, guardianship, sale, property transfer, mortgage, pawning, tax payment, guild administration, communal rights, and neighborhood norms. The burden of proof did not differ according to the type of case, and neither did procedures.

Each judge had scribes record accounts of his activities in a "register of cases" (*sicil defteri*), and during his tenure at any one court he might use more than one register. In small towns, judges had scribes record all their court's business more or less chronologically in a single notebook, moving to a new notebook when the first was full. In major cities, the norm was to use a separate register for estate inventories and perhaps another for official directives. All other records ended up together, sometimes with certain government orders in the back, in "regular" registers.[5]

When the tenure of a judge ended, his registers came to an end as well; his

[4] The executive functions included enforcing public morals and maintaining a record of official orders sent to his area. Ortaylı (1994), Imber (2002), and Gaudefroy-Demombynes (1950) discusses the functions of judges.

[5] Prior to the nineteenth century, judges had discretion on recording and categorizing. See Mandaville (1966), pp. 313-14.

successor started one or more new registers. The departing judge generally handed over his registers to his successor, or, in places with a court building, simply left them behind for storage (Faroqhi, 1997). The notebooks used as court registers vary greatly in size. A judge who opted for a thick notebook with huge pages might fit the entire record of his tenure into one book, especially if his tenure was short. If his successor started a skinny notebook, he might go through several notebooks. Many old registers must have been discarded eventually; others perished in fires, earthquakes, floods, and wars; still others must have been destroyed deliberately by individuals with something to hide.

A judge's time in any one place was limited to prevent him from developing local political ties. It could last as short as three months, but the norm was about a year, and rarely did a judge serve more than 20 months in any one post. The judges of courts located in politically sensitive places tended to be rotated more frequently than those in strategically unimportant towns, which is consistent with the political considerations that guided appointments (Ortaylı, 1994). The reassignment probability of a judge depended on his reputation. This provided incentives to minimize complaints by adjudicating consistently and fairly. Complaints about a judge posted in Istanbul would reach the Sultan more easily than those concerning judges posted in courts located far away from the capital. For that reason, too, judicial corruption would have been less common in Istanbul than elsewhere, affording an opportunity to study Islamic adjudication in a setting where it had the best chances of approaching the ideals.

For their services some judges received a salary; all were also authorized to charge litigants fees (*harç*) set by law. Their fees were usually proportional. In a commercial dispute they might collect, for instance, 2 percent of the amount at stake (Ortaylı, 1994; Bayındır, 1986; Gaudefroy-Demombynes, 1950). There appears to have been no set standard concerning the payee. The plaintiff and defendant might be expected to

share the cost. In the absence of documentation on the fee structure of the Ottoman judicial system, we do not know whether charges differed by the substance of the dispute or the amount at stake.[6] We know that the winner of a lawsuit had to pay, at a minimum, a fixed fee for a document certifying the outcome, known as a hujjet (*hüccet*).

Although judges were assigned to a jurisdiction, such as the town of Amasya or the Eyp neighborhood of Istanbul, Ottoman subjects and visitors were not required to use the court located where they lived or worked. They were free to take a dispute to a judge of their choice. In practice, then, the judges of Islamic courts were in competition for legal business.

The vast majority of the Ottoman subjects adhered to one of the three monotheistic religions. Muslims, who formed the largest group, were required to live by Islamic law. This meant that to register a contract legally, or to get a dispute adjudicated formally, they had to use an Islamic court. For their part, non-Muslims enjoyed "choice of law": though entitled to use an Islamic court, on civil matters they were free to use a court of their own choice.[7] Thus, a Greek Christian could have a debt dispute with a co-religionist litigated before an official of the Greek Orthodox Church. All litigation involving both Muslims and non-Muslims had to be handled by a Muslim judge, because of the rule that Muslims had to live by Islamic law (Kuran, 2004). This system of asymmetric legal pluralism meant that, at least in cases among non-Muslims, Muslim judges competed also with Christian and Jewish courts. Although it is certain that non-Muslims used courts of their own, they left no records, as far as is known. This is undoubtedly because in trying to minimize their

[6] It appears that as a matter of practice substantial variations existed among courts. Records of the Ottoman palace are replete with complaints about judges whose charges exceeded the authorized amount (Uzunçarşılı, 1965).

[7] All criminal matters, regardless of the identities of the accused and the victims, fell under the responsibility of Muslim officials.

tax obligations Christian and Jewish communities sought to withhold information about their financial matters from state officials.

Under Islamic law, the responsibility to deliver justice belonged to the sovereign- in the Ottoman case, the Sultan. He was free to litigate any dispute himself, and in principle anyone could take a case directly to him. In practice, he let his appointed judges try the vast majority of the lawsuits brought to an Islamic court. These judges differed in status and responsibility. Two chief judges (*kazasker*), one for Ottoman provinces in Europe and the other for the rest, handled appointments on behalf of the Sultan. Moreover, the judges of politically strategic places such as Istanbul and Cairo, like those posted in the holy cities of Mecca and Medina, ranked above the rest. The salaries of judges were tied to rank. Presumably higher ranked judges also earned more in fees by virtue of being posted to courts with exceptionally prosperous litigants.

In principle, high-ranking judges did not have greater legal authority than the rest. The youngest judge on his first assignment in a sleepy town had as much authority to deliver a verdict as a chief judge. His verdict was final, and from a doctrinal standpoint it carried as much authority as a verdict delivered by an experienced judge. Under Islamic law there exists no standardized appeals process (Shapiro, 1981). Accordingly, an Ottoman disputant could overturn an unfavorable decision only by appealing directly to the sultan. The appeals system was thus biased in favor of elites with access to the sultan's palace. For most Ottoman subjects, appealing a court decision was not a realistic option; and it was particularly costly for the residents of places located far from the sultan's palace in Istanbul. None of this implies that Ottoman judges were free to rule whimsically. As we shall see, they were subject to constraints.

2.3 Sources of Judicial Bias

The social sciences do not provide a general theory of judicial fairness. Theories of dispute resolution developed within the law and economics tradition show, disappointingly, that one cannot measure the fairness of courts within a homogenous pool of potential litigants. The key problem, as shown in a literature launched by Priest and Klein (1984) and developed by Shavell (1996), lies in heterogeneity among litigants in both information and reputational costs. Because of such heterogeneities one cannot recreate, from the subset of actually adjudicated disputes, the underlying set of potential disputes that might have gone to trial. As van Tulder and van Velthoven (2003) put it, the cases that reach trial represent only the tip of the iceberg of civil disputes.[8] Because forward-looking and utility-maximizing potential litigants may choose to settle rather than appear in court, one cannot even assess the social optimality of observed litigation and plaintiff victory rates. Questions over whether the lawsuits in a particular society contain too many frivolous cases, are socially destabilizing, and are too costly, also pose serious theoretical difficulties.

However, the social sciences have much to contribute to the analysis of judicial decision making involving disparate groups within a population. Three distinct literatures are relevant to the identification of intergroup differences in the application of justice. They involve competition among legal jurisdictions, judicial independence, and in-group bias.

2.3.1 Fee-for-Service Adjudication and Competing Legal Jurisdictions

In modern courts judges are essentially indifferent to the number of cases that come before them, because it affects neither their promotion prospects nor their compensation. Before the modern era, however, the success of judges did depend on how

[8] Other important contributions include Kessler et al. (1996), Siegelman and Donohue (1995), and Siegelman and Waldfogel (1999).

many cases they adjudicated, because they derived income at least partly through fees collected from litigants. The fee-for-service system fostered inter-court competition when multiple courts were in close proximity. This claim has been tested through a comparison of England's courts before and after the English legal reforms in the early nineteenth century. Klerman (2007) finds that under the fee-for-service compensation regime of the pre-reform period judges were more likely to rule for plaintiffs than under the salary-based compensation regime that followed.[9] He reasons that since plaintiffs decide whether to sue and also the adjudication forum, profit-maximizing judges would have tilted their verdicts in favor of plaintiffs.

This finding is obviously relevant to Ottoman courts. Although Ottoman judges received a salary from the sultan, they also collected fees from litigants. Moreover, individual plaintiffs were able to seek out judges known for their propensity to rule in favor of the plaintiff. One would expect judges to have ruled for plaintiffs more frequently than they would in the absence of a choice among courts. They would have exhibited a pro-plaintiff bias, regardless of the religious identities of the litigants.

Following Landes and Posner (1979), Klerman asks what might limit inter-jurisdictional competition from unraveling into a corner solution such that plaintiffs always win. In pre-modern England, he finds, in addition to the Chancery and the Parliament, the monarch's ability to appoint and remove judges limited the pro-plaintiff bias of courts. Biased judges reduced the legitimacy of the monarch's rule. His oversight of the legal marketplace thus constrained judges' ability to compete with each other by tilting decisions in favor of plaintiffs.[10]

[9] Building on Klerman, DiIanni (2010) adds that jurisdictional competition between courts will not necessarily generate pro-plaintiff bias if both the plaintiff and the defendant must agree on the adjudication forum.

[10] Court fees, too, limit pro-plaintiff bias. If the initiation of adjudication guaranteed the defendant's paying restitution to the plaintiff, the defendant would prefer to settle out of court in order to escape adjudication fees. The plaintiff, too, would prefer to settle, for the opportunity to bargain with the defendant over the distribution of what would have been the judge's fees. Hence, court fees, along with the ability to settle out of court, will make judges cap their pro-plaintiff bias.

As in pre-modern England, in the Ottoman Empire any pro-plaintiff bias of the courts would have been known to potential litigants. Moreover, the sultan's executive oversight would have guarded against the excesses of judicial rent-seeking. This oversight brings us to another potential source of bias in adjudication: the state's influence on judicial decisions. A sultan able to limit the pro-plaintiff bias of judges might have been able to tilt their verdicts in favor of the state in cases that affected him directly.

2.3.2 Judicial Independence

Judicial independence entails, on the one hand, the capacity to exercise judicial review and, on the other, counter-political judicial continuity. Judicial review gives courts the right to overrule executive decisions, challenge the legitimacy of the government and, under extreme circumstances, even to depose a ruler. The review process may result, of course, in the legitimization of government policies. Courts may support government policies and facilitate their enforcement (Feld and Voigt, 2003; Ramseyer and Rasmusen, 2003).

In the course of Islamic history, the judiciary's ability to challenge the sovereign's authority has waxed and waned. Initially weak due to innumerable legal controversies, its power grew during the eighth and ninth centuries as schools of Islamic jurisprudence got established and legal traditions took hold. During this period, the judiciary challenged the authority of the sovereign on occasion by attacking legal innovations as deviations from the Quran. However, well before the establishment of the Ottoman state in 1299, sultans gained effective control over the judiciary. They then solidified this control by standardizing the code of law applied in their realms, assuming sole authority over the appointment and dismissal of judges, placing religious and judicial officials on the state payroll, and binding the judiciary's wellbeing to the support of the state (Coşgel et al., 2009; Imber, 2002). The cowing of the

previously independent judiciary removed the threat of judicial review and bolstered the ability of the courts to legitimize the prevailing regime. In assuming control of the judiciary, the sultan incentivized judges to enforce imperial laws and the chief judge (*şeyhülislam*) to support the sultan whenever consulted on the legality of a decree.[11]

Counter-political continuity exists when the judicial system allows judges to remain in office following a political change. The Ottoman sultan's policy of rotating and replacing judges regularly, which was meant to decrease local loyalties and reduce corruption, limited counter-political continuity. In keeping the terms of judges short, this policy prevented the judiciary from establishing patterns that could outlive the reign of a sultan.

By the seventeenth century, then, judicial independence was essentially lacking in the Ottoman Empire. Hence, one would expect Ottoman subjects to have shown extreme caution in challenging the state in court.

2.3.3 Ingroup Bias in Judicial Decisionmaking

All courts are prone to ingroup bias, which is the tendency to give preferential treatment to people perceived as belonging to one's own group. In modern jurisprudence, ingroup bias is a recognized phenomenon that certain institutions are meant to counteract. A formal system of appeals limits a judge's ability to exercise favoritism, because having decisions overturned by a superior court exacts a reputational toll. Likewise, a norm of equal protection under the law makes judicial decision makers strive consciously to consider factors favorable to members of outgroups.

Such institutions may alleviate ingroup bias but not eliminate it. Numerous

[11] Judges may acquire greater judicial independence when they anticipate the replacement of the reigning sultan. They may begin to enforce the preferences of his successor in a display of proactive loyalty to the new sultan. For the underlying logic, see Helmke (2002, 2005) for her work on judicial independence during regime transition.

studies indicate that even in liberal societies that promote the principle of equal protection under the law judges regularly fall prey to ingroup bias. Gazal-Ayal and Sulitzeanu-Kenan (2010) and Shayo and Zussman (2010) demonstrate that Israeli judges presiding over apolitical criminal and low-stakes civil hearings exhibit persistent ingroup bias. Shayo and Zussman show also that the prevalence of ingroup bias is correlated with security-related events that heighten political tensions.

There are also modern institutions that foster ingroup bias. The jury system, whereby evidence is evaluated by the defendant's peers, promotes ingroup bias favorable to defendants facing prosecution by the state. By the same token, it can lead to ingroup bias in trials that pit an insider against an outsider. In a study of international patent enforcement in American courts, Moore (2003) finds that jury trials are more likely to exhibit xenophobic bias than trials decided from the bench.

Islamic jurisprudence requires all lawsuits to be adjudicated before a judge. This reliance on bench trials would have diminished ingroup bias relative to modern jury trials. However, in the absence of clear procedures for appealing a verdict, judges lacked professional incentives to take precautions against ingroup bias (Shapiro, 1981). Nor were judges trained to abide by a norm of equal protection in evaluating evidence. On the contrary, they learned to give greater weight to the testimony of Muslims than to those of non-Muslims. This legal tradition was built into the adjudication procedures of Ottoman courts. The transcripts of seventeenth-century trials are replete with references to seeking out Muslims specifically for opinions regarding a dispute at hand.[12]

Given that institutional pressures to counteract intentional or unintentional ingroup bias were lacking, the judges of Islamic courts were very likely to favor their co-religionists in inter-religious lawsuits. Since all judges were Muslim by design,

[12] See, for example, cases Istanbul 9 (1661) 171b/2 and 189a/3 , Istanbul 22 (1695) 105a/2, and Istanbul 23 (1696) 20a/1, all recorded in Kuran (2011b).

non-Muslims would have been at a disadvantage in lawsuits pitting them against Muslims.

2.4 The Courts of Galata and Istanbul

To test these hypotheses, we turn to the largest existing data set of transliterated and translated Ottoman court records. It includes 10,080 cases from 15 registers in the courts of Galata and central Istanbul (hereafter, simply Istanbul).[13] The registers were selected to provide a more or less uniform distribution over the entire seventeenth century.[14] In each dated account, the scribe would record the identity of the litigants—always including their religious affiliation and title, and often also their neighborhood of residence—the nature of the dispute, the evidence brought to the trial, and the verdict. These two courts were the most prominent of the 16 courts serving the Ottoman capital, which at the time had around 700,000 inhabitants. The Galata court was located near the empire's main port, and the Istanbul court near the fabled Grand Bazaar. Precisely because of their proximity to Istanbul's main commercial centers, the caseloads of these two courts consisted primarily of commercial registrations and trials. Only the trials are of interest here.

As Table 2.1 shows, these records contain 2,291 commercial trials. Its top 15 rows refer to the court registers in the data set. The register number is that assigned by the Turkish archive where the registers have been housed since 1894. Thus "Galata

[13] It is reproduced in Kuran (2011b), with full transliterations of the records in the Latin alphabet of modern Turkish, along with detailed summaries in both English and Turkish. The records are in Ottoman Turkish, which few Turkish speakers now understand, due to the high number of Arabic and Farsi loanwords. Scribes kept the records, for the most part, in an Arabic script known as broken divani (*kırma divani).*

[14] The registers were chosen to provide coverage across the seventeenth century among surviving unspecialized registers; thus, registers reserved for estate inventories or official correspondence were excluded. Over certain time spans, the registers tended to be small; to be able to check for repeat use of a given court over the periods in question, some consecutive registers were included. Big gaps exist in the series due to fires. The main motivation for constructing the sample was to see whether the standards and processes of Ottoman courts, and the economic life that they supported, changed over the course of the seventeenth century.

130" refers to the 130th register in the Galata series within the archive.

Dividing our trials by court, we find that 60.1 percent belong to Istanbul. As shown in Tables 2.2 and 2.3, the two subsamples differ in terms of the demographic composition of the litigants, and also the distribution of cases by topic.[15] Disproportionately few defendants were Muslim in Galata relative to Istanbul. Galata had a higher share of tax cases than Istanbul, but a lower share of cases concerning a *waqf* (Islamic trust). Without controlling for the differences in question, we might conclude, erroneously, that the two courts judged cases according to different standards and differed in their biases. In fact, multivariate analysis to be presented later shows that the court in which a dispute was resolved mattered significantly to the outcome.[16]

Since the vast majority of the litigants enter the records as Muslim, Christian, or Jewish, the lower shares of Muslims among the Galata litigants points to a higher share of non-Muslims. In the seventeenth century Galata had an unusually high concentration of non-Muslim residents, especially Greek and Armenian Christians.[17] That could explain why the Galata court's clientele was more heavily Christian. Yet, in contrast to a modern legal system, plaintiffs in seventeenth-century Istanbul did not need to sue in the district where the dispute arose or where they lived. They could choose among multiple courts within walking distance of one another. Christians might well have favored the Galata court because of greater convenience. Another reason could be that the Galata court had a reputation for ruling in favor of Christians more often than the competing court in Istanbul.

[15] Cases that present arguments involving more than one topic, such as a dispute over who should inherit a rental property, are counted more than once. There were also topics not listed here. For both reasons the rows of Table 2.3 do not sum to 100.

[16] The pattern could reflect year-specific fixed effects. The years of the Galata registers do not overlap with those of the Istanbul registers.

[17] Population estimates for seventeenth-century Istanbul do not distinguish between neighborhoods, so we are unable to quantify the demographic differences between central Istanbul and Galata.

2.5 Pro-Plaintiff Biases

Given that two courts competed with one another for lawsuits, their judges were all incentivized to tilt verdicts in favor of plaintiffs as a means of attracting cases. With all judges playing this game, the pro-plaintiff bias would have grown, and in the limit plaintiffs would have won all cases in every court. The limit would never be reached, of course, if potential defendants started refusing to appear before the most biased of judges. By the same token, insofar as judges were predisposed to compete for cases by favoring plaintiffs, the temptation to launch frivolous lawsuits would have increased, especially if judges made side deals with plaintiffs to share whatever could be extracted from hapless defendants.

In fact, the verdicts delivered in our two courts do appear to have favored plaintiffs. They won 59.7 percent of all cases submitted to adjudication. As already noted, it is theoretically impossible to assign a benchmark of how a court insulated from inter-jurisdictional competition would perform. Hence, one cannot say whether this plaintiff victory rate deviates from the social optimum. Yet the number is clearly below the 100 percent figure we would have found if judges decided cases with a single goal, namely, encouraging people to bring them lawsuits. Other considerations must have been in play. Fairness is, of course, the most obvious competing consideration. The sultan's executive oversight would have ensured that it played a significant role. For the sake of maintaining political stability, the sultan would have limited the profit-seeking tendencies of even the most corrupted judge.

Table 2.4 gives the breakdown of the trials in terms of the nine possible pairings among our three religious communities, and Table 2.5 the corresponding plaintiff victory rates.[18] We see that irrespective of the pairing the plaintiff wins more frequently than the defendant. The plaintiff victory rate is significantly greater than 50

[18] The figures exclude the small number of trials involving gypsies, foreigners, and new converts to Islam (*mühtedi*).

percent in intra-Muslim, intra-Christian, and intra-Jewish cases, but also when the litigants are from different faiths.[19]

Religion itself was not irrelevant to the probability of victory. Focusing on the Muslim-Christian cases, where the numbers are high enough for meaningful statistical analysis, we find that the plaintiff victory rate differs depending on the side of the Christian. In view of the debates concerning abuses against Christians, the nature of the difference may come as a surprise. The plaintiff victory rate is higher, not lower, when a Christian sues a Muslim than when the roles are reversed.[20] Equally surprising, it may seem, is that Christian plaintiffs do better when the defendant is Muslim than when he is a co-religionist (71.9% vs. 59.4%). For Muslim plaintiffs, by contrast, the defendant's religion does not matter.[21] Not even the Ottomanists who reject the strident charges of anti-Christian bias in Islamic courts would have predicted these findings. Based on casual observations, they have suggested that the Islamic courts treated Christians fairly, not that Greek and Armenian Ottoman subjects benefited from judicial discrimination in their favor.

Counter-intuitive as these findings may seem, they will not surprise students of modern litigation involving foreigners. In American courts xenophobia in adjudication goes hand in hand with a high foreign victory rate in foreign-initiated lawsuits

[19] Testing the intra-faith results in Table 3 against the null hypothesis that the plaintiff victory rate equals 50 percent, we reject the null for Muslims and Christians at the 99.9% level of statistical significance and for Jews at the 90% level (p-values: 0.00 for Muslims, 0.00 for Christians, 0.07 for Jews). We reject the null also in inter-faith cases involving Christians and Muslims (p-values: 0.00 for Christian suing a Muslim, 0.02 for Muslim suing a Christian), but cannot do so consistently in cases involving Jews or "others," in all likelihood due to their small numbers in the data.

[20] It is relatively higher in the former case at the 95% level of statistical significance (t=2.52).

[21] For Christian plaintiffs the victory rate is greater when the defendant is a Muslim at the 95% level of statistical significance (t=2.37). For Muslim plaintiffs, the null hypothesis that the rates are equivalent cannot be rejected (t=0.39). For the sake of completeness, we may compare across plaintiffs, holding the defendant's religion constant. When the defendant is Muslim, it matters whether the plaintiff is Muslim or Christian at the 99% level of significance (t=2.69). When the defendant is Christian, the plaintiff's religion does not matter (t= 0.44). The standard error for all terms in Table 3b is 0.5.

against Americans. Studying patent disputes, Clermont and Eisenburg (1996) find that American courts rule in favor of foreign firms suing American firms at a higher rate than they do in favor of domestic American firms suing other American firms.[22] Yet foreigners consider American courts xenophobic and, hence, that in patent cases pitting a foreign firm against an American firm the foregn side will generally lose. The authors reconcile this perception with their finding by invoking selection bias. Expecting American juries to be biased against them, foreign firms will sue an American firm only if they consider their case very strong. Their high victory rate thus reflects not xenophilia but the strength of their lawsuits. This interpretation is supported by Bhattacharya et al. (2007), who study the impact of litigation on stock prices as a measure of expected judicial outcomes. They find that foreign firms sued in American courts are expected to do substantially worse than American firms sued in the same courts.[23] Our finding for seventeenth-century Istanbul mirrors these findings. Like foreign plaintiffs in America today, non-Muslim Ottoman litigants did better in court than Muslims precisely because they had good reason to fear discrimination.

One may wonder whether American courts differ at all from Ottoman courts of the seventeenth century in regard to anti-outsider biases. The institutional biases of the Ottoman courts have no parallel in American courts, where American and foreign litigants enjoy identical rights concerning witnesses, and where witnesses are

[22] In lawsuits between domestic firms, the plaintiff victory rate is 64 percent. By contrast, in those that involve a foreign plaintiff and a domestic defendant, it is as high as 80 percent. The difference is significant at the 99.9% level. In lawsuits by domestic firms against foreign firms, the plaintiff victory rate drops to 50 percent. This percentage differs statistically from 64 percent, again at the 99.9% level of significance.

[23] The authors employ a Heckman two-step regression process. In related work, Moore (2003) examines the number of international patent disputes while controlling for U.S. patents held by foreign and domestic firms, thus weighing the foreign litigation rate in U.S. courts against the number of foreign-held US patents, in order to examine how often international firms sue relative to American firms. She confirms that foreigners sue disproportionately less, given the distribution of patents held. This result is consistent with the inference that sample selection bias drives the high foreign victory rate in foreign-initiated lawsuits.

not limited in terms of creed or nationality. What has not disappeared is ingroup bias, which cannot be eliminated through legislation. Like juries everywhere throughout history, American juries are drawn from people predisposed to give the benefit of the doubt to people like themselves, in other words, Americans rather than foreigners. Likewise, judges are predisposed to seeing conflicts from the perspective of American litigants.[24]

2.6 Non-Muslim Participation in Islamic Courts

If selection bias accounts for the high victory rate of Christians in their lawsuits against Muslims, they would have had special reasons to avoid Islamic courts. To determine whether Christians underutilized Islamic courts, we need to know the shares of each religious group in Istanbul's population. Although no population census was conducted in the seventeenth century, Mantran (1962) estimates that at the time Istanbul was 58.8 percent Muslim, 34.8 percent Christian, and 6.4 percent Jewish—as represented in Figure 2.1.[25]

Given these shares, let us conduct a counter-factual exercise to determine how often the three groups would have faced one another in an Islamic court if, irrespective of faith, the residents of Istanbul interact in pairs randomly. For the purpose of this exercise, we will assume that for every type of pairing (Muslim-Muslim, Christian-Muslim, etc.), interactions result in litigation with the same probability, and that all litigation takes place in an Islamic court. Under this random interactions and

[24] Bodenhausen (1955) presents experiments concerning the psychological mechanisms through which juries interpret and internalize evidence presented in trials. Evidently individual jurors tend to retain evidence that confirms their preexisting stereotypes and to discount exculpatory evidence at odds with their stereotypes.

[25] In his monumental history of Istanbul Mantran (1962) provides, on the one hand, the population size of the Christian and Jewish communities in 1690, and on the other, the Muslim and non-Muslim shares in the sixteenth century. There is no reason to believe that population shares changed significantly between the sixteenth and seventeenth centuries. By holding the relative sizes of the Muslim and non-Muslim communities constant, we are able to extrapolate the three shares.

random litigation scenario, the distribution of pairwise litigation would be as shown in the first column of Table 2.6. The intra-faith cases equal the squared shares of the communities in the population.[26] The interfaith disputes are, of course, symmetric: the share of Christian-initiated lawsuits against Muslims equals that of Muslim-initiated cases against Christians.

The second column of Table 2.6 provides the observed adjudication shares in seventeenth-century Istanbul.[27] Clearly intra-Muslim cases are vastly overrepresented relative to random adjudication, as are intra-Christian cases. The reason is that pairings did not occur randomly; Muslims interacted disproportionately with Muslims, and Christians with other Christians, resulting in disproportionately many intra-faith disputes.[28] Interactions were more likely among co-religionists because people found it advantageous to deal with individuals known to them personally; and typically they knew more about a co-religionist than about a religious outsider. For Christians, another reason was that the two major Christian groups, Greeks and Armenians, favored Islamic courts over their own communal courts as a neutral ground for adjudication.

It is in inter-faith disputes that the selection bias in question is detectable. Whereas under random interactions the share of Muslim lawsuits against Christians would equal the share with the roles reversed, in fact the two shares were starkly different: 9.9 percent of all disputes consisted of a Muslim lawsuit against

[26] For example, the Christian intra-faith share is $(0.35)^2 = 12.3\%$. If Christians and Jews had an outside option, namely, adjudication in a court of their own religious community, their intra-faith shares would be lower than those given.

[27] These calculations exclude cases involving a state official, because they involved matters between individuals and the state, rather than between private parties. They also exclude the few cases involving others (gypsies, foreigners, and new converts to Islam).

[28] For each of the nine religious pairings, the hypothesized and observed shares differ at the 99.9% level of statistical significance (t=27.66 for MvM, t=7.15 for CvC, t=15.31 for JvJ, t=16.06 for MvC, t=40.21 for CvM, t=5.73 for MvJ, t=14.04 for CvJ, and t=5.77 for JvC).

a Christian, but only 3.7 percent involved a Christian lawsuit against a Muslim.[29] This asymmetry holds regardless of whether the cases were adjudicated in Galata or Istanbul.[30]

The asymmetry in question suggests that Christians considered the courts biased against them, at least in cases in which they faced a Muslim. Two sources of institutionalized bias have already been mentioned. First, the judges and assistant judges of Islamic courts were exclusively Muslim, as were the court-appointed professional witnesses (*şuhud ül-hal*) present at every adjudication or registration procedure. These officials would have been attuned to the customs, perspectives, and aspirations of their co-religionists. As such, even if they tried to be meticulously impartial, they would have been more receptive to arguments of Muslims than to those of non-Muslims. Hence, in Muslim-Christian cases, the benefit of any doubt would have gone to the former.

Second, Muslims and Christians did not have equal rights as regards testifying in court as a litigant-invited witness. Whereas a Muslim witness could testify against anyone, non-Muslims were allowed to testify only against other non-Muslims. Our own sample of cases shows that the ban on non-Muslim witness testimony against Muslims, which became institutionalized early in Islamic history, was obeyed in seventeenth-century Istanbul. As indicated in Table 2.7, a total of 1,177 witnesses were called to testify in our 2,280 commercial lawsuits. Of these, an overwhelming majority were Muslim. Equally striking is that not one Christian or Jewish witness appears in any intra-Muslim lawsuit. In sharp contrast, Muslims represent the majority of all witnesses called to testify in intra-Christian disputes (56.6%).

The presence of institutionalized judicial biases raises the question of why Mus-

[29] The two inter-faith proportions differ at the 99.9% level of statistical significance (t=8.54).

[30] In both courts, Muslim-Christian trials exceed Christian-Muslim trials at the 99.9% significance level (t=8.28 for Galata and t=4.39 for Istanbul).

lims and Christians sued each other at all on economic matters. After all, Christians could have avoided the possibility of commercial lawsuits involving Muslims simply by doing business with other Christians. Yet potential gains from trade would often have trumped the risk of biased litigation. A content analysis of the issues over which cases arose shows that Christians and Muslims interacted most commonly through credit markets is presented in Table 2.8. Of all trials between Christians and Muslims, 65.5 percent concerned debt. Typically, the plaintiff accuses the defendant of having failed to settle a debt linked to an installment sale with an implicit interest charge; hence, in the data debt and sale are positively correlated.[31][32]

To sum up thus far, we have (1) found a higher pro-plaintiff bias for Christians than for Muslims, (2) observed that, in line with the global norm of anti-outsider judicial bias, Islamic courts were institutionally biased against non-Muslims, and (3) inferred that the anti-Muslim bias of the courts led Christian subjects to avoid suing Muslims except when their cases were particularly strong, resulting in a selection bias in the available data. Table 2.8 raises a potential problem with the first and third points. Could the figures in question be reflecting differences across the religious groups in the types of issues they brought to court? This possibility will be addressed in due course, after we deal with a possible objection to point 2.

2.7 State Officials in Islamic Courts

Because the Islamic legal system lacked a concept of legal personhood, in none of our cases does the state per se appear as a litigant. Individuals with a grievance against

[31] The coefficient of correlation between debt and sale is 0.16. It is significant at the 95% level.

[32] The plaintiff is the creditor and the defendant the debtor 94.1 percent of the time. Given that the Muslim litigant is the plaintiff in 75.6 percent of the Muslim-Christian lawsuits involving a debt, we can infer that the Muslim was more often the creditor and the Christian generally the debtor. Evidently, in the seventeenth-century, in contrast to later times, Muslims supplied more credit than Christians. This disparity does not reflect a sample selection bias, because our usual indicator is absent: even though the number of Christian-initiated lawsuits is lower than those brought by Muslims, the plaintiff win rates of the groups are statistically identical (t=1.37).

the state had to sue the state official responsible for the offending act. Likewise, individuals who shirked a civic responsibility could be taken to court by a state official rather than the state itself. The authority and responsibility of the state were thus personified. The officials who appeared in court most frequently were tax officials and estate supervisors. Palace employees sometimes showed up in connection with guild matters and communal affairs. Of the 173 state officials in our data, 12.1 percent were either Jewish or Christian, as demonstrated in Table 2.9.[33] Non-Muslims were commonly employed as tax collectors, so this is not surprising. It bears reiteration that data drawn from adjudicated cases need not mirror the underlying religious distribution of officials.

State officials appear in about 7 percent of all of our seventeenth-century disputes, a total of 158 cases. In 15 of these lawsuits they face each other, leaving 143 cases in which a state official faces a subject. State officials appear as both plaintiff and defendant, at a statistically equal frequency (Table 2.10).[34]

Turning our attention to the outcome of these adjudications, we find significant variation in the plaintiff victory rate depending upon the side of the state official. Whereas it is 59.6 percent when neither litigant is an official, it plummets to 30.3 percent when the plaintiff is an official and jumps to 85.7 percent when the defendant is an official (Table 2.11).[35] When officials face each other, the plaintiff victory rate is indistinguishable statistically from the baseline case.[36]

It appears, then, that cases between officials and subjects are very likely to be decided in favor of the latter. However, in these cases we do not observe an inequality

[33] All of these cases are recorded in Kuran (2011b), vols. 3-4, where additional statistical breakdowns may be found.

[34] We fail to reject the hypothesis of symmetry at a meaningful level of significance ($\chi^2(1) = 0.85$).

[35] The hypothesis of equality is rejected at the 99.9% significance level in both cases (t=5.05 and t=6.29, respectively).

[36] The hypothesis that the two plaintiff victory rates are equivalent is not rejected (t=1.16).

in the number of cases adjudicated. Hence, the sample selection bias argument is not applicable here.

Intriguingly, the change in the plaintiff victory rate is invariant to the religion of the subject (Tables 2.12 and 2.13). Evidently, the courts weight the testimony of a state official equally regardless of the religion of his opponent in court. This finding confounds the previous result of judicial bias against non-Muslims. Irrespective of religion, when a non-official agrees to confront a state official in court, his case is sufficiently strong to withstand any pro-state judicial bias. Indeed, this may account for the symmetry displayed in Table 2.10. In disputes involving a state official the sample selection bias is so great that subjects go to court only if certain of victory. Such cases are sufficiently rare to maintain statistical symmetry.

Finally, we turn our attention to substance of the cases that involve a state official. Table 2.14 shows the topics covered in adjudications whose litigants included at least one state official. Tax, inheritance, and debt disputes involve officials more frequently relative to the baseline where no official is involved. Tax disputes typically entailed disagreements between tax collectors and subjects over tax obligations; in some of the cases a subject complains of harassment by an official when he has already paid his tax to another collector. Inheritance disputes stemmed from the Ottoman law that assigned the estates of heirless individuals to the state. The state's confiscation of an estate would be challenged by a person claiming to be a rightful heir. Similarly, debt adjudications will involve an official because unpaid taxes are considered a debt to the state.

It is curious to observe the tremendous predictive power that the presence of a state official brings to the ability to forecast the outcome of adjudication. As a rule, judges decide against litigants employed by the state. Given the judiciary's lack of independence from the executive, this observable outcome is almost certainly caused by extreme sample selection bias. No individual who chooses to appear in

court opposite an official would do so unless highly certain of winning. However, this result does not mirror that found in our analysis of anti-minority bias, since the number of observed cases brought against the state is nearly identical to the number of cases the state brings against private individuals. While this is most likely caused by extreme sample selection bias, we are unable to test this hypothesis adequately due to the lack of extra-judicial data pertaining to the cases settled out of court.

2.8 Documented Contracts and Insurance against Judicial Bias

We already know that Ottoman courts did more than adjudicate lawsuits. Of the 10,080 records in our 15 court registers, 6,494 comprised the registration of a contract or settlement. Ottoman subjects registered agreements in court to have a record in writing as insurance against misunderstandings. When an agreement was entered into a court register, each side received a copy. Either copy, or the record in the register itself, could be consulted in the event of a dispute. Could the institutionalized biases of the Islamic courts be mitigated through documentation of agreements? One might expect non-Muslim Ottoman subjects to have alleviated the risks of pro-Muslim litigation by documenting their commercial interactions with Muslims. Likewise, subjects of all faiths might have lessened the dangers of pro-state litigation by documenting their transactions involving the state. In fact, some of our litigants did present documents to bolster their testimony. As Table 2.15 indicates, a document was introduced by one side or the other, or both, in 15.2 percent of all trials. The table also shows that the presentation of a document massively increased the chances of victory. The plaintiff win rate, which is 60.2 percent when the sides make their cases without reference to any documentation, jumps to 83.9 percent when the plaintiff alone introduces a document in support of his case. Equally striking, it plummets

to 7.2 percent when the defendant alone presents a document.[37]

By themselves, these figures do not establish the commonness of documenting commercial transactions. A potential plaintiff's decision to seek a judicial ruling would have been influenced by whether he had supportive documentation. By the same token, he would be motivated to avoid court fees by settling out of court insofar as the document would convince the defendant that he would lose. A better proxy for the extent of document use is the rate at which defendants introduced documents, because it was not they who sought adjudication. But defendants, too, could settle, and they could deter a lawsuit simply by showing a potential plaintiff their relevant documents. Nevertheless, potential defendants probably ended up in court at a higher rate if they had documentary evidence likely to exonerate them. Thus, the share of trials involving documents probably overstates the level of commercial documentation in seventeenth-century Istanbul.

We are still left with the question of why documentation use was so low. The courts charged for documents. The judges of our two courts collected a 2 percent ad valorem fee registering an estate settlement and between 8 and 30 aspers for a document certifying a court registration or verdict (*hüccet*); lesser fees were collected by their assistants (Uzunçarşılı, 1965). All in all, the cost of registering a debt or sale contract corresponded to what a skilled worker made in one to three days.[38] For small transactions, these fees alone would have discouraged registration. Low literacy rates—no greater than 10 percent for any group—would also have limited the demand for documentation. Yet another reason to forego documentation was that in Islamic jurisprudence documents themselves lacked evidentiary value in the absence of corroboration by witnesses to their preparation. The witnesses could charge for

[37] The two differences (83.9 percent vs. 60.1 percent) and (7.2 percent vs. 60.1 percent) are both statistically significant at the 99.9% level (t=8.69 and t=19.61, respectively).

[38] Özmucur and Şevket Pamuk (2002) estimate that in Istanbul a skilled worker made around 22.5 aspers per day at the start of the seventeenth century and about 36.9 aspers at the end.

their services.

In a modern economy, a signed and notarized contract signals that the signers accepted its contents and that the terms of a transaction were set in advance. In a lawsuit it will serve as evidence that the parties understood their duties associated with the transaction. In the present context, the document includes a list of witnesses to its creation and indicates that they could resolve any conflict as to what was agreed. If parties to a court-registered contract end up in court, and one side introduces the contract, it serves notice that witnesses are available to speak to its validity and testify to its particulars. The other side might concede the case at that point.[39] In the absence of a concession, witnesses will be called, and it is their testimony that will clinch the case for the document presenter.[40] The upshot is that registering a contract in court was not simply a matter of drafting and recording its terms. A cadre of mutually agreed witnesses had to be found. And a special challenge was present whenever the sides included one or more Muslims along with one or more non-Muslims. The rules of Islamic jurisprudence put non-Muslims at a disadvantage in documenting such contracts. The degree of trust would have been higher between co-religionists than between members of different communities Hence, we would expect the use of documentary evidence to be greater in inter-faith cases than in intra-faith cases. While the raw numbers tentatively support this hypothesis—document use in 16.3 percent of all inter-faith cases versus 14.3 percent of all intra-faith cases—the paucity of observations precludes statistical significance.

Perhaps the most important reason for the low rate of documentation observed in our registers is not the cost of documentation itself but that in the seventeenth century the Ottoman economy had not yet begun the transition from personal to

[39] In our data set, one or more documents is introduced in 351 of 2,287 trials. In 184 (52.9%) of these, no witnesses were called because the other side conceded the case.

[40] Uzunçarşılı (1965). The theme appears in one of our registers: Istanbul 22 (1695), 93a/1.

impersonal exchange. Commercial organizations were tiny and short-lived, and business took place largely among acquaintances (Kuran, 2011a). Irrespective of religion, people enforced contracts largely through reputation-based means, turning to courts only as a last resort and in extraordinary circumstances. Indeed, extrapolating from the average number of disputes in our registers, we find the probability of any given business transaction leading to litigation to be around 0.05 percent.[41] In borrowing from acquaintances, or buying an object from a seller with whom one has interacted repeatedly, the expected benefit of documentation is limited. In fact, it may damage the relationship itself by signaling mistrust.

Nevertheless, document use should be more prevalent in the court records when the risk of losing is greater. That risk should peak on matters involving a state official, since the courts are predisposed to rule in favor of the state. Table 2.16 shows that document use does indeed rise when a state official is involved, relative to the full sample of cases.[42]

In Table 2.16, 47 of the 158 (29.7%) cases that involve a state official, documentary evidence is presented to the court. As with the full sample, the side presenting a document has a much greater chance of winning. Moreover, the extreme sample selection bias discussed in Section 2.9 still results in the subject winning the majority of all cases involving a state official. But the value of a document remains high. When a subject submits documentary evidence, he wins 95.8 percent of the time, as opposed to 74.8 percent when no documentary evidence is submitted. Similarly, when the state official submits documentary evidence, he is much more likely to win.[43]

[41] This estimate is based on the assumption that every adult participates in 20 important commercial interactions per year and that the adult population of Istanbul amounted to 66.7% of 700,000 residents.

[42] Documentary evidence appears in 29.7 percent of all trials involving a state official. In those that do not involve an official, documentary evidence is submitted in only 13.6 percent of all cases. The difference is significant at the 99.9% level (t=4.32).

[43] The two differences (95.8 percent vs.74.8 percent) and (33.3 percent vs.74.8 percent) are both

2.9 Multivariate Analysis of Judicial Decisionmaking

Having examined the institutional biases in the Ottoman judicial system, noted the apparent asymmetries in court participation among Istanbul's religious communities, and explored the role of documents as insurance against pro-plaintiff and pro-state bias, we are left to question to whether the differing plaintiff victory rates between religious communities—specifically the asymmetries involving Muslim-Christian trials—were driven by inter-group differences concerning the substance of trials. Muslims and Christians might have appeared in debt or waqf cases at disproportionate rates, and any judicial bias may have varied according to topic. A related factor that could sway a judge is the presence of an official among the litigants.

The relative weights of such influences can be disaggregated through a logistic regression framework. This exercise will enable us to separate the impact of the litigants' religion, their status in relation to the state, the substance of the dispute, and document use. To facilitate interpretation, we provide the results through odds ratios.

2.9.1 Estimation Model

The logistic regression estimates the equation $\text{Prob}(y_{it} = 1|\beta, \sigma^2) = \text{Prob}(x'_{it}\beta + z'\eta + \epsilon_i > 0)$ where ϵ is distributed logistically with mean zero. More specifically, we decompose the vector $x'_{it}\beta$ as equal to

$$x'_{it}\beta = a'_{it}\delta + x'_{it}\gamma + x'_{it}\xi$$

where $y_{it} = 1$ when when the adjudicator of case i in year t rules in favor of the plaintiff and $y_{it} = 0$ otherwise; and a_i, b_i, and c_i are vectors that identify, respectively, whether either or both litigants are Muslim, whether either litigant is employed as a

statistically significant at the 99.9% and 99% level, respectively (t=3.58 and t=3.41).

state official, and which litigant introduced documentary evidence at the trial. The vector z consists of year-specific dummy variables to control for fixed effects on cases adjudicated in the same year.

Since the religion-specific variables are not interacted with the employment or document variables, we are assuming that the weight of documentary evidence in the judge's decision making is invariant to the faith of the party introducing the evidence. In addition, we are assuming that the testimony of state-employed individuals is weighed by the judge consistently irrespective of their religion.

2.9.2 Specifications

Under these assumptions, we present in Tables 2.17 and 2.18 estimates for four specifications, each of which adds a new class of controls. (1) controls for religion of the litigants; (2) adds controls for cases where the litigants include an official; (3) adds controls for the dispute topic; and (4a-c) control for documentary evidence. Of the final trio, (4a) simply adds controls to specification (3); while (4b) and (4c) estimate specification (3) on cases that either include or exclude documentary evidence.

In all the regressions, our reference specification involves a non-Muslim plaintiff and a non-Muslim defendant. Hence, the odds ratio for the variable "Defendant Muslim" indicates the change in the relative probability of the judge ruling for the plaintiff when the defendant is Muslim rather than non-Muslim. The specifications determine whether the above-discussed plaintiff victory rate differences in Muslim-Christian trials are robust when controlling for factors that might have influenced judicial decisions.

Religion of Litigants

Interestingly, in no specification does the plaintiff's religion affect the judge's decision making process. Consistently, this variable turns out statistically insignificant. Fur-

thermore, the religion of the defendant is significant at the 99% level in specification (1), and at the 95% level in (2), (3), and (4a). Hence, without controls for documentary evidence, a judge is more likely to rule in favor of the plaintiff when a Christian sues a Muslim than in the reverse case. This mirrors the results in Table 2.15.

To interpret the lack of statistical significance in specification (4a), we estimate specification (3) twice more while limiting our dataset first to cases without documentary evidence and then to those with documentary evidence. The estimated results of specification (4b) are similar to those presented in (3). By contrast, in specification (4c) the evidence of sample selection bias indicative of prejudicial bias is no longer significant. Evidently, non-Muslim apprehensions about encountering a prejudiced judge are alleviated when documentary evidence is present. Although documentation cannot change the judge's attitudes toward non-Muslims, it keeps him from disregarding persuasive evidence, thus limiting the impact of his ingroup bias. That is why, in specification (4a), which controls for documentary evidence, the previously observed effect of non-Muslims suing Muslims only when certain of a win is diminished.

Finally, the statistical significance of the variable "All Muslim" helps interpret the odds ratio of "Defendant Muslim." In each specification it indicates that when a Muslim sues another Muslim, the probability of a plaintiff victory returns to the baseline specification of adjudication among non-Muslims. Evidently, judges were just as likely to rule in favor of the plaintiff in intra-Muslim cases as in cases among non-Muslims.[44]

[44] When both litigants are Muslim, the Plaintiff Muslim, Defendant Muslim, and All Muslim indicator variables are equal to 1. In this case, the statistically negligible coefficient of Plaintiff Muslim plays no role, but the coefficient on All Muslim is statistically the negative of that of Defendant Muslim. In each specification, the sum of the two coefficients is statistically equal to 0 (χ^2=0.01 for specification (1), χ^2=0.35 for specification (2), χ^2=0.81 for specification (3), χ^2=2.41 for specification (4), χ^2=0.83 for specification (5), and χ^2=2.16 for specification (6), all with 1 degree of freedom.)

Presence of State Officials

Specifications (2)–(4c) all control for the presence of state officials among the litigants. The controls reflect the expectation that judges would have avoided putting themselves in conflict with state authority. Each of these specifications points to a strong anti-state bias. Specifications (2)–(4a) indicate that judges are about 4 times more likely to rule in favor of the plaintiff when the defendant is a state official. Moreover, the plaintiff is 3.3 times more likely to lose a case if he is a state official, all else being equal.

As with the high plaintiff victory rate for minorities, these results probably reflect a sample-selection bias. Potential defendants threatened with a lawsuit by a state official would have opted to settle out of court unless very confident of prevailing in court. Similarly, potential plaintiffs would have refrained from suing a state official before a state-appointed judge unless their case was very strong.[45] In brief, the observed lawsuits involving a state official represent the disputes in which the official case was particularly weak.

In specification (4c), the results for trials involving state officials differ from those of the preceding specifications. For cases where documentary evidence is presented, a plaintiff who is an official has the same probability of winning as a plain subject, irrespective of religion. In other words, document-holding officials who initiate adjudication have no greater chance of winning than any other document holder. Revealingly, these results are not repeated in instances where the defendant is an official. Since state officials are more likely than subjects to possess documentary evidence, when we limit the sample to trials where documents are submitted as evidence the plaintiff will have to be highly certain of victory to challenge an official; in itself documentary evidence does not provide a high likelihood of winning since the

[45] Unfortunately these hypotheses cannot be tested directly without data on settlement rates, which are unavailable.

defendant may be able to rest his case on documents of his own. However, in the sample restricted to cases with documents, plaintiffs suing an official are 11 times more likely to win, relative to our baseline specification.

Topic of Dispute

When the dispute topic is controlled for, the estimates of the earlier regressions remain statistically significant and constant (specifications (3)–(4c)). Evidently, the qualitative nature of the dispute, while undoubtedly germane to the adjudication process, does not make the judge weigh the religion of the litigants differently. The last four specifications suggest also that judges were more likely to rule for the plaintiff when adjudicating disputes over guild business and property transactions.

Documentary Evidence

Specification (4c) offers overwhelming support to our earlier observation that documents carry significant weight in trials. According to Table 2.18, if the plaintiff introduces documentary evidence to the court, his odds of winning increase four-fold. More dramatically, when a defendant challenges the plaintiff's account through documentary evidence, the judge is 25 times less likely to rule in favor of the plaintiff! As proposed earlier, putting contracts in writing provides substantial insurance against breach of contract. Whatever the biases of the Ottoman judicial system, people could boost their chances of winning a commercial lawsuit merely by having a written document to support their case.

2.10 Conclusions and Implications for the Present

Since no existing country has a fully functioning Islamic legal system, studying its past applications can provide the necessary clues as to whether it satisfies basic conditions of the rule of law. Through an examination of seventeenth-century Ottoman

court data, we have shown that lack of judicial independence from the government and discrimination against non-Muslims may be among the consequences of instituting a legal system modeled on past applications of Islamic law.

As practiced within the Ottoman Empire, Islamic law was characterized by institutionalized judicial biases against religious minorities and lack of judicial independence from executive authority. Courts were biased against non-Muslims in practice as well. We drew this inference from a comparison of lawsuits involving a non-Muslim plaintiff and a Muslim defendant with lawsuits in which the roles are reversed. Though relatively fewer, inter-faith lawsuits of the first kind exhibit a greater plaintiff victory rate, which suggests that non-Muslims preferred to settle claims outside the court unless reasonably confident of winning. The institutionalized biases of the courts must have been among the major reasons. The result holds up in a multivariate analysis that controls for involvement by a state official, the issue in dispute, and the use of documentary evidence. In lawsuits pitting a subject against a state official, the plaintiff victory rate is much higher when the lawsuit is initiated by the former. It appears that subjects were reluctant to sue a state official unless very confident of prevailing, though the asymmetry in numbers lacks significance.

Legal reformers of the nineteenth century invoked these factors, along with the unsuitability of Islamic courts to modern financial and commercial practices, in campaigning for secular commercial courts. They considered traditional Islamic courts to be contributors to the region's deepening economic problems. By the 1850s, they had managed to establish secular commercial courts in Istanbul, Cairo, and Alexandria. Insofar as their justifications were grounded in fact, and the secular courts did adjudicate more impartially, the new courts would have contributed to jump-starting the Middle East's catch-up process. Native Christians and Jews were becoming increasingly important economic players, and leveling the playing field in their favor would have stimulated trade, especially internal trade involving Muslims. Studying

the records of the commercial courts should enable the testing of whether the biases identified here did, in fact, diminish.

One might object that conditions in the twenty-first century are different enough that the foregoing findings are irrelevant to the re-imposition of Islamic law. Another possible objection is that the Ottoman courts of the seventeenth century are unrepresentative of genuinely Islamic courts that existed only in the early centuries of Islam, in the religion's Arab heartland. Both of these objections involve empirical matters. Did the courts of early Islam not discriminate against non-Muslims? Did their judges avoid favoritism toward state officials? If one or both of these questions has a negative answer, one would want, of course, to know why practices differed. With respect to the first objection, one would want to know exactly what differs now. Is it that judges are insulated from political pressures? Have the proponents of Islamization agreed to abrogate the longstanding practice of treating Muslim testimony as inherently superior to that of non-Muslims?

The proponents of reinstituting Islamic law have tended to avoid such questions by focusing on ideal conditions that are unlikely ever to be attained. In the Islamic society of their imagination, all public officials have infused an Islamic morality, which makes them averse to corruption and favoritism. Accordingly no judge rules in favor of a state official merely for fear of state reprisals, and no state official would expect him to do so anyway.[46] Yet there is no evidence that morality alone can eliminate the biases in question. In no known society are public officials immune to material incentives. Moreover, according to Muslim accounts widely accepted among Islamists, no known Muslim society after the year 661, the end of Sunni Islam's canonical golden age, has been free of corruption. These accounts beg the question of how the biases identified in this paper would be eliminated. There is an

[46] For examples of social designs that treat Islamic morality as an antidote to corruption, see Afzal-ur-Rahman (1980); Chapra (1992), and Naqvi (2003).

additional reason why today's proponents of Islamic law avoid addressing the possibility of injustices toward non-Muslims: the presumption, discredited here through econometric techniques, that historically Islamic courts ruled impartially on cases involving non-Muslims.

There is a difference in conditions that would mitigate the biases identified in this paper, were courts similar to those of seventeenth-century Istanbul to go into operation. The use of documentation is much more widespread in the modern Middle East than in the time period covered in our empirical work. Documentary evidence mitigated the biases inherent in Islamic justice. Subjects could protect themselves against pro-state biases and non-Muslims against pro-Muslim biases simply by documenting their transactions. It is impossible, however, to write a complete contract that accounts for every possible contingency. Conflicts on matters that could not have been anticipated are inevitable, as are disagreements over contract wording. Hence, documentation alone would not eliminate the biases.

Economic globalization has raised the importance of judicial independence. As one of the determinants of investment risk, it affects capital flows. In particular, capital generally flows toward countries where the courts are relatively independent of the executive branch of government. Sooner or later, therefore, the supporters of an Islamic legal system will have to come to terms with potential costs of a lack of judicial independence. They will need to deal with the structure of courts staffed by corruptible judges who are responsive to material incentives.

Insofar as globalization promotes economic interactions across national and religious boundaries, it also raises the potential costs of pro-Muslim judicial biases. The adoption of a legal system that weights witness testimony by the religion of witnesses would pose an impediment to commerce and thus economic growth. Companies setting up shop in a country under Islamic law would expect compensation for the risks incurred by their employees. In countries with a substantial non-Muslim population,

such as Egypt, where estimates of the Christian Coptic population ranges between 10 and 20 percent, pro-Muslim judicial biases would take an economic toll both by weakening minorities and by impeding cross-religious commercial interactions. In predominantly Muslim countries without a substantial non-Muslim minority, such as Saudi Arabia, similar costs could arise from judicial biases against socially unpopular or officially disfavored Islamic sects. The very logic of disregarding non-Muslim testimony could be used to justify the discounting of, say, Shii testimony.

The demand for reinstituting Islamic law receives a sympathetic hearing from some intellectuals and policy makers committed to preserving or reviving local cultures. Insofar as multiculturalism is itself desirable, one can make a case for using the Islam's rich legal heritage as a basis for legal reforms in the Islamic world. Nevertheless, there are tradeoffs that should be recognized and considered. Imposing Islamic law could harm Muslims economically by restricting their participation in the global economy.

Table 2.1: Trial Records by Register

Register	Years	Obs	Percent
Galata 25	1604	139	6.1
Galata 27	1604–05	155	6.8
Istanbul 1	1611–13	78	3.4
Istanbul 2	1615–16	50	2.2
Galata 41	1616–17	40	1.7
Galata 42	1617	85	3.7
Istanbul 3	1617–18	143	6.2
Istanbul 4	1619	107	4.7
Istanbul 9	1661–62	549	24.0
Istanbul 16	1664–65	163	7.1
Galata 130	1683	201	8.8
Galata 145	1689–90	247	10.8
Istanbul 22	1694–96	172	7.5
Istanbul 23	1696–97	111	4.8
Galata Total		913	39.9
Istanbul Total		1378	60.1
Total		2,291	100.0

Table 2.2: Adjudications by Court – Muslim Share of Litigants

Court	Obs	Muslim Plaintiff	Muslim Defendant***
Galata	913	71.4	63.0
Istanbul	1378	75.9	70.9
p-value		0.02	0.00

***: Difference between Galata and Istanbul significant at 99.9% level.

Table 2.3: Adjudications by Court – Shares of Categories

Court	Obs	Tax**	Waqf**	Rent	Property	Sale	Guild
Galata	913	6.7	12.8	8.8	21.0	26.1	3.5
Istanbul	1378	3.6	17.3	8.9	19.1	27.3	2.8
p-value		0.00	0.00	0.89	0.26	0.52	0.37

***: Difference between Galata and Istanbul significant at 99% level.

Table 2.4: Number of Trials by Religion

Plaintiff	Defendant			
	Muslim	Christian	Jewish	All
Muslim	1411	223	41	
Christian	96	377	7	
Jewish	20	29	20	
All				2224

Table 2.5: Plaintiff Victory Rate by Religion

Plaintiff	Defendant			
	Muslim	Christian	Jewish	All
Muslim	59.0	57.9	65.9	
Christian	71.9	59.4	57.1	
Jewish	65.0	55.2	70.0	
All				59.7

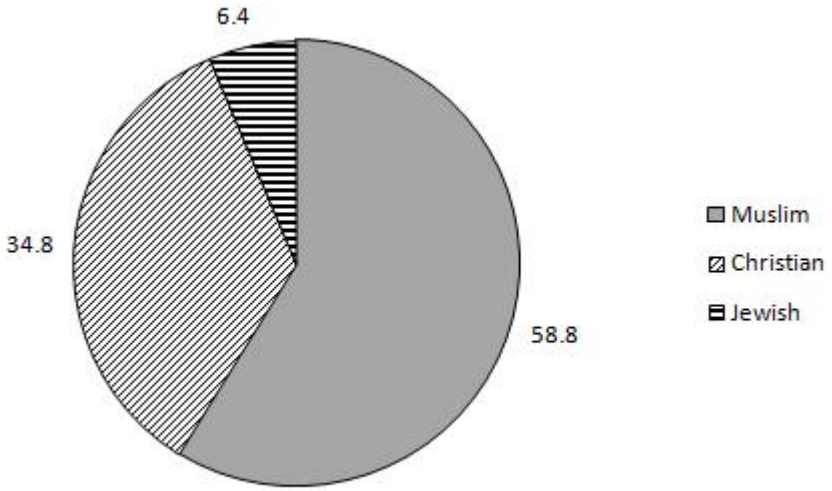

FIGURE 2.1: Population Shares of the Religious Communities in Seventeenth-Century Istanbul

Table 2.6: Share of Private Adjudications by Religion of Litigants

Litigants		(1)	(2)
Plaintiff	Defendant	Random	Observed
Muslim	**Muslim**	**34.6**	**63.3**
Muslim	Christian	20.5	9.9
Muslim	Jewish	3.7	2.0
Christian	Muslim	20.5	3.7
Christian	**Christian**	**12.1**	**18.1**
Christian	Jewish	2.2	0.3
Jewish	Muslim	3.7	0.9
Jewish	Christian	2.2	1.0
Jewish	**Jewish**	**0.4**	**0.9**
Total		100	100

Table 2.7: Religious Distribution of Litigant-Selected Witnesses

Litigants		Witnesses			
Plaintiff	Defendant	Muslim	Christian	Jewish	Total
Muslim	**Muslim**	**741**	**0**	**0**	**741**
Muslim	Christian	62	12	0	74
Muslim	Jewish	42	0	0	42
Christian	Muslim	39	16	0	55
Christian	**Christian**	**99**	**76**	**0**	**175**
Christian	Jewish	0	4	0	4
Jewish	Muslim	5	0	0	5
Jewish	Christian	9	5	0	14
Jewish	**Jewish**	**0**	**0**	**13**	**13**
All Above Combination		997	113	13	1,123
All Three Religious Groups		27	0	0	27
Other Combinations[a]		23	4	0	27
All Trials	Number	1,047	117	13	1,177
	Percent	89.0	9.9	1.1	100.0

[a]: Lawsuits in which one or both sides includes a foreigner (*müstemen*), gypsy (*çingene*), or recent convert to Islam (*mühtedi*).

Table 2.8: Categories as Percentage of Inter-Faith Adjudication

Category	(1) Christian Suing Muslim	(2) All Cases with Christians and Muslims	(3) Muslim Suing Christian	t-test (1) = (3)
Guild	6.3	3.8	2.7	1.31
Debt	53.1	65.5	70.9	2.97
Waqf	14.6	17.9	19.3	1.05
Property	22.9	11.0	5.8	3.72
Tax	7.3	6.0	5.4	0.62
State	19.8	11.9	8.6	2.51
Sales	36.5	30.1	27.4	1.58
Partnership	29.2	21.0	17.5	2.20
Inheritance	29.2	20.7	17.0	2.92
Obs	96	319	224	

Note: Columns sum to more than 100% because many cases involve several of these topics simultaneously.

Table 2.9: State Officials by Religion

	Muslim	Christian	Jewish	All
Number	152	7	14	173
Percent	87.9	4.0	8.1	100.0

Table 2.10: Distribution of Litigants by State Affiliation

	Defendant		
Plaintiff	Subject	Official	All
Subject	2,066	77	
Official	66	15	
All			2,224

Table 2.11: Plaintiff Victory Rate by State Affiliation

	Defendant		
Plaintiff	Subject	Official	All
Subject	59.6	85.7	
Official	30.3	73.3	
All			59.7

Table 2.12: Plaintiff Victory Rate According to Involvement of State Officials — By Plaintiff's Religion

	Defendant		
Plaintiff	Subject	Official	t-test
Muslim	57.9	83.3	5.54
Christian	60.9	85.0	2.80
Jewish	32.3	–	–

Table 2.13: Plaintiff Victory Rate According to Involvement of State Officials — By Defendant's Religion

	Defendant		
Plaintiff	Muslim	Christian	Jewish
Subject	60.7	59.9	64.6
Official	37.5	33.3	100.0
t-test	3.22	2.96	5.92

Table 2.14: Plaintiff Victory Rate According to Involvement of State Officials — By Defendant's Religion

Category	Subject	Official	t-test
Guild	2.9	7.0	1.99
Debt	58.7	38.0	5.14
Waqf	15.9	12.7	1.78
Property	19.7	15.8	1.29
Tax	2.8	32.3	7.86
Sales	27.5	20.9	1.94
Partnership	19.7	22.8	0.89
Inheritance	28.1	55.7	6.76
Obs	2,066	158	

Note: Columns sum to more than 100% because many cases involve several of these topics simultaneously.

Table 2.15: Document Use in Trials

	No Document	Presented by Plaintiff	Presented by Defendant	Presented by Both	Total
Plaintiff Win Rate	60.3	83.9	7.2	21.1	59.6
Obs	1,938	218	111	19	2,286

Table 2.16: Document Use in State Trials

	No Document	Presented by Subject	Presented by Official	Presented by Both	Total
Subject Win Rate	74.8	95.8	33.3	80.0	72.8
Obs	111	24	18	5	158

Table 2.17: Odds Ratio of Logistic Regression on Plaintiff Victory: No Document Controls

VARIABLES	(1) All	(2) All	(3) All
Plaintiff Muslim	0.968	1.025	1.064
	(0.153)	(0.163)	(0.172)
Defendant Muslim	2.017***	1.809***	1.708**
	(0.452)	(0.413)	(0.393)
All Muslim	0.508***	0.513**	0.519**
	(0.133)	(0.136)	(0.138)
Plaintiff State Official		0.364***	0.350***
		(0.0905)	(0.0945)
Defendant State Official		4.405***	4.628***
		(1.346)	(1.429)
Waqf			0.801*
			(0.101)
Sale			1.115
			(0.114)
Property			1.410***
			(0.171)
Tax			0.980
			(0.239)
Guild			2.315***
			(0.684)
Year-Specific Dummies	Yes	Yes	Yes
Observations	2,287	2,287	2,287
Pseudo R^2	0.04	0.05	0.06

s.e. in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 2.18: Odds Ratio of Logistic Regression on Plaintiff Victory: Document Controls

VARIABLES	(4a) All	(4b) No Doc	(4c) Doc
Plaintiff Muslim	1.062	1.062	1.144
	(0.181)	(0.190)	(0.490)
Defendant Muslim	1.545*	1.674**	2.209
	(0.371)	(0.434)	(1.224)
All Muslim	0.520**	0.470**	0.626
	(0.145)	(0.141)	(0.415)
Plaintiff State Official	0.312***	0.282***	0.546
	(0.0930)	(0.0944)	(0.297)
Defendant State Official	4.458***	3.551***	11.09***
	(1.473)	(1.244)	(7.817)
Waqf	0.787*	0.712**	1.354
	(0.107)	(0.105)	(0.407)
Sale	1.101	1.026	1.592*
	(0.119)	(0.117)	(0.443)
Property	1.653***	1.615***	0.609
	(0.215)	(0.222)	(0.194)
Tax	1.246	1.359	0.456
	(0.340)	(0.413)	(0.221)
Guild	2.096**	1.731	4.090***
	(0.714)	(0.692)	(2.002)
Plaintiff Document	3.530***		
	(0.718)		
Defendant Document	0.0450***		
	(0.0172)		
All Document	0.152***		
	(0.0968)		
Year-Specific Dummies	Yes	Yes	Yes
Observations	2,287	1,936	345
Pseudo R^2	0.13	0.06	0.14

s.e. in parentheses

*** p<0.01, ** p<0.05, * p<0.1

3

Health Shocks and Compliance in Chile's Defined-Contribution Pension Program

3.1 Introduction

Chile's adoption of a defined-contribution retirement savings program in 1981 has been heralded by some economists as a model for maintaining long-run social protection in middle-income countries (The World Bank, 1994; Gill et al., 2005). By divesting itself of the demographic and financial risk associated with a pay-as-you-go defined-benefits pension program, Chile has offered its citizens competitive returns on retirement savings in conjunction with risk-mitigating retirement programs that limit market volatility and guarantee a mandatory minimum pension to those who regularly contribute to their retirement savings program.

In a perfectly enforceable system, the defined-contribution system would credibly commit all individuals to adhere to the requirements of the program: Everybody must contribute about 10 percent of their income to their retirement account to finance his or her pension and, after contributing for 20 years, individuals are entitled to the full range of risk mitigation programs the system offers to protect savers from

market risk. However, Chile is characterized by a large informal sector that lacks the informational infrastructure to credibly commit all workers to the mandatory pension contribution requirement.

Like most middle-income countries, a large subsection of the population is able to bypass legal requirements by finding employment in the large informal sector of the economy. When individuals fail to contribute regularly, they may not qualify for minimum pension guarantee that limits the financial risks of the defined-contribution retirement system. In extreme situations when ill-timed market volatility erodes retirement savings upon reaching retirement age, individuals who failed to comply with the mandatory savings requirement may be forced to receive their pension from the defined-benefit social safety program designed for the indigent at an income level lower than the guaranteed minimum pension for fully compliant individuals. When this occurs, the sustainability of adopting a defined-contribution system are weakened due to the noncompliance and potential under-saving of individuals.

There is a longstanding line of concern over defined-contribution pension programs for specifically this reason (Orszag and Stiglitz, 2001; Barrientos, 1996). Becker et al. (2009) highlights the uncertainty inherent in macroactuarial forecasts involved in implementing Kazakhstan's defined-contribution pension program, which are highly sensitive to uncertain foundational assumptions on savings rates and compliance. Furthermore, Acuña and Iglesias (2001) maintain that coverage under Chile's defined-contribution pension program has failed to reach the heights under the previous pay-as-you-go retirement system — with the ratio of contributors to individuals signed up to the program falling from 73.6 percent in 1982 to 50.9 percent in 2000, which the authors attribute to periodic unemployment and individuals who leave the labor market due to health and family reasons. Finally, Barr (2002) underscores the importance of good governance to sustain a defined-contribution pension program: A government must be able to credibly maintain economic stability, regulate finan-

cial markets, and possess the administrative capacity to collect taxes and enforce contributions.

Given the systemic challenges caused by noncompliance, it is curious that there has been no rigorous analysis of the failure to comply with the mandatory pension contribution requirement. This paper posits a model designed to explore individuals' decision to comply with the mandatory pension contribution requirement. It accounts for the risk of personal and family negative health shocks that reduce household income, potentially forcing individuals to choose between complying with the law and maintaining household consumption.

The remainder of the paper is organized as follows. In Section 3.2, the institutional nature of the Chilean pension and health insurance markets is explored and contextualized. Following that, Section 3.3 provides a formal model that initially examines the relationship between retirement savings and health insurance, eventually extending the model to consider the longevity effect of health insurance and how it relates to retirement savings decisions. Section 4.3 introduces the underutilized EPS dataset that allows for the econometric analysis of individuals' retirement savings decisions. Section 3.5 discusses the methodology employed to analyze these decisions while acknowledging and controlling for potential endogeneity in the model's specification. Finally Section 4.4 presents the findings of the paper.

3.2 Chile's Institutional Context

3.2.1 The Chilean Pension System

The current incarnation of the Chilean pension system was created in 1981 during the rule of General Augusto Pinochet. Prior to these reforms, Chile utilized a multitude of pay-as-you-go pension systems for its labor and military constituencies that have been characterized as fragmented and inefficient (Piñera, 2004). Following the 1981 reforms, the structure of the national pension system was reformed to become a three-

pillar social protection system composed of A) a pay-as-you-go system to care for the indigent and provide a mandatory minimum pension for those who contributed for 20 years to the second pillar, B) the Administradoras de Fondos de Pensiones (AFPs), a system of privately provided pension funds in which individuals are required to invest a share of their income, and C) a voluntary program through which individuals could invest more than the mandatory minimum of their earnings for retirement. This paper focuses on the second pillar.

The AFPs are similar to mutual funds in which individuals are required to invest about 10 percent of their income until about US$1,500.[1] These investments are invested in funds of the contributor's choosing. They accumulate based upon the performance of the portfolios the AFPs manage. In this defined-contribution retirement savings program, the performance of the AFPs is partly contingent upon the asset classes in which they may invest (Kotlikoff, 1999). Srinivas and Yermo (1999) note that the assets in AFPs have been increasingly diversified over time as regulators remove barriers to investing in equities, corporate debt, and international assets. The performance of the funds is regulated by the Chilean government, which measures AFP returns against each other and requires that when an AFP's returns exceed a centrally determined cutoff value that tracks the average AFP returns, excess earnings must be saved in a reserve fund to supplement the AFP's performance when earnings are below the permitted performance band. In this manner, AFPs' equity risk is mitigated and the volatility of returns on pension savings are reduced greatly.

Individuals are permitted to change which AFP their savings are invested in once a year. Switching was a regular occurrence until the Chilean government in 1997 required individuals to submit their request to change AFPs in person in an

[1] Individuals must spend about 2-3 percent of their income in fund-specific administrative fees and to finance disability and survivor insurance in addition to saving 10 percent of their income (Arenas de Mesa et al., 2008).

attempt to reduce the frequency of pension switching to limit administration costs. Mitchell et al. (2007) find that the 1997 regulatory changes reduced switching, but add that the level of switching was comparable to trading in US 401(k) retirement plans.

Although the law stipulates that all income-earning individuals who are not self-employed must contribute to their pension programs, there are several avenues by which Chileans evade this duty. First, Chile has a sizeable informal labor sector. Because a firm in the informal labor sector does not report income to the government, it need not withhold and transfer its employees' mandatory share of earnings to the pension funds. Second, individuals employed in the formal sector and otherwise required to contribute to their savings may reach an agreement with their employer to underreport income and hence undercontribute to their pension savings. However, in the analysis that follows, we treat all individuals who contribute some share of their income as complying with the requirements of the AFP system.

It is the aim of this paper to explain the variation in compliance with the requirements on retirement savings by examining the roles that health insurance and health shocks play in individuals' savings decisions. Previously, Arenas de Mesa et al. (2008) found that Chileans contribute roughly half of their working life — more frequently than under the previous pay-as-you-go system according to Gill et al. (2005). Furthermore, this paper will build on Arenas de Mesa et al. (2004), which finds that the decision to comply with the mandatory contribution requirement is correlated with health insurance status, by modeling the decision to invest in retirement savings and examining role health shocks play in individuals' decisionmaking. To examine how individuals' health status affects their retirement savings decisions, we must examine the Chilean health insurance market.

3.2.2 Health Insurance in Chile

Health insurance in Chile can be described as "mandatory but not universal" (Sapelli and Torche, 2001). While the law requires all individuals to purchase health insurance at a minimum cost of 7 percent of total income,[2] the coverage purchased by that fee is at the discretion of the individual. In general, there are two health insurance sectors in Chile: The public insurance program, which is called the *Fondo Nacional de Salud* (FONASA); and the private insurance sector, whose firms are designated *Instituciones de Salud Previsional* (ISAPRE).

A review of these two parallel systems will illuminate the characteristics of each program's participants. Sapelli and Torche (2001) show that while fees for public insurance — which are based solely on the income of the individual purchasing insurance independent of the number of dependents covered, preexisting health conditions, and age — are generally lower than those for privately supplied insurance, a third of Chile's richest quintile is enrolled in the public program and a tenth of Chile's poorest households purchases private insurance. This mixing is most likely caused by the public sector insurance program serving as an insurer of last resort when individuals receive a significant negative health shock and transition from private insurance to the FONASA for medical care (Holst et al., 2004; Barrientos, 2002).

Nevertheless, as Holst et al. (2004) show, the vast majority of FONASA contributors earn less than US $400 a month. By restricting our attention to the effects of *public* health insurance on pension compliance, we are over-representing the neediest and most likely to be informally employed Chileans. Yet this restriction allows us to understand the cost of health insurance because by law the fee for public health insurance is limited to a fixed percentage of total income, rather than subject to market forces stemming from the composition of individual insurance policies which are

[2] Chileans may enroll in public insurance at a fixed rate of 7 percent of income up to about US $200 a month (Sapelli, 1999)

unobservable to the econometrician. However, Chile's public health system remains competitive against the private system by segmenting its subscribers by income into several policy groups that offer varying degrees of services. The public health system comprises four groups: Group A, which covers the indigent and working poor; Group B, which covers the lower class; and groups C and D, which cover the middle and upper class Chileans who choose to enroll in public insurance (Barrientos and Lloyd-Sherlock, 2000).

Group A provides those covered under these plans access to public health facilities only. Groups B, C, and D grant access to public facilities, but also provide opportunities to go out of network and receive medical care from private clinics which accept public insurance.[3] Thus FONASA provides varying degrees of coverage for individuals based upon their income. As demonstrated in Table 3.1, the the higher-income public insurance policies provide less coverage for medical costs than the lower-income policies to finance the cost of subsidizing the poorer subscribers to the public insurance system while remaining competitively priced compared to private insurance policies. Also of note, the 7 percent charge is waived for individuals covered in Group A who fall below the poverty line.

For this paper's analysis, we will focus on the public insurance program. This is due to the general uniformity in coverage within each policy group that public insurance provides, conditional upon which region an individual resides.[4] This paper's purpose is to examine how these two institutions interact. In brief, we seek to determine whether the health insurance mandate enforces or weakens the ability of individuals to comply with the mandatory pension contribution requirement. However, before turning to the data, it is necessary to formalize a model of the

[3] When going out of network, the share of liability covered differs from that in Table 3.1. Individuals are covered for between 50 and 75 percent of the cost of medical services (Sapelli, 1999)

[4] Bruce (2000) finds significant variation in the provision of public health by region.

individual's allocation of scarce resources between the two institutions.

3.3 Modeling Retirement Savings and Health Insurance Decisions

3.3.1 Baseline Specification

The analysis of deciding whether to contribute to the pension plan is conducted using a three-period model. In this model, individuals decide whether to participate in public health insurance in the first period and, contingent on that decision, decide whether to contribute to their retirement program in period two. Finally, individuals consume their retirement savings in period 3. Each individual is described by a metric of health $\lambda \in (0, 1]$, where higher values of λ correspond to better health and a smaller risk of health shocks.

For the purposes of this model, we will treat the first period's income as fixed at y_1 and the second period income y_2 as a stochastic variable with a support such that $y_2 \in [0, 1]$ is distributed according to the probability density function $f(y_2|p, \lambda, \alpha)$, where $p \in \{0, 1\}$ is binary variable that indicates whether the individual is uninsured ($p = 0$) or insured ($p = 1$) and $\alpha \in (0, 1]$ is a measure of the quality of health insurance.

Although income in period two is stochastic, it is important to make certain assumptions about the relative distributions of y_2. Namely, for all individuals $\lambda < 1$ and insurance systems $\alpha > 0$, $f(y_2|p = 1, \lambda, \alpha)$ first-order stochastically dominates $f(y_2|p = 0, \lambda, \alpha)$, which is to say,

$$F(y_2|p = 0, \lambda, \alpha) \geq F(y_2|p = 1, \lambda, \alpha)\, \forall y_2,\ \forall \alpha < 1$$

The underlying rationale is that health shocks negatively impact income in period two, and for all individuals who are not perfectly healthy, being insured increases the likelihood of receiving a higher draw of y_2.

Therefore, the individual needs to choose p and c to maximize the lifetime utility $V(p,c)$:

$$
\begin{aligned}
V(p,c) =& u\left[[y_1]^{(1-p)}\left[(1-\gamma)y_1\right]^p\right] + \\
& \beta\left[\int u\left[[(1-\xi)y_2]^c\,[y_2]^{(1-c)}\right] f(y_2|p=0,\lambda,\alpha)dy_2\right]^{(1-p)} + \\
& \beta\left[\int u\left[[(1-\gamma-\xi)y_2]^c\left[(1-\gamma)y_2\right]^{(1-c)}\right] f(y_2|p=1,\lambda,\alpha)dy_2\right]^{p} + \\
& \beta^2\left[\int u\left[[\xi r y_2]^c + d\right] f(y_2|p=0,\lambda,\alpha)dy_2\right]^{(1-p)} + \\
& \beta^2\left[\int u\left[[\xi r y_2]^c + d\right] f(y_2|p=1,\lambda,\alpha)dy_2\right]^{p}
\end{aligned}
\tag{3.1}
$$

where $c \in \{0,1\}$ is a binary variable that indicates whether the individual contributes to his pension program ($c=1$) or does not contribute ($c=0$), $\gamma < 1$ is the cost of public insurance, $\xi < 1$ is the proportion of income that is required to be contributed to the pension system, r is the return on capital invested in the pension system — which is held constant for the purposes of this model — and β is the inter-temporal discount rate. All individuals are guaranteed a minimum level of retirement income by the state which is defined as $d > 0$.

In Chile, both the mandatory contribution to the pension system and the fee for admission to the public insurance program are denominated as a percentage of total income. The cost of public insurance in Chile is equal to 7 percent of total income ($\gamma = 0.07$) and the mandatory minimum contribution to the pension program is 10 percent of total income ($\xi = 0.10$).

The main thrust of this analysis is to examine under what conditions individuals contribute to the pension program in Chile. In this model, the decision to contribute is decided in period two, after the realization of y_2. Therefore, individuals have

already chosen whether to purchase insurance ($p = 1$) or not ($p = 0$).

An uninsured individual will contribute to his pension if and only if

$$u\left[(1-\xi)y_2\right] + \beta u\left[\xi r y_2 + d\right] \geq u\left[y_2\right] + \beta u\left[d\right], \tag{3.2}$$

and an insured individual similarly chooses to contribute if the following holds

$$u\left[(1-\gamma-\xi)y_2\right] + \beta u\left[\xi r y_2 + d\right] \geq u\left[(1-\gamma)y_2\right] + \beta u\left[d\right]. \tag{3.3}$$

To arrive at a closed-form solution, let $u(\cdot) = \log(\cdot)$. Under this specification, let us define $\bar{y}_2^p$ as the income at which an individual is indifferent between contribution and not contributing to his pension, given that $p \in \{0, 1\}$. Under the assumption of log-utility and utilizing equations (2) and (3), it is possible to define the income above which individuals will contribute to the pension program and below which they will fail to contribute, given their insurance status, as

$$\bar{y}_2^0 = \frac{\left[\frac{1}{(1-\xi)}\right]^{\frac{1}{\beta}} d - d}{\xi r} \tag{3.4}$$

$$\bar{y}_2^1 = \frac{\left[\frac{(1-\gamma)}{(1-\gamma-\xi)}\right]^{\frac{1}{\beta}} d - d}{\xi r}. \tag{3.5}$$

where $\bar{y}_2^0$ is the pension contribution cutoff value for uninsured individuals and $\bar{y}_2^1$ is the pension contribution cutoff value for individuals who are enrolled in a health insurance program.

The existence of a cutoff value in period two income predicts that particularly damaging income shocks that reduce wealth below the threshold value $\bar{y}_2^p$ will prompt individuals to fail to comply with the legal requirement to contribute to their pension plan. Both insured and uninsured people are at risk of receiving a sufficiently negative health shock that results in a failure to contribute. However, the probability

of receiving these health shocks will differ according to the distribution of income contingent upon whether the person is insured, personal health, and the quality of the insurance program.

From equations (4) and (5), it is straightforward to demonstrate that $\bar{y}_2^1 > \bar{y}_2^0$. That is to say, there exists a range of income $y_2 \in [\bar{y}_2^0, \bar{y}_2^1]$ where insured individual will not contribute to retirement savings when uninsured individuals will. Hence, retirement savings and health insurance are substitutes within this income region, all else being equal.

This curious result is driven by the cost of insurance. Since Chile's public health insurance program is not costless, two identical individuals who receive the same income draw $y_2 \in [\bar{y}_2^0, \bar{y}_2^1]$ but differ in their insurance status will differ in their willingness to contribute to the pension system due to the diminished contemporaneous utility caused by the consumption lost in paying for insurance.[5] While society's sickest individuals will fail to save for retirement (even though the probability of mortality in this model is nil), those who have chosen to purchase insurance will fail to save after receiving a weaker negative health shock than those who are uninsured.

However, just because some of the ill and uninsured will save for retirement when the ill and insured cannot afford to do so, the rate of failing to comply with the requirements of a defined-contribution retirement system will not necessarily be higher for the insured than the uninsured. Stochastic dominance is not a suitably strong condition to determine whether insured individuals are more likely to default on their savings requirement than uninsured individuals. Since the threshold values $\bar{y}_2^1 > \bar{y}_2^0$, it is impossible to determine the relative conditional probability that an individual fails to contribute to his pension program conditional on his health and

[5] This model treats insurance as a binding commitment while contributing to retirement savings is non-binding. While this assumption drives the finding that low draws of income result in a failure to contribute to retirement savings instead of a failure to pay one's insurance bill, the fact that the low income draw is caused by a negative health shock in this model eliminates the incentive to drop healthcare upon receiving a low draw, lending tractability to this key assumption.

insurance status:

$$F(\bar{y}_2^1|p=0,\cdot,\cdot) \geq F(\bar{y}_2^0|p=0,\cdot,\cdot) \geq F(\bar{y}_2^0|p=1,\cdot,\cdot)$$
$$F(\bar{y}_2^1|p=0,\cdot,\cdot) \geq F(\bar{y}_2^1|p=1,\cdot,\cdot) \geq F(\bar{y}_2^0|p=1,\cdot,\cdot) \tag{3.6}$$

$$F(\bar{y}_2^0|p=0,\cdot,\cdot) \gtrless F(\bar{y}_2^1|p=1,\cdot,\cdot)$$

Rather, it is necessary to specify the scale by which the cumulative distribution function (CDF) $F(y_2|p=0,\lambda,\alpha)$ differs from $F(y_2|p=1,\lambda,\alpha)$ for each value of y_2. Therefore, let

$$F(y_2|p,\lambda,\alpha) = \begin{cases} 0 & \text{if } y_2 \leq 0 \\ y_2^{p\alpha(1-\lambda)+\lambda} & \text{if } 0 < y_2 < 1 \\ 1 & \text{if } y_2 \geq 1 \end{cases} \tag{3.7}$$

where λ is the metric of the general healthiness of an individual and α measures the quality of health insurance. When an individual is perfectly healthy ($\lambda = 1$), the distribution wealth that he could receive in period two is equivalent to the uniform distribution between the points $[0,1]$. However for any $\lambda < 1$, one observes from Figure 3.1 that the probability of receiving a negative health shock (lower value of y_2) increases as λ decreases.

In Figure 3.1, there are three CDFs depicted corresponding to values of $\lambda \in \{0.3, 0.5, 0.7\}$. It can be shown that the expected value of y_2 is increasing as $\lambda \to 1$, as greater health causes the CDF to converge to the uniform distribution on the support [0,1], which has a corresponding expected value equal to 0.5. Table 3.2 the expected values of y_2 corresponding to each values of λ depicted in Figure 3.1.

Extending the depiction of the CDF defined in Equation 3.7, we can demonstrate the effect of insurance on the CDF of a particularly unhealthy individual from Figure 3.1 characterized by $\lambda = 0.3$. As a rule, for any given $\lambda < 1$ and $0 < \alpha < 1$, being insured ($p = 1$) decreases the probability of a large income loss contingent

on receiving a negative health shock. In Figure 3.2, the baseline specification is the uninsured individual with $\lambda = 0.3$ from Figure 3.1. The additional curves are the CDF of the same individual upon purchasing insurance characterized by $\alpha \in \{0.3, 0.7\}$. As expected, being insured stochastically dominates being uninsured, with the expected value of y_2 increasing as $\alpha \to 1$, as demonstrated in Table 3.3.

Utilizing this specification of the CDF of period two income (y_2), it is possible to estimate the probability that an individual does not contribute to his retirement program conditional upon one's insurance status, health status, and the quality of available health insurance. Therefore, the probability that an uninsured person fails to contribute to his pension program is the probability that $y_2 < \bar{y}_2^0$ or,

$$P(c = 0|0, \lambda, \alpha) = F(\bar{y}_2^0|0, \lambda, \alpha) = \left[\frac{\left[\frac{1}{(1-\xi)}\right]^{\frac{1}{\beta}} d - d}{\xi r} \right]^{\lambda} \tag{3.8}$$

and the probability that an insured person fails to contribute to his pension program is likewise $y_2 < \bar{y}_2^1$ or,

$$P(c = 0|1, \lambda, \alpha) = F(\bar{y}_2^1|1, \lambda, \alpha) = \left[\frac{\left[\frac{(1-\gamma)}{(1-\gamma-\xi)}\right]^{\frac{1}{\beta}} d - d}{\xi r} \right]^{\alpha(1-\lambda)+\lambda} \tag{3.9}$$

Now that the conditional probabilities are fully specified, it is possible to examine the conditions under which uninsured individuals are more likely to renege on their commitment to contributing to retirement programs than insured individuals due to receiving a sufficiently damaging negative health shock.

$$P(c=0|0,\lambda,\alpha) > P(c=0|1,\lambda,\alpha) \Longleftrightarrow \tag{3.10}$$

$$\frac{\lambda}{(1-\lambda)}\left\{\frac{\log\left[\frac{\left[\frac{1}{(1-\xi)}\right]^{\frac{1}{\beta}}d-d}{\xi r}\right]}{\log\left[\frac{\left[\frac{(1-\gamma)}{(1-\gamma-\xi)}\right]^{\frac{1}{\beta}}d-d}{\xi r}\right]}-1\right\}=\bar{\alpha}<0<\alpha$$

In other words, for any quality of insurance α, the proportion of uninsured individuals who default on their retirement contribution obligation will be larger than the proportion of insured individuals who default.[6] This result may be interpreted as evidence that for any quality of insurance $\alpha \in (0,1]$, purchasing insurance reduces the probability of receiving a poor second-period income due to a negative health shock. For any individual who is not perfectly healthy ($\lambda < 1$), the magnitude of the reduced probability is sufficient to compensate for the marginal cost of health insurance, which is depicted as the ratio of utilities within the brackets in Equation 3.10.

3.3.2 Incorporating Pre-Retirement Mortality Risk

The specification of the model offered in Equation 3.1 contains only one stochastic component: The propensity for a health shock in period two. Under this specification, an individual fails to contribute to his state-mandated retirement program due to the unexpected cost of a contemporary health shock. While these short-term debilitating health shocks could have a tremendous impact on one's propensity to save, it is necessary to consider how one's health impacts the probability of reaching retirement age. Simply said, a competing reason for failing to contribute to a retirement program is that an individual's health is so poor that he does not expect to

[6] This result comes from assuming an interior solution to the problem, in which both $\bar{y}_2^0$ and $\bar{y}_2^1$ are in [0,1]. From this assumption, it follows that the ratio of the two logarithms is less than one and therefore the cutoff value $\bar{\alpha}$ is less than zero.

reach retirement age.

It is possible to incorporate the risk of pre-retirement mortality as a probability of reaching retirement age $\phi(p, \lambda, \alpha)$ which is a function of whether one has insurance, the overall level of health of an individual, and the quality of the insurance program in which he is enrolled. It is reasonable to assume that,

$$\begin{aligned} \phi(1, \lambda, \alpha) &> \phi(0, \lambda, \alpha) && \forall\ 0 < \lambda \cdot \alpha < 1 \\ \frac{\partial \phi}{\partial \lambda} &> 0 \\ \frac{\partial \phi}{\partial \alpha} &> 0 && \forall\ \lambda < 1 \text{ and } p = 1 \end{aligned}$$

which implies that for any level of health or insurance quality, being insured increases the probability of reaching retirement age. The rationale of this assumption comes from the increased access to medical care that insurance brings. A voluminous literature exists on the role that preventative care plays in improving health outcomes. The probability of reaching retirement age increases also as one's level of health λ increases. Finally, as the quality of insurance (α) increases, the probability of reaching retirement age increases.

Therefore, by altering the individual's maximization problem in Equation 3.1 to include the probability that the individual reaches retirement age $\phi(p, \lambda, \alpha)$, we arrive at the following:

$$
\begin{aligned}
V(p,c) =& u\left[[y_1]^{(1-p)}\left[(1-\gamma)y_1\right]^p\right] + \\
& \beta\left[\int u\left[[(1-\xi)y_2]^c\,[y_2]^{(1-c)}\right] f(y_2|p=0,\lambda,\alpha)dy_2\right]^{(1-p)} + \qquad (3.11)\\
& \beta\left[\int u\left[[(1-\gamma-\xi)y_2]^c\left[(1-\gamma)y_2\right]^{(1-c)}\right] f(y_2|p=1,\lambda,\alpha)dy_2\right]^{p} + \\
& \phi(p,\lambda,\alpha)\left\{\beta^2\left[\int u\left[[\xi r y_2]^c+d\right] f(y_2|p=0,\lambda,\alpha)dy_2\right]^{(1-p)} + \right. \\
& \left. \beta^2\left[\int u\left[[\xi r y_2]^c+d\right] f(y_2|p=1,\lambda,\alpha)dy_2\right]^{p}\right\}.
\end{aligned}
$$

Including the possibility of pre-retirement mortality in Equation 3.11, complicates the decision of whether to invest in health insurance. No longer does insurance provide a larger expected lifetime utility through the second and third period consumption level alone. Now the individual must take into account the risk of not reaching his final period of consumption and thus forgo even the minimum of retirement consumption provided by the minimum guaranteed pension d.

As earlier, an uninsured individual in period two characterized by health status λ and given the opportunity to purchase health insurance of quality α will choose to contribute to his pension fund if and only if

$$
u\left[(1-\xi)y_2\right] + \phi(0,\lambda,\cdot)\,\beta u\left[\xi r y_2 + d\right] \geq u\left[y_2\right] + \phi(0,\lambda,\cdot)\,\beta u\left[d\right],^{7} \qquad (3.12)
$$

and an insured individual will choose to contribute to his retirement program if and only if,

$$
u\left[(1-\gamma-\xi)y_2\right] + \phi(1,\lambda,\alpha)\,\beta u\left[\xi r y_2 + d\right] \geq u\left[(1-\gamma)y_2\right] + \phi(1,\lambda,\alpha)\,\beta u\left[d\right]. \qquad (3.13)
$$

[7] Recall that the quality of insurance α is not at issue when p=0.

By assuming log-utility once again, one can calculate a closed-form solution for the cutoff point $\hat{y}_2^p$ over the support of y_2, below which an individual will not contribute to his retirement program, contingent upon whether he is enrolled in insurance, *ceteris paribus*,

$$\hat{y}_2^0 = \frac{\left[\frac{1}{(1-\xi)}\right]^{\frac{1}{\phi(0,\lambda,\alpha)\,\beta}} d - d}{\xi r}, \tag{3.14}$$

$$\hat{y}_2^1 = \frac{\left[\frac{(1-\gamma)}{(1-\gamma-\xi)}\right]^{\frac{1}{\phi(1,\lambda,\alpha)\,\beta}} d - d}{\xi r}. \tag{3.15}$$

Under the previous model, it was possible to declare that the cutoff value $\bar{y}_2^0 < \bar{y}_2^1$. This is because the difference between the two values was proportional to the prices charged for participation in the insurance program. However, the decision to participate in the insurance program is no longer solely a function of the improved expected value of income in the second period and the price of the insurance program. It must now take into account the improvement in probability of surviving to retirement. Hence, the relative positioning of the cutoff values $\hat{y}_2^p$ on the spectrum [0,1] is no longer explicit. Rather, it is characterized by

$$\hat{y}_2^1 > \hat{y}_2^0 \Longleftrightarrow \frac{\phi(0,\lambda,\alpha)}{\phi(1,\lambda,\alpha)} > \frac{\log\left[\frac{1}{(1-\xi)}\right]}{\log\left[\frac{(1-\gamma)}{(1-\gamma-\xi)}\right]}. \tag{3.16}$$

By the definition of $\phi(p,\lambda,\alpha)$, we know that $\frac{\phi(0,\lambda,\alpha)}{\phi(1,\lambda,\alpha)} < 1 \quad \forall\lambda < 1$. Thus, we arrive at the interpretation that dependent upon:

1. The cost of purchasing health insurance and participating in the the retirement program, and

2. The comparative increase in the probability that an individual will live to retirement upon purchasing health insurance,

there will exist some region on the possible spectrum of y_2 wherein health insurance and investment in retirement benefits are either substitutes (when Equation 3.16 binds) or complements (when Equation 3.16 does not hold.)[8]

If we normalize $y_2 = 1$, the right-hand side of Equation 3.16 becomes the ratio of the marginal cost of contributing to a pension program when uninsured to the marginal cost of contributing to a pension program when insured. The left-hand side is the ratio of probability of reaching retirement age given that one is uninsured to the probability of reaching retirement age when insured. In short, when the probability of reaching retirement age does not differ greatly between being insured and uninsured $\left(\frac{\phi(0,\lambda,\alpha)}{\phi(1,\lambda,\alpha)} \to 1\right)$ then the uninsured will be more willing to contribute to their retirement fund when suffering a serious health shock than an insured individual would be under the same situation.

To illustrate this point, recall that for a given individual characterized by λ and a given health insurance system characterized by α, when $\hat{y}_2^1 > \hat{y}_2^0$, the potential values of y_2 can be represented on the line segment given in Figure 3.3.

As one can see in Figure 3.3, the condition given in Equation 3.16 is satisfied, meaning that the probability that an uninsured individual lives to retirement is not different enough to overcome the comparative costs of paying for insurance and contributing to the retirement program. For the comparatively wealthy and healthy, there is no tradeoff between contributing to one's retirement fund and being insured. By contrast, for the less well off, there exists a region where only the uninsured are willing to save for retirement when suffering an extreme health shock.[9] Since the

[8] For the rest of the intuition that Equation 3.16 can either bind or not bind, note that $\frac{\log\left[\frac{1}{(1-\xi)}\right]}{\log\left[\frac{(1-\gamma)}{(1-\gamma-\xi)}\right]} > 0$

[9] To be clear, the probability of living until retirement age is independent of the severity of the

risk of mortality is only slightly different than when the individual is insured, the uninsured will still be willing to contribute to his pension fund when extremely sick.

However, when Equation 3.16 does not hold, we can interpret this situation as one in which the uninsured are much more likely to perish prior to reaching retirement age. In this case, when the uninsured individual suffers an extreme health shock ($y_2 < \hat{y}_2^0$), he is not willing to contribute to his retirement fund because his probability of living to consume it is sufficiently low. This is demonstrated in Figure 3.4.

In Figure 3.4, we see the opposite situation than was presented in Figure 3.3. In this example, the insured have a much higher probability of reaching retirement age than the uninsured, all else being equal. Once again, the well off need not agonize over whether to contribute to their retirement fund. However, for those who receive a negative health shock that lands them within the region $[\hat{y}_2^1, \hat{y}_2^0]$, the uninsured will not contribute to their retirement fund because they have a low probability of living to consume their savings, whereas the insured will still continue to save. Of course, those who suffer the most extreme health shocks will not save for retirement regardless of their insurance status, much as in the previous example.

Finally, we return to the question posed in the previous section: What is the probability that an individual fails to contribute to his pension program given his insurance status? Previously, the analysis was decided by the inequality $\bar{y}_2^1 > \bar{y}_2^0$. However, in this case, the inequality is contingent upon the relative probabilities of reaching retirement age, as demonstrated by Equation 3.16.

In the case where Equation 3.16 holds, it is necessary to define a CDF for the stochastic variable y_2 since, as demonstrated in Equation 3.6, stochastic dominance

specific health shock an individual receives in period two. The realized value of y_2 is correlated with the probability of living until retirement only in the sense that those who have a low value of λ are both likely to receive a poor draw of period two income and have a low probability of reaching retirement. The value of y_2 provides no addition information to the individual since he already knows his overall level of health λ; rather, the low draw of y_2 forces a tradeoff between consuming the entirety of one's income in period two or saving a fraction of it to potentially consume it the next day.

is not a sufficient condition to compare the relative probabilities of defaulting on one's contribution requirements. Therefore, it will be necessary to utilize Equation 3.7 to determine the conditional probabilities:

$$P(c=0|0,\lambda,\alpha) = F(\bar{y}_2^0|0,\lambda,\alpha) = \left[\frac{\left[\frac{1}{(1-\xi)}\right]^{\frac{1}{\phi(0,\lambda,\alpha)\beta}} d - d}{\xi r}\right]^{\lambda} \tag{3.17}$$

and

$$P(c=0|1,\lambda,\alpha) = F(\bar{y}_2^1|1,\lambda,\alpha) = \left[\frac{\left[\frac{(1-\gamma)}{(1-\gamma-\xi)}\right]^{\frac{1}{\phi(1,\lambda,\alpha)\beta}} d - d}{\xi r}\right]^{\alpha(1-\lambda)+\lambda}. \tag{3.18}$$

Under this specification, it is possible to restate Equation 3.10 as

$$P(c=0|0,\lambda,\alpha) > P(c=0|1,\lambda,\alpha) \Longleftrightarrow \tag{3.19}$$
$$\frac{\lambda}{(1-\lambda)}\left\{\frac{\log\left[\frac{\left[\frac{1}{(1-\xi)}\right]^{\frac{1}{\phi(0,\lambda,\alpha)\beta}} d-d}{\xi r}\right]}{\log\left[\frac{\left[\frac{(1-\gamma)}{(1-\gamma-\xi)}\right]^{\frac{1}{\phi(1,\lambda,\alpha)\beta}} d-d}{\xi r}\right]} - 1\right\} = \bar{\alpha} < 0 < \alpha$$

From Equation 3.19, one sees that as in the previous example, the probability that an uninsured individuals fails to contribute to his retirement program is greater than the probability that an insured individual fails to contribute.[10]

However, when Equation 3.16 does not hold — and therefore $\hat{y}_2^1 < \hat{y}_2^0$, stochastic dominance is a necessary and sufficient condition to prove $P(c=0|0,\lambda,\alpha) > P(c=0|1,\lambda,\alpha)$ without specifying a form for the CDF of y_2 $F(y_2|p,\lambda,\alpha)$, since

[10] The intuition for this claim is that $\phi(1,\cdot,\cdot) > \phi(0,\cdot,\cdot)$, and so the fraction of the left-hand side of the equation cannot be greater than 1.

$$
\begin{aligned}
\hat{y}_2^0 > \hat{y}_2^1 \Longrightarrow F(\hat{y}_2^0|0,\lambda,\alpha) > F(\hat{y}_2^1|0,\lambda,\alpha) \qquad & \text{By non-decreasing nature of CDF.} \\
F(\hat{y}_2^1|0,\lambda,\alpha) > F(\hat{y}_2^1|1,\lambda,\alpha) \qquad & \text{By stochastic dominance.} \\
F(\hat{y}_2^0|0,\lambda,\alpha) > F(\hat{y}_2^1|1,\lambda,\alpha). &
\end{aligned} \tag{3.20}
$$

Therefore, one may assert that the uninsured are more likely to default on the required contributions to their retirement savings than the insured when insurance is a complement of retirement savings, as opposed to when insurance is a substitute of retirement savings — which requires the utilization of a CDF that allows for measuring the scale by which $F(y_2|1,\lambda,\alpha)$ stochastically dominates $F(y_2|0,\lambda,\alpha)$, as proffered in Equation 3.7.

Ultimately, these theoretical specifications elucidate the question: Are public health insurance and mandatory defined-contribution pension programs complements or substitutes? For many members of society, this question is unnecessary — the wealthiest and healthiest amongst the population receive enough income to meet these two obligations regularly. However for the poorest, who are among those most vulnerable to health shocks and least likely to save for retirement, does the Chilean government's policy of separating enrollment in public health insurance and the public pension program force individuals to choose between these two vital public services? This question is analyzed and answered in Section 4.4. First, however, we shall examine the data that are necessary to resolve this puzzle.

3.4 Chile's EPS Dataset

3.4.1 The Survey

Following Chile's pension program reforms in the 1980s, attention turned to monitoring how individuals were transitioning from the previous defined-benefits system to the new defined-contribution program. Aiming to monitor this transition, the Fac-

ultas Economía y Negocios at Universidad de Chile developed the Social Protection Survey (*Encuesta de Protección Social (EPS)*).

The initial round of the EPS, which was published in 2002, was designed to survey a representative sample of all pension contributors. Subsequent rounds, published in 2004, 2006, and 2009, sought to develop a national sample of individuals regardless of whether they are affiliated with Chile's pension program. In the analysis here, the data are drawn principally from the 2006 survey and incorporates lagged variables from the 2004 round.

The EPS surveys are divisible into two parts: The longitudinal questionnaire and the labor history section. The organization of these two sections differ markedly. The longitudinal questionnaire mirrors many other surveys in that a household head is contacted periodically (2004 and 2006) and administered the same survey. The survey questions generally ask about pension awareness and practices, health and health insurance status, family income, savings, and education. Many of the questions in this section are retrospective over the past year or two and, when combined with previous rounds' responses, form a standard panel data framework with one observation per survey round represents one distinct individual.

The labor history section, on the other hand, is a unique facet of this longitudinal data. It allows us to observe each job that an individual has between the two rounds of the survey. Rather than indicate a single observation per individual per year, the labor section of the EPS gives each individual as many observations as he had labor transitions over the time period. Transitions include periods of employment, unemployment, job seeking, and inactivity. For each transition, the survey reports start and end dates.

During the periods of employment, the labor history section reports the level of pension compliance for each job. Accordingly, it distinguishes between contributing the required amount, under-contributing, and noncompliance, although for the

purposes of this paper, all levels of contribution will be considered full compliance since it is impossible to observe the scale of under-contributing that individuals are reporting. In addition, it is possible to observe the annualized labor income for each position and the sector/role of the surveyed individual in the firm, such as: owner, self-employed, employee in private sector, employee in public sector, and family employee without remuneration.

Incorporating both sections of the EPS, one may obtain a complete picture of individual's employment status, health insurance status, and level of compliance with Chile's mandatory pension program. In the subsequent analysis, the unique structure of the labor history section will be supplemented with variables from the longitudinal section to distinguish the career transitions of individuals by their survey-consistent and job-invariant descriptive variables to identify variations among households. However before embarking on this approach, we shall turn our attention to the cross-sectional statistics that provide a better sense of the population observed in the EPS data.

3.4.2 Summary Statistics

Examining the 2006 round of the EPS, one can develop a greater understanding of the degree of compliance with the institutional requirements of every Chilean household. By law, individuals in Chile must be insured and contribute to their pension program so long as they are not self-employed. Since savings and insurance decisions vary by the sex of a household member, the subsequent analysis is restrict to the males in the EPS dataset, who number 7,071 out of the total observations of 14,425 (49%). This decision is predicated on the fact that women are more likely to have gaps in their work histories for births and child care, adding a level of complexity that we have not yet modeled. Throughout the analysis that follows, we will be restricting our attention to the variation between the uninsured and those who are publicly insured

through the FONASA system as described in Section 3.2

Focusing on the males in the sample, one can explore variations between the insured and uninsured in Chile. Judging from Table 3.4, it is immediately apparent that the age and income of the uninsured differ from those of the uninsured. In sharp contrast to the US health insurance market, the uninsured in Chile are on average wealthier than the insured.

This finding illustrates a fundamental difference between the high-income economic context of the United States and the middle-income context in Chile. The availability of public health insurance in middle-income countries such as Chile provides individuals with an inexpensive option to acquire healthcare. Wealthy individuals, however, have a greater outside option when it comes to health insurance, since in addition to the better funded private health clinics to which private insurance provides access, wealthy individuals have the option of traveling abroad for higher-quality health care outside the domestic insurance system. Since this option holds only for the wealthiest individuals, it is helpful to examine where the income of the uninsured in our sample lies in the spectrum of publicly available health insurance. Recalling that public health insurance is divided into different groups that provide varying degrees of coverage by the income of the subscriber, we note in Table 3.5 that the uninsured are indistinct from those who enroll in Group D — the highest-income publicly provided insurance available in Chile.[11]

Since we are examining the effects of health shocks on pension compliance, it is necessary to define how we measure a health shock. By incorporating lagged health status variables from the 2004 EPS dataset into the 2006 round of the longitudinal survey, one can construct two versions of a health shock. In both rounds of the data, the surveyors ask every household member to describe their health state subjectively

[11] As one would expect, the average income is increasing in each public insurance policy since individuals are placed in a public policy according to their income level.

on a scale that goes from "very poor" to "excellent." First, we construct the variable for personal health shock as when the principal respondent of the survey records his state as one of the lowest two status (very poor or poor) in 2006 when in the previous survey round his status was one of the top four. This variable captures the most extreme changes in health. We find that the prevalence of personal health shocks varies with the individual's level of insurance, as demonstrated in Table 3.6

Table 3.6 makes clear that the EPS data contain few incidences of extreme personal health shocks. Therefore, it is useful to create another measure of extreme health shocks that allow us to examine another form of illness. We define family health shocks as an indicator variable that equals one when the change between the average family health status between the two periods is negative and greater than one standard deviation away from the mean of the change. Interestingly, family health shocks are not statistically correlated with the individual's level of insurance.

Turning our attention to the labor history section of the EPS, let us examine how individuals' labor and pension compliance decisions appear in the data. Focusing on the males in the sample who are not self-employed and who are younger than the retirement age of 65, we find that over the observed period between January 2004 and September 2007 (when the final survey of the "2006" round was completed), the average length of time an employed individuals stayed in his job was 20.34 months. While the maximum number of jobs observed for one individual was 9, we find that the median individual in the sample held a single job throughout the entire labor history.

With the average person in the data holding 1.49 jobs over the 32.30 months of observation covered in the labor history section of the EPS, we observe 585 incidents wherein a salaried male under the age of retirement who holds a job fails to comply with the mandatory pension contribution requirements. The remainder of this paper will focus on analyzing the variation in deciding not to comply with the mandatory

pension investment requirement. Given the availability of each individual's work history, it is possible to examine how long an individual is observed until failing to comply. Finally, since an observed individual could fail to comply and then return to full compliance, our analysis allows for people to reenter the pool after failing to contribute.

In all, the 585 instances in which an individual fails to comply average out to 0.25 failures per individual. However, since individuals can fail more than once, this average is misleading. Indeed, as demonstrated in the empirical probability density function (PDF) of the number of instances where an individual fails to comply presented in Figure 3.5, the median individual in the data complies with the mandatory pension contribution requirement throughout his time under observation. Furthermore, about 15 percent of individuals fail to contribute for at least one period of employment over the period in which they are observed.

Since individuals are able to fail more than once when observed by the econometrician, it is illustrative to discuss the empirical recidivism rate observed in the EPS data. Since the PDF in Figure 3.5 is normalized as the total number of failures for each individual in the sample, it is possible to assess what the probability of re-failure to comply. For instance, the probability that an individual fails to contribute more than once conditional upon having failed once is simply computed as the sum of the probabilities that the total number of failures is equal to 2 through 6 divided by the probability that an individual has failed at least once (since each person who has failed 2 or more times has failed once already). Utilizing this methodology, we find that an individual who has failed to comply once has a 21.6 percent chance of failing to comply again in the future.[12]

[12] It is possible to continue this methodology to examine the probability that an individual who has already failed n times will fail more than n times. As one would expect, the recidivism rate increases with the number n from 21.6 percent when $n = 1$ to 35.1 percent when $n = 2$ and 51.5 percent when $n = 3$. However, after attaining this maximum, the recidivism rate begins to decrease

In keeping with the model specified in Section 3.3, we are interested especially in the variation in pension contributions when distinguishing individuals by insurance status. By computing the hazard ratio, which is interpreted as the instantaneous probability of failing to comply with the mandatory pension contribution requirement, conditional upon complying in the past, we can observe from Figure 3.6 that insured individuals are less likely to fail to comply than uninsured individuals.

To further distinguish the relative rates of noncompliance by insurance status, it is possible to divide the publicly insured into four different groups. As discussed in Section 3.2, each public insurance group is covered for different levels of liability according to the insured individual's income. Hence, the members of Group A are the poorest and disabled individuals covered by publicly provided insurance and Group D are the wealthiest individuals covered by the public insurance program.

Figure 3.7 makes clear that compliance with the mandatory pension savings program varies within those who are publicly insured. While there is little distinction between the members of Groups B, C, and D, Group A is indistinguishable from the uninsured. This is a curious, given that, as shown in Table 3.5, the uninsured are most similar to the wealthy in Group D — not to the destitute of Group A.

In order to analyze the variation in pension compliance with greater detail, it is necessary to introduce a more nuanced econometric specification—Cox's Proportional Hazard Model. In addition to introducing this model in the following section, we will demonstrate some of the challenges to implementing the model and offer potential restrictions to limit the effects of endogeneity. We shall also propose a test for robustness.

---

to 23.5 and 25.0 percent for $n = 4$ and 5, respectively. This result is logical, as each failure can span many months and the ability to fail again after 3 incidents decreases since each individual is observed for a limited period of time.

3.5 Empirical Strategy

As described in Section 4.3, the 2006 EPS dataset can be conceptualized as an uneven panel where individual i holds job j, with the number of jobs unequal between the individuals in the sample. This framework lends itself to examining the time to failure where a person/job pair (i, j) endures for time $0 < t_{ij} \leq T_i$, where T_i is the total time person i is observed in the data.

If we define J_i as the total number of jobs that individual i holds over the entire period of observation $[0, T_i]$, we can divide the data that describe individual i into two groups. Let x_{ij} be the job-specific data describing the pair (i, j). This data vary by the job j held by individual i and are constant throughout the period $(t_{ij-1}, t_{ij}] \ \forall j \in [1, J_i]$. Drawn from the labor history section of EPS, these job-specific variables include labor income and the role played by the individual in the firm at which he works.

In addition, let z_i be the individual-specific data that are constant throughout the period of observation. These variables are created from the longitudinal data in the EPS. They can include variables created by subtracting individual i-specific variables from the 2004 round of the data from individual i's response in the 2006 data. These variables include individual i's age, insurance status, and health status.

Since we are interested in examining how long person i complies with the mandatory pension requirement, we will apply a time-to-failure analytical methodology to the data to explain the variation of the dependent variable y_{ij}, which is takes the values

$$y_{ij} = \begin{cases} 0 & \text{when individual } i \text{ contributes when employed at job } j \\ 1 & \text{when } i \text{ fails to contribute when employed at job } j \end{cases}.$$

If we define $\psi_i(t, X_{ij})$ has the hazard of individual i failing to comply when at job

j — that is, the probability that $y_{ij} = 1$ at time t_{ij} conditional upon $y_{ik} = 0 \ \ \forall 0 < k < j$ — with $X_{ij} = [x_{ij}, z_i]$, it is necessary to define a solution for $\psi_i(t, X_{ij})$. Follow Cox (1972), we define the probability as

$$\psi_i(t, X_{ij}) = \psi_0(t) \exp(X_{ij}\beta). \tag{3.21}$$

This specification assumes that the time-component of the hazard rate is independent of the characteristics describing the person/job pair (i, j). This standard assumption is fundamental to Cox's Proportional Hazard model because it allows for the formation of a Hazard Ratio, which is independent of estimating the specific form of the time-variant probability $\psi_0(t)$. Furthermore, since the focus of our analysis is on the variation between individuals, assuming that the time-component of the hazard rate is independent from the person/job pair is akin to controlling for in a standard linear regression technique with either a trend or time-specific fixed effects invariant accross individuals.

$$\begin{aligned} \text{Hazard Ratio} &= \frac{h_0(t) \exp(x_{ij}\beta_1 + z_i\beta_2)}{h_0(t) \exp(x_{kj}\beta_1 + z_k\beta_2)} \\ &= \exp((x_{ij} - x_{kj})\beta_1 + (z_i - z_k)\beta_2) \end{aligned} \tag{3.22}$$

where $k \neq i$.

By sidestepping the issue of defining a specific form for $\lambda_0(t)$, this model enables us to examine the determinants of the variation in the failure to comply with the mandatory pension contribution requirement by focusing on the regressors X_{ij}, which vary with individual i and job j and control for regional variation through the use of dummy variables. Finally, we control for time-specific fixed effects by defining t as the month/year period of observation rather than a measure of the time in sample, which controls for date-specific correlation.

3.5.1 Controlling for Endogeneity

The model's three-period structure requires individuals to choose their insurance status prior to choosing whether to contribute to their pension fund. This model corresponds to our ability to observe the data, since the decision to comply with the pension fund requirement is made on a job-by-job basis while insurance status is decided once at some point over the period $t \in [0, T_i] \quad \forall i$. Since the decision to enroll in health insurance may be made in the interior of this set, we must control for the endogeneity introduced whenever individuals choose simultaneously whether to comply with the pension requirements and enroll in health insurance.

Two alternative methods exist for controlling this endogeneity. First, one can analyze the decision to comply with the mandatory pension requirements controlling for the lagged insurance status of individuals. This data is readily available from the 2004 EPS sample. It was recorded at the end of the 2002-2004 period observed by the surveyors, implying that the 2004 values for insurance status from the starting point for the 2004-2006 period analyzed in this paper. However, during the period 2004-2006 under observation, it appears that a considerable number of individuals surveyed in both 2004 and 2006 transitioned to public insurance from their previously uninsured state, as noted in Table 3.8.

These transitions indicate that during the period 2004-2006, many previously uninsured individuals enrolled in Chile's public insurance program. This supporting evidence of endogeneity leads to a second method for predetermining health insurance status. By restricting our data to those individuals whose insurance status remained constant during the period 2004-2006, one can identify the roles of insurance status and health shocks in determining compliance with the mandatory pension system. As evident from Table 3.9, the 76.2 percent of individuals observed during both the 2004 and 2006 rounds of the EPS survey maintained their insurance status between

the two periods.

Under this restriction, however, the observed individuals are no longer representative of the Chilean population, as in the first specification without any restrictions to control for endogeneity. Indeed, it is impossible to distinguish whether the individuals who remain in the sample have re-examined their health insurance decisions over the period and elected to remain in the same program or have chosen not to re-examine their insurance status, and therefore are less informed about their access to healthcare. Therefore, we will turn to a last test of robustness at the conclusion of analysis to further control for endogeneity through a generalized method of moments framework.

By restricting the data to those individuals whose insurance status remains constant, we are able to control for the endogeneity inherent in simultaneously choosing a dependent and independent variable. Finally, for the purposes of getting a clearer picture of how investment decisions are made, it is necessary to impose the following restrictions upon the data: Limiting the age of individuals to those well below retirement (less than 60), studying only salaried employees at private institutions who are beholden to the mandatory pension requirements, and examining only those individuals observed in both the 2004 and 2006 datasets.

3.6 Results

3.6.1 Cox's Proportional Hazard Model

As discussed in Section 3.5, we begin our empirical analysis by applying Cox's Proportional Hazard Model to examine the determinants of time-to-noncompliance, defined as the first period wherein individuals fail to contribute to their mandatory pension savings. Limiting our attention to salaried employees of private businesses under 60 years of age, we estimate three specifications of the model to analyze the effects of insurance status and health shocks on time-to-noncompliance (Tables 3.10 and 3.11).

In all instances, Specification (1) is the base model that uses contemporaneous insurance status as observed at the end of the period of observation. Specification (2) is the first method employed to limit endogeneity by using lagged values of insurance status from the beginning of the period of observation. Specification (3) relies on the restricted method of controlling for endogeneity, wherein insurance status is held constant by estimating the model on those individuals whose insurance status is constant throughout the period 2004–2006. Finally, Specification (4) uses the same control for endogeneity as in Specification (3), but it examines the effect of health shocks without controlling for whether the individual is insured. Specification (4) is estimated solely for comparison's sake with the test for robustness employed in Section 3.6.2 and is disregarded in the analysis in this section.

Since Cox's Proportional Hazard Model compares similar individuals' time-to-noncompliance, the estimates on the regressors are interpreted as odds ratios. These are bound between 0 and infinity, with numbers above one (of value $x > 1$) indicating that a unit increase in the value of the regressor increases the probability of failing to comply, conditional upon having complied in all previous periods by the value of the estimated coefficient (the hazard rate is multiplied by x).

From Table 3.10, we note that across the board, wealthier individuals are less likely to fail to comply with the mandatory pension contribution requirements. However, an individual's insurance status remains a significant predictor of pension compliance — the insured are significantly less likely to fail to comply with the mandatory pension contribution requirement. While the point estimate on insurance status varies between the different specifications, we find that in all instances the direct effect is such that public insurance enables individuals to comply with the mandatory pension contributions longer than those who are uninsured.

Yet, the interpretation of the coefficient on insurance status must take into account the existence of health shocks. In Table 3.10, the regressions specification

controls for personal health shocks. In this model, the effect of personal health shocks on time-to-noncompliance shows up statistically insignificant, most likely due to the small number of instances in which personal health shocks affect individuals, as evidenced by Table 3.6. Therefore, since we are controlling for personal health shocks and their interaction with insurance status, we must temper our interpretation of the coefficient on insurance status with the knowledge that this coefficient represents the marginal impact of insurance status without considering whether an individual receives a health shock.

To examine the impact of health insurance on time-to-noncompliance conditional upon being sick, we note that the total effect of receiving a health shock when insured is equivalent to the estimated coefficient of insurance status times that of the interacted term. We can test whether this value of multiplied coefficients differs from the coefficient of the personal health shock variable by computing the small-sample F-statistic for each of the three specifications in Table 3.10. In all cases, we fail to reject the null hypothesis that the effect of health insurance is statistically distinct from being uninsured when an individual receives a personal health shock.[13] However, since Chile's health insurance system provides nearly full coverage for health care costs, this finding may result from wages lost due to illness rather than the failure of the public insurance program to cover health care costs, a hypothesis we will further test when controlling for income level and the share of liability covered by each insurance program.

Finally, Specification (4) identifies the aggregate effect of personal health shocks. Here we note that when aggregated, personal health shocks significantly reduce to time-to-noncompliance — indicating that individuals who suffer a personal health shock are much more likely to fail to comply with the mandatory contribution require-

[13] Specification (1): $F(1, 2353) = 0.22$, p-value $= 0.64$; Specification (2): $F(1, 2353) = 0.39$, p-value $= 0.53$; and Specification (3): $F(1, 2027) = 2.37$, p-value $= 0.12$.

ment regardless of insurance status, even though health insurance status significantly prolongs the time-to-noncompliance.

Table 3.11 examines a similar model as that in Table 3.10, the difference being that it controls for family health shocks rather than personal health shocks. From Table 3.7, we know that the incidence of family health shocks exceeds those of personal health shocks, which allows us to get a greater sense for the effects of health shocks on time-to-noncompliance. Examining Table 3.11, we find once again a negative relationship between enrollment in health insurance and time-to-noncompliance. Interestingly, a negative relationship emerges between suffering a family health shock and time-to-noncompliance when an individual is uninsured. The non-negativity of this estimate belies intuition in the theoretical framework established in Section 3.3. Perhaps for individuals in a household with more than one working adult, an extreme negative health shock that reduces the income of another adult in the household increases the importance of the household head's income. This household head may respond by insuring against his own death through contributions to his pension fund.

As in the previous model, the effect of health insurance on time-to-noncompliance is negative. However, when allowing for health shocks, the combined effect of receiving a family health shock while insured is not statistically distinct from receiving a family health shock while uninsured.[14] Once again, we interpret this finding as evidence that lost wages from illness prevent insured individuals from complying with the mandatory pension contribution requirement.

Nevertheless, examining Specification (4), we find that the health insurance significantly prolongs the time-to-noncompliance while family health shocks have no effect. The statistical insignificance of family health shocks is obviously driven by the disparate effects of a family health shock conditional upon insurance status as

[14] Specification (1): $F(1, 2353) = 0.56$, p-value $= 0.45$; Specification (2): $F(1, 2353) = 1.95$, p-value $= 0.16$; and Specification (3): $F(1, 2027) = 2.07$, p-value $= 0.15$.

demonstrated in Specification (3). However, by comparing this specification against that depicted in Table 3.10, we can note the unambiguous difference between different kinds of health shocks — where personal health shocks, although rarer than family health shocks, unambiguously affect the decision to contribute to retirement savings more than family health shocks.

Although we are able to control for labor income in the specifications depicted in Tables 3.10 and 3.11, it is possible to observe the effects of differing wealth and liability coverage by separating insurance status by the differing constituent public insurance programs. However, due to the small number of observations of personal health shocks, this approach is tractable only in the case of family health shocks. As demonstrated in Table 3.11, controlling for family health shocks is not without its peculiarities when it comes to the impact of family health shocks on the uninsured — the curious findings depicted earlier carry over into the results displayed in Table 3.12.

By regressing the time-to-noncompliance on the different insurance programs, one can control simultaneously for the wealth of the individual and the level of health care coverage supplied to each policyholder. The first result that is revealed under this specification is the finding that the poorest individuals, who are in insurance Group A, do not behave differently than the uninsured in terms of time-to-noncompliance when in good health. This result was previously revealed in Figure 3.7, where we noted that the uninsured and Group A had similar hazard functions. However, even though they behave similarly when in good health, we note that they differ in their time-to-noncompliance upon receiving a family health shock.

In Specifications (2) and (3), we see that individuals in Group A who receive a family health shock have a shorter time-to-noncompliance (that is, are more likely to fail to contribute to the mandatory retirement savings program) than uninsured individuals who receive a family health shock. Furthermore, even though we noted in Table 3.5 that uninsured individuals have a similar income to the wealthy individuals

insured in Group D, we find that their behavior differs both in the baseline time-to-noncompliance when individuals do not receive family health shocks and in the case where family health shocks alter individual's behavior. Comparing the uninsured and those in Group D in times of good health, we find that members of Group D are significantly more likely to contribute to the mandatory retirement savings program than the uninsured. However, in the event of a family health shock, members of Group D are significantly less likely to comply with the law (greater hazard rate) than the uninsured — who are more likely to contribute when faced with a family health shock.

In all cases, we find that the insured are more less likely to contribute to their mandatory retirement savings program when suffering a family health shock than the uninsured. Particularly noteworthy is the unusually large coefficient on the interacted term of Group C and family health shock. This curious result most likely results from Group C having the lowest frequency of family health shocks of all the insurance groups.

When testing the equality of the coefficients that indicate which insurance group individuals are enrolled in, it is possible to reject the hypothesis that all the insurance groups have the same effect at the 99.9 percent significance level in each specification. Since this statistic tests equality in periods of good health, it posits significant wealth effects between insured individuals in differing groups. Given the disparity between the indigent members of Group A and the wealthy members of Group D, this is a reasonable result.

Yet when testing the equality of the interacted terms, most specifications fail to reject the null hypothesis that all individuals incur the same impact. In principle, this finding could mean that the quality of care is uniform between the differing insurance groups, though the literature on the Chilean health system suggests otherwise. A more plausible interpretation is that the variable of family health shocks accurately

reflects the cost of receiving a health shock, that is equally devastating regardless of the income level of the family.

Finally, we note in Specification (4) that when family health shocks are aggregated, the coefficient is insignificant once again. As in Table 3.11, this result is caused by the differing effects of family health shocks on the insured and uninsured. Yet once again, Specification (4) drives home the significance of insurance status reducing the propensity of individuals to fail to contribute to the mandatory pension system. In each of these models, we find that enrollment in the state-run health insurance program increases the likelihood of contributing to one's pension program. In addition, we find that when individuals receive health shocks, the reduction in income is frequently large enough to cause individuals to withhold contributions, with the notable exception of the uninsured. Therefore, it appears that the cutoff value for the insured is larger than that of the uninsured, indicating that the condition in Equation 3.16 holds and that while individuals with health insurance are more likely to save for retirement, the two goods are substitutes when individuals suffer either self or family health shocks.

3.6.2 Testing for Robustness

In Section 3.5, the issue of endogeneity in the specification of this model was discussed in the context of selecting pre-determined regressors. However, in an effort to avoid endogeneity resulting from a misspecification of the model, it is useful to test for the model's robustness utilizing an instrument variable regression approach.

This technique cannot be utilized when fitting a survival analysis model. Hence, it is necessary to redefine the dependent variable that we are examining. In the tables that follow, the analysis will utilize a generalized method of moments framework to fit a linear estimation on the proportion of time he fails to contribute to his pension program out of the entire time the individual is observed in the data. Note that since

this is a linear regression model, the estimated coefficients cannot be interpreted as above. Now they represent the unit marginal change in the proportion of time an individual fails to contribute to his pension program.

For the GMM specifications, we instrument for health shocks and insurance status by using the following variables:

- Lagged insurance status and lagged health indicators,
- Indicator variables that note whether the individual smokes, drinks, or has been to a doctor in the past year,
- Indicator variables that note whether the individual needs help performing daily tasks, is obese, disabled, or injured from work, and
- Continuous measures of Body Mass Index (BMI) and expected time until retirement.

Examining each instrument in detail, we find that many of them are significantly correlated with individuals' insurance and health status. As presented in Table 3.13, we note that insurance status is significantly correlated with its lagged value, lifestyle choices such as drinking and smoking, whether one has gone to the doctor in the past year, whether one has been injured on the job, expectations about the future, and the measure of an individual's BMI.

For both kinds of health shocks, we find that the indicator variables that signal whether the individual has gone to the doctor in the past year, whether the individual needs help performing basic household tasks, and whether the individual is disabled are all highly statistically significant. Furthermore, we find that the variable that indicates whether an individual smokes is correlated with personal health shocks but not family health shocks.

In each specification, the independent variables are similar to those in the previous subsection. However, due to the limiting number of potential instruments, we are unable to estimate coefficients for the interacted insurance status and health shock regressors as before; therefore, for the sake of comparison, the results depicted in Tables 3.14 and 3.15 are most similar to Specification (4) found above.

In Specifications (1) and (2) of Table 3.14, we examine the effect of personal health shocks on the percentage of total time in which the individual fails to contribute to the mandatory pension program. As depicted in Table 3.10, we find, once again, that controlling for household per capita income and age, individuals enrolled in the state health insurance program are less likely to fail to contribute to the required pension program.

Unlike in Table 3.10, the personal health shock is not statistically significant in explaining the variation of time spent failing to comply with the mandatory pension law. However, an examination of the strength of the instrument variables indicates that there is little explanatory power to instrument for health shocks in any specification. Although we were unable to estimate the survival time analog for Specifications (3) and (4) in the survival analysis framework, we find that even when disaggregating the different insurance groups, each group is individually and negatively correlated with the share of noncompliance time — further bolstering the findings in the previous subsection.

Returning to our analysis of family health shocks, we note in Table 3.15 that much as in Table 3.14, the effect of family health shocks is statistically insignificant. However, in each specification, enrollment in health insurance increases the amount of time than an individual complies with the mandatory pension contribution requirement. In every case, we can support the findings in the proportional hazard model when examining the role of health insurance. To sum up, by turning to the instrument variable regression framework, it has been possible to lend further sup-

port to the initial findings of this paper, while seeking to further reduce the potential for endogeneity to bias our findings.

3.7 Analysis and Conclusions

The aim of this paper has been to examine the effects and interactions of Chile's dual social protection requirements. By requiring all individuals both to contribute to the defined-contribution retirement savings program and to purchase health insurance, Chile attempts to enforce two distinct and costly social protection mandates. Yet given the constraints on the government of Chile's ability to monitor citizens' compliance, many individuals fail to comply with these two requirements.

While the research presented in this paper focuses on the role that extreme health shocks plays on decisions to contribute to the mandatory pension program, savings and retirement decisions are almost certainly influenced by minor health shocks as well. However, in the case of minor health shocks, individuals would conceivably continue contributing to their pension program while allowing their health insurance status to lapse temporarily. Because the public health insurance system has open enrollment, minor health shocks may not force people to sacrifice long-term planning for immediate health costs since the economic impact and expected duration of minor health shocks is certainly less damaging than the health shocks examined in this paper. With this in mind that future research will expand the scope of analysis to differing forms of health shocks.

Furthermore, as the data management of the EPS improves through new iterations of the survey, it will soon be possible to track individual household members between the surveys instead of just the representative household respondent. Hence, the unit of analysis can move from household health shocks to individual health shocks that control for the familial role of the individual. It will be possible to explore whether children's health shocks prompt a similar shift from long-term re-

tirement savings as adult household members' health shocks.

An additional implication of this paper's findings is that the ability to create long-run macroeconomic savings profiles to model the intergenerational sustainability of pension programs must take into account the prevalence of health impacts. Macroactuarial forecasts on pension programs' financial wellbeing currently incorporate, in addition to standard macroeconomic variables, demographic factors such as working-age population growth, birth and mortality rates, and labor participation rates. However, this paper contends that labor participation and pension compliance rates are, in part, a public health issue that can be altered by the quality and availability of health insurance, in addition to the morbidity rates associated with extreme negative health shocks at the personal and household level. While incorporating household-level health data into macroactuarial forecasts is an unreasonable level of granularity, it is a natural extension of this paper that public health trends should be incorporated into future research in generating individual and societal savings projections when examining the sustainability of pension systems — especially those systems such as defined-contribution pensions with incomplete monitoring that characterize middle income countries' social protection programs.

Since health shocks generally pose a more immediate risk than securing income for one's retirement, it is not surprising that the level of compliance is greater for the mandatory health insurance requirement than the mandatory pension contribution requirement.[15] Indeed, when surveying individuals who fail to contribute to their pension program on what steps would increase the likelihood of contributing, 68.6 percent of respondents in the EPS (2006) proposed linking the benefits of compliance to health benefits. The analysis within this paper is partially motivated by exam-

[15] The differing degrees of noncompliance are demonstrated by noting that 33.0 percent of all individuals fail to comply with the mandatory pension contribution requirement during the period they are observed while only 14.1 of individuals are uninsured during the same period. The difference between these two statistics is significant at the 99.9 percent level of significance.

ining the potential impact of linking pension contributions and health insurance by measuring the interactions between these two mandatory programs.

We have seen that individuals who purchase mandatory health insurance are more likely to comply with the requirement that they contribute to their pension programs regularly. In this instance, it is possible to consider these individuals generally law-abiding, as they comply with the two requirements imposed on them from the government of Chile. In contrast to this group, those relatively wealthy individuals who opt out of both the public or private Chilean insurance programs are generally non-compliant with the law. It is worth reiterating that this noncompliance decision is rational and forward-looking, as demonstrated by the models presented in Section 3.3. In other words, these individuals are not profligate scoff-laws, but rather utility maximizers who choose not to comply.

Yet while we find significant positive correlation between the decision to enroll in health insurance and the decision to comply with the mandatory pension contribution requirement, this relationship begins to break down when individuals become ill. In that event, they are much less likely to contribute to their pension program even when insured. Furthermore, the likelihood of an insured individual complying with the mandatory pension contribution requirement while suffering a health shock is statistically indistinguishable from the likelihood of an uninsured individual complying, *ceteris paribus*. This finding holds when individuals suffer a family health shock as well.

Since Chile's public insurance program covers between 80-100 percent of all medical costs, contingent upon the income of the insured individual, the failure to contribute one's pension program when suffering an extreme health shock most likely is caused by the foregone wages associated with illness rather than the cost of medical treatment. In the case of individuals suffering a personal health shock, the decision not to contribute to the mandatory pension system is probably twofold: individu-

als may be unable to afford the 10 percent of wages that must be contributed to the pension program; and individuals suffering from an extreme health shock may choose not to contribute since the health shock reduces the probability that they will reach retirement age. While it is impossible to differentiate between these two causes, it is worth pointing out that only the budget constraint-hypothesis holds when individuals suffer a family health shock.

By recalling that family health shocks significantly reduce the likelihood of complying with the mandatory contribution requirement, it is possible to argue that extreme health shocks have a substantial income effect. If pension compliance and health insurance benefits were linked, individuals who suffered an extreme health shock that substantially reduced household income may be forced to fail to purchase health insurance, compounding the cost of the health shock.

Even when examining the effect of family health shocks on pension compliance by insurance policies, as we do in Table 3.12, we find that even the wealthiest insured individuals are less likely to comply when suffering an extreme health shock. This indicates that wealthier individuals — those most at risk of failing to purchase health insurance — would be less likely to participate in Chile's retirement system if pension compliance and access to health insurance were linked.

Ultimately, we find that under normal circumstances health insurance complements retirement savings. The model presented in this paper proposes two mechanisms to account for this relationship. First, individuals who purchase health insurance are less likely to receive an extreme health shock and therefore more likely to have sufficient income to save for retirement. Second, individuals with health insurance have greater access to medical treatment and therefore more likely to live to retirement age, increasing the need for retirement savings.

It is only when individuals suffer an extreme health shock that the behavior of law-abiding Chileans who comply with the dual mandates of health insurance and

pension contribution become indistinguishable from those who fail to comply with the health insurance mandate due to the negative income shock associated with an extreme health shock. When this occurs, individuals rationally choose not to contribute to their pension program, highlighting that when illness forces individuals to choose between these two complements, they will fail to comply with the mandatory retirement savings requirement.

Table 3.1: Insurance Coverage by Policy

Public Insurance Group	Share of Liability Covered
Group A	100%
Group B	100%
Group C	90%
Group D	80%

Source: Sapelli (1999)

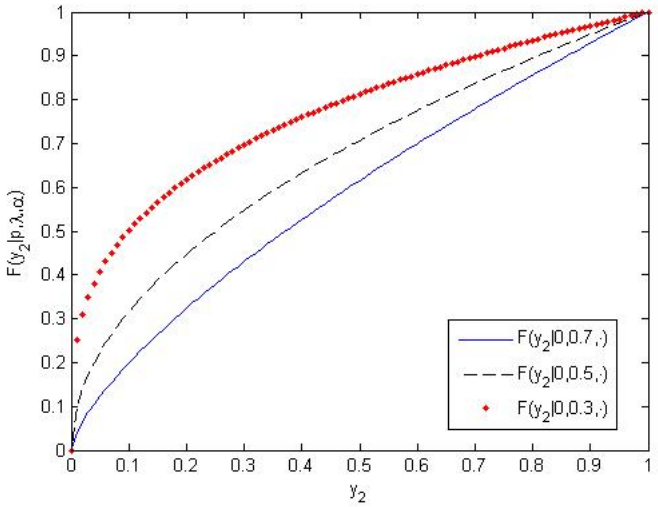

FIGURE 3.1: CDF of y_2 for Uninsured Individual with Health λ

Table 3.2: Expected Values of y_2 Given λ and $p = 0$

λ	$E(y_2\|p = 0, \lambda, \cdot)$
0.7	0.42
0.5	0.36
0.3	0.23

Table 3.3: Expected Values of y_2 Given $p = 1$, $\lambda = 0.3$, and α

α	$E(y_2\|p = 1, \lambda = 0.3, \alpha)$
0.7	0.45
0.3	0.34
0.0	0.23

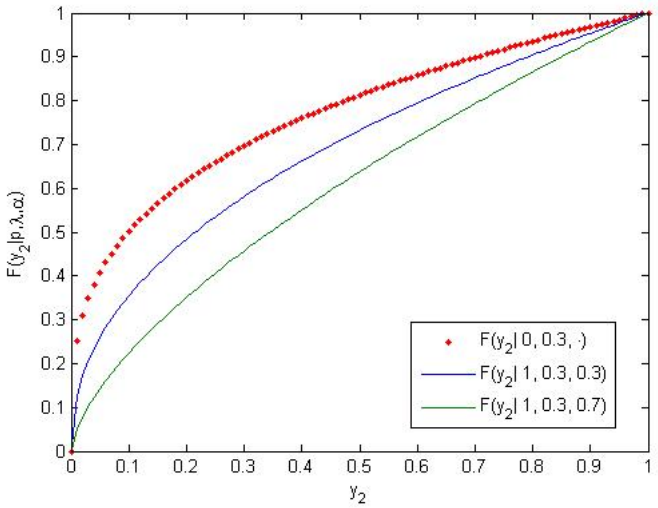

FIGURE 3.2: CDF of y_2 for Individual ($\lambda = 0.3$) With and Without Insurance α

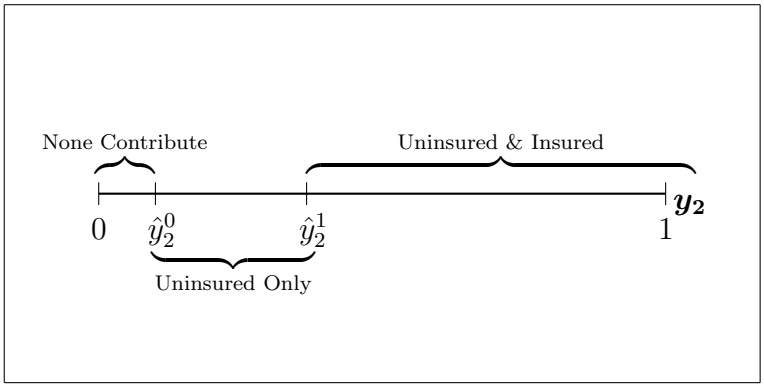

FIGURE 3.3: Pension Contribution Regions on the y_2 Numberline when $\hat{y}_2^1 > \hat{y}_2^0$

Table 3.4: Insured vs. Uninsured Individuals

	Uninsured	Insured	T-Test
Age	33.46	35.67	2.37
HH Size	5.27	5.30	0.18
HH Income per capita	1,245,675	769,181	2.97
Personal Health	2.79	2.74	1.73
Family Health	2.72	2.70	1.08
Obs	934	6,137	
Weighted Proportion	13.6%	86.4%	

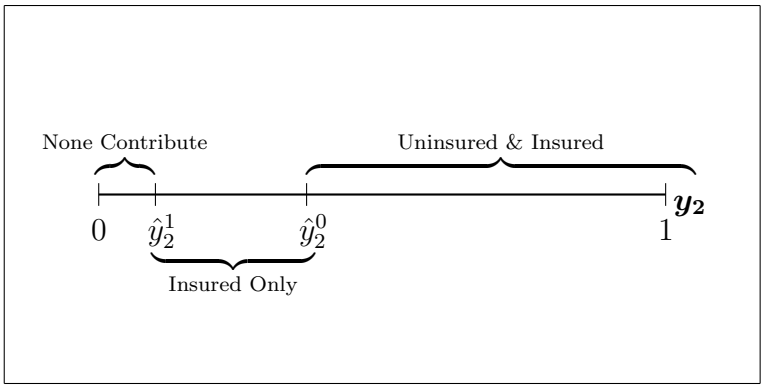

FIGURE 3.4: Pension Contribution Regions on the y_2 Numberline when $\hat{y}_2^1 < \hat{y}_2^0$

Table 3.5: Insured vs. Uninsured Individuals by Insurance Group

Insurance	HH Income per capita	Weighted Share	T-Test
Uninsured	1,316,923	13.6%	–
Group A	547,928	26.6%	4.01
Group B	762,545	29.9%	2.80
Group C	775,573	16.8%	2.75
Group D	981,187	13.1%	1.64

Table 3.6: Personal Health Shocks by Insurance Group

Insurance	Number of Personal Health Shocks	Obs	Weighted Incidence of Personal Health Shock
Uninsured	8	383	1.7%
Group A	45	854	4.7%
Group B	31	1220	1.8%
Group C	12	760	1.0%
Group D	9	545	1.0%

Pearson $\chi^2(4) = 36.13$

Table 3.7: Family Health Shocks by Insurance Group

Insurance	Number of Family Health Shocks	Obs	Weighted Incidence of Family Health Shock
Uninsured	63	383	16.8%
Group A	150	854	16.7%
Group B	172	1220	12.7%
Group C	92	760	11.8%
Group D	62	545	14.4%

Pearson $\chi^2(4) = 12.94$

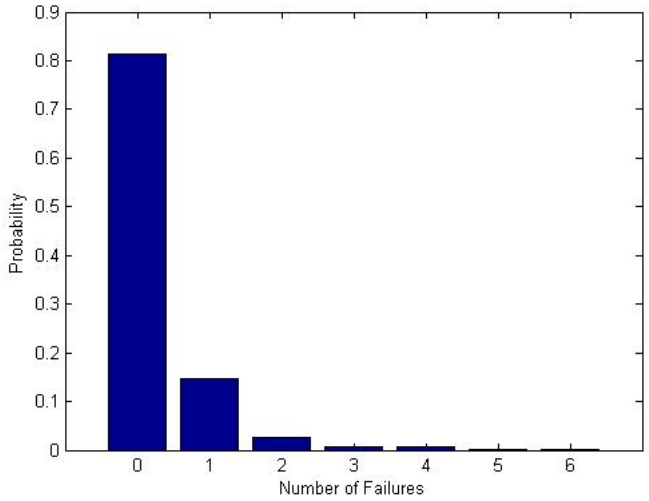

FIGURE 3.5: Empirical PDF of the Instances an Individual Fails to Contribute.

Table 3.8: Transition Probabilities for Insurance Status

	2004	
2006	Uninsured	Public
Uninsured	30.7%	7.6%
Public	69.3%	92.4%

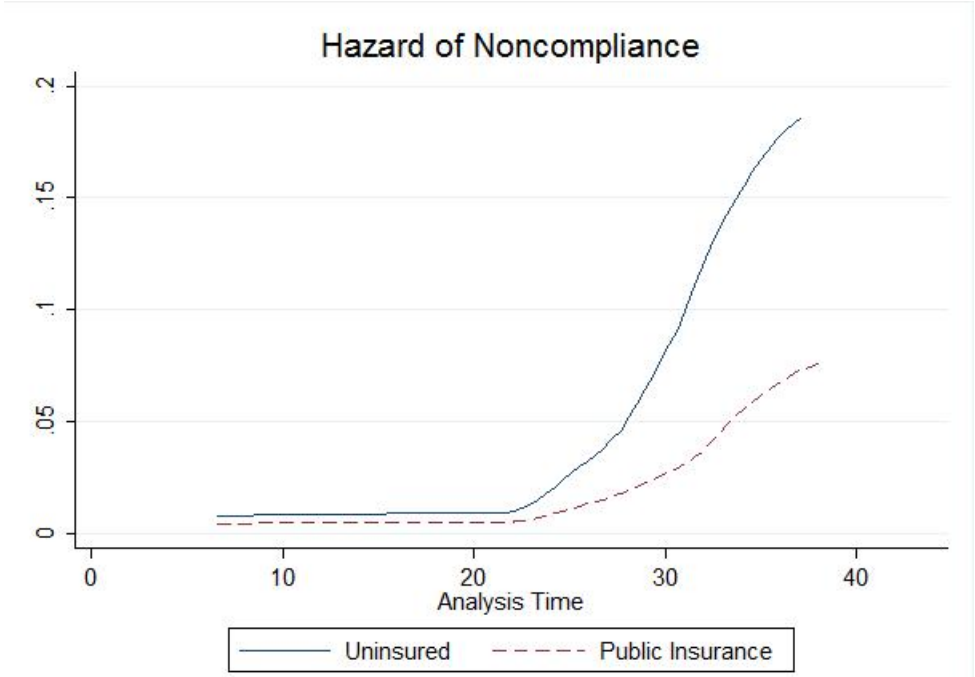

FIGURE 3.6: Hazard Rate by Insurance Status

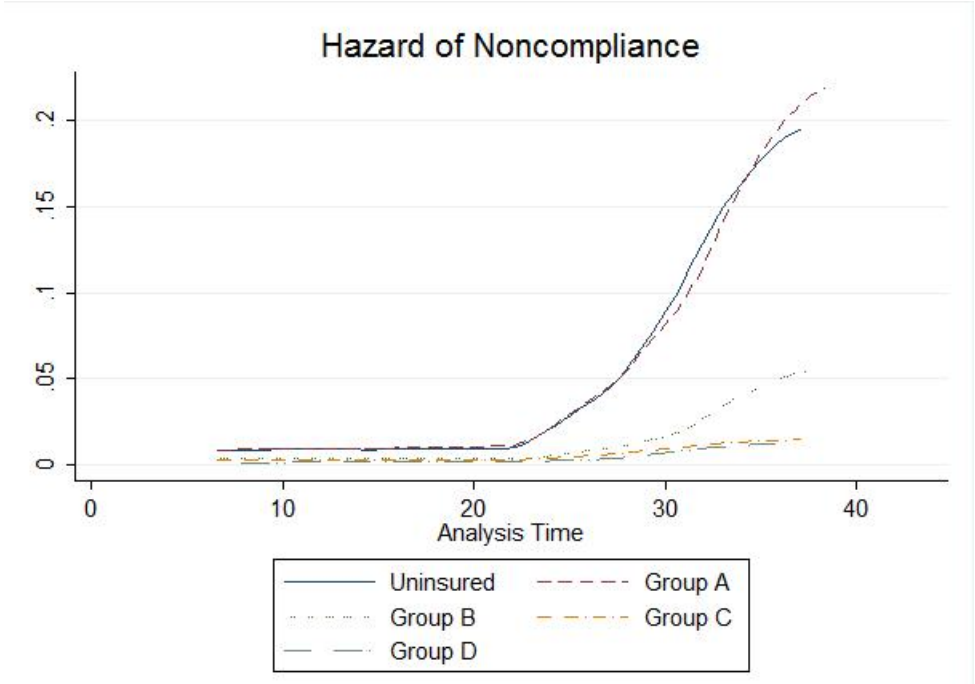

FIGURE 3.7: Hazard Rate by Insurance Group

Table 3.9: Proportion of Individuals by Insurance

	2004	
2006	Uninsured	Public
Uninsured	8.0%	5.6%
Public	18.2%	68.2%

Table 3.10: Proportional Hazard Model on Time to Noncompliance: Personal Health Shocks

	(1)	(2)	(3)	(4)
VARIABLES	Base	Lag	Lag Equal	Lag Equal
Labor Income	0.997***	0.997***	0.997***	0.997***
	(0.000882)	(0.000880)	(0.000970)	(0.000882)
Age	0.838***	0.843***	0.815***	0.837***
	(0.0373)	(0.0353)	(0.0436)	(0.0373)
Age^2	1.002***	1.002***	1.002***	1.002***
	(0.000561)	(0.000526)	(0.000656)	(0.000561)
Insured	0.364***	0.582***	0.344***	0.373***
	(0.0686)	(0.0962)	(0.0824)	(0.0685)
Personal Health Shock	1.235	1.642	0.668	2.380***
	(0.469)	(0.647)	(0.188)	(0.667)
Ins X Personal Health Shock	2.395*	1.733	4.602***	
	(1.147)	(0.824)	(1.924)	
Region-Specific Dummies	Yes	Yes	Yes	Yes
Observations	3,514	3,514	3,006	3,006

s.e. in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 3.11: Proportional Hazard Model on Time to Noncompliance: Family Health Shocks

VARIABLES	(1) Base	(2) Lag	(3) Lag Equal	(4) Lag Equal
Labor Income	0.997***	0.997***	0.997***	0.997***
	(0.000868)	(0.000865)	(0.000910)	(0.000949)
Age	0.835***	0.842***	0.801***	0.811***
	(0.0371)	(0.0348)	(0.0403)	(0.0426)
Age^2	1.002***	1.002***	1.003***	1.003***
	(0.000562)	(0.000522)	(0.000620)	(0.000649)
Insured	0.319***	0.513***	0.243***	0.338***
	(0.0623)	(0.0881)	(0.0497)	(0.0762)
Family Health Shock	0.450**	0.427**	0.187***	0.838
	(0.172)	(0.169)	(0.0997)	(0.226)
Ins X Family Health Shock	2.494**	2.905**	6.834***	
	(1.091)	(1.293)	(4.017)	
Region-Specific Dummies	Yes	Yes	Yes	Yes
Observations	3,514	3,514	3,006	3,006

s.e. in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 3.12: Proportional Hazard Model on Time to Noncompliance: Family Health Shocks

VARIABLES	(1) Base	(2) Lag	(3) Lag Equal	(4) Lag Equal
Labor Income	0.998**	0.998***	0.999	0.998
	(0.000869)	(0.000803)	(0.000964)	(0.00110)
Age	0.887***	0.846***	0.837***	0.853***
	(0.0393)	(0.0325)	(0.0421)	(0.0487)
Age^2	1.001**	1.002***	1.002***	1.002
	(0.000561)	(0.000491)	(0.000632)	(0.000717)
Insured Group A	0.786	1.051	0.717	0.969
	(0.191)	(0.197)	(0.157)	(0.216)
Insured Group B	0.219***	0.511***	0.166***	0.221***
	(0.0541)	(0.107)	(0.0381)	(0.0573)
Insured Group C	0.168***	0.282***	0.0794***	0.176***
	(0.0551)	(0.0818)	(0.0372)	(0.0966)
Insured Group D	0.0451***	0.241***	0.00725***	0.0114***
	(0.0161)	(0.101)	(0.00544)	(0.00733)
Family Health Shock	0.426**	0.422**	0.155***	0.768
	(0.175)	(0.163)	(0.0882)	(0.294)
Ins A X Family Health Shock	2.128	2.638**	7.897***	
	(1.050)	(1.230)	(5.211)	
Ins B X Family Health Shock	3.244**	1.232	7.129***	
	(1.784)	(0.674)	(5.257)	
Ins C X Family Health Shock	4.628*	7.321***	33.56***	
	(3.653)	(4.661)	(30.69)	
Ins D X Family Health Shock	7.568**	2.727	11.45*	
	(6.365)	(2.019)	(15.96)	
Region-Specific Dummies	Yes	Yes	Yes	Yes
Observations	3,203	3,514	1,578	1,578
Test Equality of Groups (p-value)	0.00	0.00	0.00	0.00
Test Equality of Ins X Shock (p-value)	0.28	0.05	0.26	

s.e. in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 3.13: Estimated Correlations between Instrumented and Instrument Variables

	Instrumented		
Instruments	Insured	Personal Health Shock	Family Health Shock
Lag Insured	0.312***	0.010	0.019
Lag Health	-0.001	0.046**	0.015
Smoker	-0.044**	-0.026	0.002
Drinker	-0.047**	-0.054***	0.009
Doctor	0.113***	0.078***	0.084***
Needs Help	0.022	0.107***	0.057***
Obese	0.034*	0.013	0.005
BMI	0.018	-0.012	-0.016
Disabled	0.020	0.175***	0.081***
Time to Retirement	-0.042*	0.002	0.014
Work Injury	0.047**	0.016	0.020

Table 3.14: OLS and IV Regressions on Share of Noncompliance Time: Personal Health Shocks

VARIABLES	(1) OLS	(2) GMM/IV	(3) OLS	(4) GMM/IV
Personal Health Shock	0.108	0.678	0.0327	-0.220
	(0.0775)	(0.696)	(0.0691)	(1.360)
Insured	-0.480***	-1.176***		
	(0.0683)	(0.264)		
Insured Group A			-0.238***	-1.028
			(0.0846)	(1.040)
Insured Group B			-0.492***	-1.207*
			(0.0717)	(0.678)
Insured Group C			-0.538***	-0.935*
			(0.0683)	(0.551)
Insured Group D			-0.570***	-1.553*
			(0.0663)	(0.815)
Age	-0.0318***	-0.0205**	-0.0248***	-0.0223*
	(0.0101)	(0.00937)	(0.00952)	(0.0118)
Age^2	0.000424***	0.000275**	0.000337***	0.000296**
	(0.000126)	(0.000112)	(0.000119)	(0.000144)
Income Per Capita	-4.58e-07	-5.48e-07	-3.91e-07	-2.23e-07
	(2.99e-07)	(4.35e-07)	(2.92e-07)	(5.85e-07)
Constant	1.184***	1.761***	1.062***	1.698**
	(0.208)	(0.305)	(0.196)	(0.743)
Region-Specific Dummies	Yes	Yes	Yes	Yes
Observations	1,952	1,239	1,952	1,239
R-squared	0.208		0.293	

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3.15: OLS and IV Regressions on Share of Noncompliance Time: Family Health Shocks

VARIABLES	(1) OLS	(2) GMM/IV	(3) OLS	(4) GMM/IV
Family Health Shock	-0.00296	0.404	0.00213	0.269
	(0.0415)	(0.252)	(0.0414)	(0.263)
Insured	-0.481***	-1.383***		
	(0.0679)	(0.331)		
Insured Group A			-0.238***	-1.259*
			(0.0843)	(0.733)
Insured Group B			-0.493***	-1.272**
			(0.0715)	(0.569)
Insured Group C			-0.539***	-1.220***
			(0.0681)	(0.416)
Insured Group D			-0.571***	-1.599**
			(0.0661)	(0.756)
Age	-0.0319***	-0.0179	-0.0247***	-0.0179
	(0.0100)	(0.0111)	(0.00936)	(0.0126)
Age^2	0.000427***	0.000245*	0.000337***	0.000241
	(0.000125)	(0.000130)	(0.000116)	(0.000155)
Income Per Capita	-4.62e-07	-8.17e-07	-3.91e-07	-6.24e-07
	(2.98e-07)	(5.95e-07)	(2.92e-07)	(6.74e-07)
Constant	1.185***	1.864***	1.060***	1.785***
	(0.206)	(0.350)	(0.193)	(0.553)
Region-Specific Dummies	Yes	Yes	Yes	Yes
Observations	1,952	1,239	1,952	1,239
R-squared	0.207		0.293	

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

4

Pension Liquidity, Access to Credit, and Housing: The Rationale of Noncompliance in Chile's Defined-Contribution Pension Program

4.1 Introduction

Chile's move to a defined-contribution pension system in 1981 ushered in a significant shift in the way countries and international institutions approached the provision of old-age retirement benefits (The World Bank, 1994; Gill et al., 2005). This much-remarked-upon program was, among other things, a meaningful ideological statement in favor of an ownership culture that encouraged its citizens to increase their engagement with the domestic financial sector of the economy. Viewed from this perspective, the adoption of a defined-contribution pension system is a public policy tool to marry the futures and fortunes of the Chilean public to the greater economy.

The policymakers who worked to promote an ownership culture by linking private citizens' retirement benefits to the financial market did not stop there. They also embraced the longstanding tradition of Chile's involvement in the real estate market

to encourage individuals to purchase homes. Policymakers bolstered homeownership by encouraging savings, affordable home construction, and the development of the domestic mortgage market. Yet, as is often the case, these two government policies — both designed with the same economic and cultural aims — have the potential to undermine one another in practice. This paper examines how the common objectives of these two government programs designed to increase long-term personal savings work at cross purposes. Specifically, they encourage individuals to violate the letter of the law by withholding mandatory pension contributions in order to purchase a home.

The motivation and structure of Chile's defined-contribution pension program is straightforward to understand. As outlined in Chapter 3, Chile's publicly-mandated pension program is privately administered by Administradoras de Fondos de Pensiones (AFPs) that invest savings in a diverse set of assets. All employed individuals are required to contribute 10 percent of their income each pay period to an AFP of their choice. These funds are invested in various assets that includes domestic and international equity, government debt, and various private fixed income assets — including mortgage backed securities (Srinivas and Yermo, 1999). Policymakers believe that by investing long-term retirement savings in a variety of financial assets, pension funds will accumulate at a faster rate than if savings were limited solely to sovereign debt.

However, a major critique of this funded pension system framework stems from the question of individuals saving enough for retirement (Orszag and Stiglitz, 2001; Barrientos, 1996). There are several policies in place to mitigate this risk. Chile provides a minimum pension guarantee for individuals who contribute regularly to their pension programs. All pension funds are required to insure against business cycle fluctuations by saving excess returns in a separate fund to supplement future market under-performance. Yet, one of the greatest threats to the accumulation of

retirement savings comes from regulators' inability to enforce the mandatory pension contribution requirement in the presence of competing demands for individual resources, including expenditures on durable goods.

In the sections that follow, this paper will demonstrate how rational forward-looking individuals may choose to disregard the mandatory pension contribution requirement in an effort to increase their short-term liquid savings prior to undertaking significant housing decisions. In what follows, we shall see how the competing institutions that underlie the Chilean pension system and mortgage market drive individuals to choose between compliance and homeownership, potentially undermining the ability of individuals to save for retirement. However, before addressing these points, we will further explore the institutional context of Chile's housing market to better understand individuals' economic decisionmaking.

4.2 The Chilean Housing Market

Studying the Chilean housing market calls for examining the role of the government. Throughout the 20th century, the government of Chile was intimately involved in the development and financing of the domestic housing market (Lira, 1994; Cummings and DiPasquale, 2002). In an effort to promote homeownership, Chilean policymakers created a series of subsidies and programs to encourage individuals to purchase property and bolster banks' lending to finance housing purchases.

Alvayay and Schwartz (1997) present a detailed summary of Chile's policies and government programs designed to encourage homeownership. Yet, despite these programs, the key impediment for individuals planning to purchase a home is sizeable down payment required by lenders. Government restrictions on lending limit mortgages to between 50 and 75 percent of the total value of a home. These restrictions create a credit constraint that potential homeowners can overcome only by accumulating a sizeable down payment. Policymakers employ two main avenues to expand

homeownership while maintaining the strict requirements on lending: Protecting lenders from inflation and providing public sector subsidies to low-income families.

All housing-related financial assets are valued in a theoretical currency called *unidad de fomento* or UF, which is updated daily to incorporate fluctuations in the consumer price index over the previous month. By valuing homes, mortgages, and collateralized mortgage obligations called *letras hipotecarias* or LHs in UFs rather than Chilean pesos, the Bank of Chile protects lenders in the housing market a real rate of return insulated from inflation. Since inflation favors debtors — as they can repay a loan with devalued currency — the expansion of UF-denominated housing loans shifts the inflationary risk from private banks and spurs private-sector lending.

However, Lira (1994) and Pardo (1999) challenge the efficiency of the Chilean housing finance system. They highlight the large interest-rate spreads that banks capture between regularly supplying mortgages at about 9 percent and selling LHs at 6 percent to pension funds. In particular, Lira (1994) argues that since LHs are primarily held by pension funds, there is a legislated inefficiency when individuals who are purchasing a home must continue to contribute to their pensions while borrowing from private banks. Especially since the private sector mortgage will be purchased by the pension fund at a later date.

For poorer households who are unable to acquire housing through the private sector, policymakers offer various subsidies to facilitate homeownership. Some programs seek to expand the available units of low-income housing by auctioning off the right to build housing developments to private construction companies. The government then arranges purchasers for the increased housing stock by offering mortgages to low-income households at interest rates on par with those offered by the private sector. However, in analyzing these urban development programs, Cummings and DiPasquale (2002) argue that these programs introduce sizeable inter- and intra-regional distortions that reduce labor mobility and concentrate low-income home-

owners in newly constructed large housing developments located on the periphery of urban areas.

More traditional direct subsidies are also offered by the government to lower-income households purchasing their first home (Pardo, 2001). For instance, a household can apply for a subsidy contingent upon household income, the location and value of the house being purchased, and the amount the household has already saved in dedicated savings accounts administered by banks. These dedicated accounts are designed specifically to bolster personal savings for home purchases. Subsidy administrators favor the accumulation of savings in these accounts when choosing which households will receive subsidies, providing an incentive for personal savings for housing that is unique in Latin America (Ferguson, 2004).

Broadly speaking, the Chilean government's primary goal in developing the domestic housing market is to expand homeownership in general and for lower-income households specifically. Ruprah and Marcano (2005) find that this policy has been generally successful, as homeownership increased from 61.5 percent of the population in 1990 to 70.3 percent in 2003. They highlight that although this expansion has affected individuals of all income levels, the greatest gains were in the lower two quintiles. Having started from a baseline of both regionally- and globally-high homeownership, policymakers offered a collection of subsidies, incentives, and financial mechanisms to simultaneously increase the housing supply, individual savings, and the attractiveness of collateralized mortgage obligations to institutional investors. In the process, the expansion of the private pension system was tied to the deepening of the housing finance market.

4.3 Summary Statistics of Chile's EPS Dataset

4.3.1 Homeowners and Renters

Turning our attention to the 2006 survey of households, it is possible to begin building a comprehensive picture of how households differ in their housing decisions, access to capital, and degree of pension compliance. Throughout our analysis, we will limit our observations to males of the sample, who number 7,071 out of a total of 14,425 obserations (49%). Our sample is further restricted by limiting it to individuals who are less than 60 years of age — far from retirement at 65 — and for whom we have employment data.

For starters, it is possible to differentiate households who own their own home from those who rent or live in a home provided by family members or employers. Turning our attention to Table 4.1, we see some key facts about the attributes that differentiate homeowners from renters.

An initial look indicates that our data mirrors the findings in Ruprah and Marcano (2005) in which 70.3 percent of households are homeowners in 2003. Homeowners, on average, are older than renters. In addition, total household savings are larger for homeowners than for renters. These findings are most likely intertwined, due to Chile's strict limits on mortgage financing which require individuals to save about 25 percent of the value of a home in order to purchase property, thus requiring individuals to spend more time working and saving if they want to purchase a home than if they want to rent, *ceteris paribus*.

What we shall call the "full credit" variable is a key indicator that will be referenced multiple times below, so it deserves our attention now. In the EPS, all individuals are asked the question, 'In the event of an emergency in which you had to borrow six months' income, what percentage of this amount could you borrow within a month?' On average, the respondent claims that he can borrow 33.6 per-

cent of his six-month income.[1] Throughout this paper, the response to this question is interpreted in two ways: First, we define 'Full Credit' as an indicator variable for when individuals claim access to 100 percent of their six-month salary; secondly, we define 'Partial Credit' as a continuous variable for the percentage of six-month's salary defined between 0 and 100 percent. In the analysis that follows, individuals' responses to this question is treated as a measure of access to capital and financing. Judging from Table 4.1, we note that homeowners have far greater access to credit than renters — an understandable result when medium- and high-income households who cannot qualify for state subsidies must rely on a mixture of high savings and the private mortgage market when purchasing a home.

Finally, it is worthwhile to draw attention to where renters and homeowners coincide. In general, household size and total household income per capita are indistinct between these two groups of individuals. This curious finding suggests that the main difference between renters and buyers lies in the ability to accumulate savings and find credit to purchase a home. Furthermore, the difference in savings between the two groups are probably similar despite the similar earnings profiles speaks to the benefits of accessible credit, as individuals with access to credit are able to invest money in dedicated housing-related savings accounts without fear of having to accommodate short-run income shocks.

Throughout the period during which individuals in the sample are under observation, they may move from home to home or from renting to owning their dwelling. Since this paper seeks to study the rate of pension compliance when an individual is *planning* on moving to a new home, we must identify when an individual is anticipating purchasing a home. To this end, we resort to a key question in the EPS dataset. In each round, the interviewer asks the head of the household how long

[1] The standard deviation for this response is 0.437. However, since most individuals respond with either 0 or 100 percent, the distribution is far from normal and is largely concentrated at the endpoints of the support.

he has lived in his current home. The 2006 EPS dataset allows us to observe the contribution decisions of an individual over the period 2004–2006. However, if we only controlled for individuals who moved into a new home over this period, we would risk biased results stemming from the endogeneity resulting from a household simultaneously choosing whether to move and whether to contribute to their pension program. Therefore, we will identify households that are planning on purchasing a new house by their response in the 2009 round of the EPS. In the analysis that follows, we examine individuals who bought a new home over the period 2006–2009, since they were *planning* to purchase a home over the period of 2004–2006.

4.3.2 Access to Credit

Upon purchasing a house with mortgage financing, individuals who are already required to contribute to their pensions become doubly-constrained by their mortgage obligations. Of course, there are lasting consequences to defaulting on one's mortgage, so individuals whose finances are beholden to these twin mandatory payment structures have a strong incentive to insure their income through some form of credit.

When asked if individuals would have access to six months of emergency credit if needed, we observe that 26.1 percent of respondents answer positively. Moreover, individuals' response to this question does not change significantly if they are planning to purchase a new home. To be clear, respondents to the question about credit are answering in 2006. When we break down these responses by whether the individual will buy a home over the period 2006–2010, we find that individuals who bought a home over this period cannot be distinguished from individuals who did not buy a home solely by their reported access to credit.[2]

However, individuals' access to credit does vary once we examine those who are

[2] We cannot reject the null hypothesis that the two populations have the same mean response ($t = 0.94$).

planning to purchase a home. In this case, we find statistically significant variance in access to six months' wages of credit, which is correlated with the method of financing the household will utilize to purchase the new house. These results are presented in Table 4.2, most notable stemming from Pearson's test statistic ($\chi^2(4) = 14.7$).

Table 4.2 also shows that individuals who will purchase a home outright without financing make up about 20 percent of all homebuyers. Interestingly, when compared with the 44 percent who purchase their home using mortgage financing, the difference in access to credit between the two groups is not statistically significant.[3] This counters the hypothesis that those who purchase a home in cash do so because they do not have access to credit.

A large number of new homeowners acquired their home through "other" sources of financing. By and large, these individuals bought a homes through one of the many government subsidy programs. While the analysis in this paper focuses on homeowners who purchase their homes without utilizing government subsidies, it is necessary to briefly examine the large percentage of individuals who received subsidies. The main findings are depicted in Table 4.3.

Strikingly, income per capita and savings appear to be indistinct from one another. To make sense of this result, one must differentiate between subsidy programs. When looking at the distribution of income in these two groups, one finds that the mean of the incomes are noticeably distinct from one another. However, turning to Figure 4.1, we can graph the distribution of total household income per capita by whether the household received a subsidy. Clearly, the long right tails of the distributions are very similar. This overlap — which also exists in the distribution of total household savings — accounts for the failure to reject the hypothesis that the means differ by subsidy status. The long right tails of income variables are common

[3] Comparing those who purchased their home using through mortgage financing and in full, the two are statistically indistinguishable ($t = 1.53$). This result is robust when testing the equivalency between the two groups when analyzing gradations of access to credit.

in economies with high income-inequality.[4]

There is a driving forces that leads key economic variables such as per capita income and savings to overlap despite the difference in subsidy status — there are numerous kinds of subsidies offered by the Chilean government. In short, subsidies can be broken into two groups; subsidies for low-income individuals to acquire publicly constructed housing and subsidies for middle-income individuals to increase access to home financing. Of course, with only 119 observations in which a new homeowner has used a subsidy, any further subsidy-specific analysis runs the risk of reducing the data to too fine a level to compare between groups. However, since different subsidies target different income levels within Chile, analyzing income-related variables by whether one has or has not received a subsidy is too coarse to accurately observe the differing characteristics involved.

Yet, returning to Table 4.3, one observes that the proportion of individuals who have access to emergency credit to replace six-month's salary differs significantly between the two groups. By and large, individuals who purchase a house without a subsidy are more likely to have access to emergency credit than those who purchase a home via a subsidy. This should come as no surprise, since individuals who are able to purchase a home without government aid are doing so through a mix of personal savings and privately acquired mortgage financing—access to which is strongly correlated with access to credit from formal lending institutions.

Ultimately the aim of this paper is to examine the role that access to emergency credit plays once an individual has purchased a house and must continue to contribute to his pension plan. For this reason, the analysis presented in Section 4.4 will focus on individuals who are planning on purchasing a home in the next three years and do so without utilizing any government subsidies. The key question is how does

[4] Most estimates of Chile's Gini coefficient are about 0.54 (Solimano and Torche, 2008), which signifies a strong degree of income inequality and explains the distribution of income observed in Figure 4.1.

purchasing a home change one's ability to comply with the mandatory pension contribution requirement when controlling for credit. In order to assess this question, it is necessary to examine the baseline rate of compliance in the data before controlling for competing factors.

4.3.3 Pension Compliance

The Chilean government requires all individuals employed in the private sector to contribute a share of his or her income to an individual retirement account. With the notable exception of self-employed individuals who are absolved of this requirement and some public-sector employees — most notably in the military — who have access to alternative publicly provided pension systems, the majority of Chileans are required by law to contribute about 10 percent of their regular income to their pension program.

This mandatory contribution requirement is usually implemented by requiring formal-sector employers to withhold a portion of individuals' paychecks and contribute these funds to the pension fund of the employee's choosing. Although employers are responsible for collecting these funds, they do not contribute a share of the retirement funds directly, in contrast to the US social security program.

In practice, however, individuals are often able to approach their employers and request that their wages not be withheld. This is the main manner by which individuals avoid contributing to their pension program. Indeed, when asked about why they have failed to contribute to their pension program, 38 percent of those who have not regularly contributed to their pensions replied that they and their employer had decided not to comply with the contribution requirement, either at the employer's request, the employee's request, or as a mutual understanding. Essentially, due to the Chilean government's either inability or unwillingness to monitor businesses' compliance with the mandatory pension contribution program, individuals are able to opt

out of the legal requirement by reaching an understanding with their employer.

Understandably, individuals may choose to contribute during one portion of their work history and not during another. In this chapter, the data that are employed in the analysis that will follow are drawn from 2,907 individuals. As discussed earlier, it is possible to observe each individual as he moves from job to job. On average, we observe each person for about 32 months, during which he holds 1.5 jobs. However the total number of jobs an individual holds varies, and the median individual holds one job throughout the period of observation. The largest number of jobs that one observe individual holds is nine. Due to the independent variables that will be employed at various times throughout the analysis in this chapter, these summary statistics differ from those presented in Chapter 3.[5]

Since individuals can hold more than one job over the period of observation, the 2,907 unique individuals are represented by 4,307 records which correspond to each job/individual pairing in the sample. However even though we observe 4,307 records, the incidence of failing to comply with the mandatory pension contribution requirement is as low as 628 observed failures over the entire sample. This corresponds to about 0.22 failures per person, with the median number of failures observed per person being 0. Yet even though most individuals comply with the pension requirements throughout their period of observation, there is a large degree of variance in individuals' propensity to default on their pension requirements, with the largest number of failures to comply by a single individuals being 9 times during the period of observation.[6]

In Chapter 3, we showed how individuals will choose to contribute to their pension based upon their health or the health of those who live in their household. Here,

[5] For instance, the analysis in Chapter 3 excluded individuals who hold private health insurance who are included in the sample here, which accounts for the different number of observations.

[6] For a full treatment of the observed distribution of failures and the rate of recidivism, see Subsection 3.4.2 and Figure 3.5.

however, we will show how individuals can choose to disobey the mandatory pension contribution requirement due to their long-term financial planning and access to credit. Given the structure of the data described in Subsection 3.4.2, the most straightforward method to analyze the propensity to fail to comply with the mandatory contribution requirement is to estimate the hazard rate, which is interpreted as the marginal probability of failing to contribute to one's pension at time t conditional upon having contributed in the past.

Hazard rates vary and generally increase over time. This hazard rate can be estimated for the entire sample, as is presented in Figure 4.2. Note that at its maximum, the probability of failing to contribute to one's pension program conditional upon having contributed throughout the entire period of observation is slightly above 2 percent. This graphical depiction of the hazard rate is a broad measure of the sample-wide propensity to fail to comply with the the legal requirements that drive the Chilean pension plan. Yet, the hazard rate in Figure 4.2 is a naive specification, because it treats all individuals within the sample identically without distinguishing them by the key characteristics that were discussed above: whether the contributors are planning on purchasing a home and whether they have access to emergency credit.

As discussed above, it is likely that access to credit will play a significant role in a household's decisionmaking. While the government requires individuals to save a fixed proportion of their income in their pension plan every billing period, personal and household-level economic shocks always pose a risk that could jeopardize a household's financial security. Since the government does not provide a legal pathway for individuals who are planning a capital-intensive purchase to defer pension contributions, a major purchase can force an individual to choose between contributing to his pension program and investing in alternative long-term assets. An extreme case of this was studied in Chapter 3 when examining the tradeoff between purchasing

health insurance and saving for retirement.

Yet this tradeoff between financial diversification and pension compliance can be lessened by a household's access to credit. Conceivably, a household with access to credit can avoid legal risk of defaulting on its pension contributions by drawing on its credit to finance major purchases. To test this hypothesis, we turn to Table 4.4 to examine the relative hazard of noncompliance by estimating the hazard rates for two subsections of the total observed population — those with full access to credit and those without it.

Judging from Table 4.4, we see that those without access to credit are about 1.4 times more likely to fail to contribute to their pension than those who have access to credit equal to six months' salary. In short, Cox's relative hazard is interpreted as a bivariate odds ratio which compares the likelihood of failure—that is, noncompliance—over the variation of an individual-specific variable.[7] Furthermore, Pearson's test statistics confirms that access to credit is a strong predictor of an individual's likelihood to comply with his pension contribution requirements.

Another method of analysis that supports Table 4.4 and underscores its findings is portrayed in Figure 4.3. Just as in Figure 4.2, Figure 4.3 displays the estimated evolution of the hazard rate over time for the data. However in this case, Figure 4.3 divides the sample into two groups; those with access to six months' salary worth of emergency credit and those who do not have access to emergency credit. Dividing the sample in this way heightens the effect that access to credit has on individuals' ability to comply with the mandatory pension contribution requirement. As is evident from the figure, individuals with full access to credit are far less likely to default on their pension obligations than individuals without access to credit.

So far, we have shown that the ability of an individual to comply with the manda-

[7] Like all odds ratios, a point estimate greater than 1 indicates that an individual is more likely to fail and an estimate between 0 and 1 indicates that an individual is less likely to fail as the independent variable increases.

tory pension contribution requirement is affected by that individual's access to credit. Since pension savings are illiquid, every period in which someone chooses to comply with the savings requirement is an additional step towards committing present-day earnings to financing one's retirement. Since this commitment is irreversible, compliance comes with a liquidity cost in addition to the cost of foregone contemporaneous consumption. This liquidity cost is driven by the risk of economic shocks that could imperil short-run financial security — though this risk is mediated by access to credit. Hence, individuals with access to credit can commit their present-day earnings to retirement since they are able to finance alternative investment vehicles through the private credit market.

The key alternative investment vehicle that we are examining throughout this paper is housing. As stated above in Section 4.2, Chilean government policy has dedicated a large amount of institutional expertise and capital to encouraging homeownership for both political and economic reasons. However this investment has not been reflected in the institutional framework that shapes Chile's mandatory pension program. Due to the policies that regulate housing finance in Chile, individuals must amass a significant down payment prior to acquiring a mortgage when purchasing a home. For many individuals, it is likely that this down payment is difficult to save when contending with the government-mandated pension savings requirement. Therefore, we must be able to identify individuals who plan to purchase a home and assess whether they are more or less likely to fail to comply with the mandatory pension savings program.

It is helpful to consider two variables that indicate whether an individual is planning on purchasing a house. Since we are utilizing data from 2004–2006 to examine pension compliance, it is possible to incorporate data from the 2006–2009 EPS dataset to see if any of the households under observation in the 2006 round purchased a new home in the three years that follow that period of observation. In Table 4.5,

we examine two overlapping groups; individuals who purchased a new home in the period 2006–2009 and individuals who did so specifically through a new mortgage over the same period.

In both cases, we find that individuals who acquired a new home are more likely to comply with the mandatory pension contribution requirement than individuals who did not purchase a new home. This result—which is reproduced graphically in Figure 4.4—is not entirely intuitive, given the institutional framework discussed above and individuals' incentive to default on their pension obligations while saving for a new home. However referring back to Tables 4.1 and 4.3, it is clear that controlling for whether an individual is planning to purchase a house alone is a powerful yet indistinct variable to assess pension compliance. Purchasing large durable goods necessitates access to credit, as discussed above, and therefore any bivariate analysis that focuses housing purchases without simultaneously controlling for income and credit variables is bound to yield significant results that nevertheless muddy the analysis.

Therefore, it is necessary to expand our analysis to a multivariate setting. In the next section, we utilize a similar estimation strategy as presented in Subsection 3.5, which allows us to control for an individual's propensity to default on his mandatory pension contribution requirement while controlling for job- and individual-specific variables. It is important to note that many of the issues brought up in Chapter 3 that contend with the endogeneity of health insurance and pension compliance are happily avoided in the analysis that follows, since the decision to purchase a house takes place after the period in which the econometrician observes individuals' pension compliance. By regressing pension compliance on future values of homeownership, we effectively instrument for any potential endogeneity that would flow from the simultaneity of housing and pension decisionmaking.

4.4 Results

Cox's Proportional Hazard Model

As discussed in Chapter 3, the objective of utilizing a Cox's proportional hazard model framework is to analyze the time-to-noncompliance — which is defined as the time an individual is observed complying with the mandatory pension contribution requirement until that time in which he first fails to contribute—while controlling for job and personal attributes which vary over time and between individuals, respectively. Once again, it is important to highlight that the analysis conducted in this section is examining the contribution decisions of males between the ages of 18 and 59 who are employed in the private sector and who have not received any government subsidies directed at encouraging homeownership. These restrictions focus the analysis on non-governmental workers who are at least five years away from retirement and who make their housing decisions based upon the private sector home financing and credit markets.

Turning to the results from the proportional hazard model in Table 4.6, one notes that there are four specifications utilized to test the effects of purchasing a new home on an individual's propensity to comply with the mandatory pension contribution requirement. Specifications (1) and (2) differ from specifications (3) and (4), in that they directly control for whether the individual has purchased a new home, rather than whether the individual has purchased that home with a mortgage. Furthermore, specifications (1) and (3) control for whether an individual has access to credit to supplement six months' salary—a condition that throughout this paper has been referred to as full access to credit—while specifications (2) and (4) allow for partial access to credit, which namely asks individuals to answer what percentage of six months' salary they have access to via formal and informal credit channels.

The results in Table 4.6 are presented as odds ratios, which allows us to interpret

the estimated coefficients as a measure of likelihood of defaulting on one's mandatory pension contribution requirement. For instance, in each specification, the odds ratio on the Age variable is roughly 0.85, which means that for every additional year of age, the odds of an individual failing to contribute to his pension program decrease by a factor of 0.15. Said another way, older individuals are more less likely to fail to comply—that is, more likely to comply with the contribution requirement—all else being equal.

Unsurprisingly, it also appears that labor income is correlated with pension compliance. Individuals who earn more money are less likely to fail to comply with the mandatory contribution requirement. This result is identical to that which was reported in Chapter 3, and serves as a crucial sanity-check on the specification. However, the most important findings are those that follow.

Taking specification (1) as our baseline, it is possible to paint a comprehensive picture on the role that purchasing an unsubsidized home plays on one's propensity to comply with the mandatory pension contribution requirement. Broadly speaking, purchasing a home is not correlated with pension compliance, as demonstrated by the odds ratio on the "New House" variable. This indicator variable is equal to 1 when the individual in our data has purchased a new home in the period between 2006 and 2009.

In specifications (3) and (4), however, it is curious that the odds ratio presented for New Mortgage indicates that an individual is significantly more likely to comply with the contribution requirement when compared to an individual who has not received a new mortgage, when the motivating question of this paper is whether individuals suffer a financial constraint when purchasing a home that prevents them from contribution to their pensions. However, it is vital to note that since we are restricting our sample to individuals who were able to purchase a home without utilizing any government subsidies and we are controlling for individuals' access to

credit and method of home financing, the "New Mortgage" variable most likely is picking up the correlation between acquiring a new mortgage and the fact that individuals who are able to finance a home without public funds are 'highly financial-literate types' who are more likely to contribute to their pension program than those who are generally unaffiliated with credit markets, banks, and savings plans.

When examining the impact that access to credit has on one's propensity to comply with the mandatory pension contribution requirement, the finding is as expected: Individuals who have access to credit are more likely to comply with the contribution requirements than individuals without access to credit. This finding holds whether we control solely for individuals who have access to six full months of credit, as demonstrated in specifications (1) and (2), as well as in cases where a more nuanced measure of credit is used that allows for responses along the continuum of credit from zero to all six months' income worth of emergency credit, as in specifications (3) and (4).

This result should come as no surprise. Access to credit allows individuals to smooth their consumption and savings paths in the face of short-term stochastic shocks. Individuals who are able to borrow up to six months' worth of income in an emergency are less pressed to save money by failing to contribute to their pension programs when suffering a negative income shock. This finding reproduces the univariate survival analysis presented in above in Table 4.4, which demonstrates that access to credit significantly reduces one's propensity to fail to comply with the mandatory pension contribution plan, however by extending the findings to a multivariate framework, the result is strengthened as the same conclusion holds in am multivariate framework.

In turning to the next handful of variables, we highlight the primary contribution of this paper. In combination, the odds ratios estimated on the interacted variables "Full Credit X Mortgage" and "Income X Full Credit X Mortgage" present an im-

portant picture on the interaction between homeownership and retirement savings. Controlling for when an individual buys a new home and for one's access to credit, we still find a large and significant increase in the propensity of an individual to fail to contribute to his mandatory pension program upon buying a house through mortgage financing, even when enjoying full access to credit. In specification (1), we are confronted by a surprisingly large estimate. It appears that an individual is five hundred times more likely to fail to contribute to his pension program when he is planning on financing a new house within three years of being surveyed if he has full access to credit.

Now, it is important to underscore the interpretation of the odds ratios presented in Table 4.6, since the result presented above is a strikingly large number. Cox's proportional hazard model allows the econometrician to measure the variation in individuals' propensity to fail to contribute to their pension program by assuming that all individuals have the same intrinsic hazard function that evolves over time. By taking the ratio or *proportion* of two individuals' probability of noncompliance, the hazard function falls out of the equation and the main drivers of two individuals' relative probability of noncompliance are the individual- and job-specific time-invariant variables that define the multivariate analysis.

In the case of this model, we find that when comparing individuals with full access to credit who will purchase a new home via a mortgage in the next few years to individuals with full access to credit who will not purchase a new home, the ratio of the probability of noncompliance is about 500. However, returning to the table of results, one notes that this high point estimate is accompanied by a large standard error. This standard error implies a 95-percent confidence interval of 13 to 20,084 — an incredibly wide margin. While the point estimate remains positive and statistically significant at the 99.9 percent level, the small number of observations of individuals who will purchase a house with a mortgage and also have

full access to emergency credit causes our results to take on large point estimates, while maintaining the sign that is hypothesized above. The stringent criteria for inclusion in the model is lessened in specification (2), as individuals no longer need full access to credit to be included in the analysis — a change that provides much more reasonable estimates.

This odds ratio is highly dependent upon the specification under examination, as evident in specification (2)-(4) in Table 4.6. However, in each case, an individual who is going to purchase a home within three years after being surveyed through a mortgage is more likely to default on his mandatory contribution requirement than an individual who is not about to move. This result almost certainly comes from the requirement that individuals with access to credit can finance no more than 50 to 75 percent of the value of a new home, requiring those who are saving to purchase a new home to accumulate a large down payment when they do not have access to housing subsidies. In order to accumulate this sizeable down payment, many individual find it necessary to forego their mandatory pension contribution requirement.[8]

Finally, it is worth noting that while individuals who are planning on purchasing their home through mortgage financing often times must choose not to contribute to their pension programs in order to amass a sufficient down payment, this will not hold across all new homebuyers. It bears repeating that the analysis presented in this paper focuses on individuals who are purchasing their home without the aid of the numerous kinds of government subsidies available to encourage homeownership. Furthermore, individuals who earn more are not compelled to choose between saving for retirement and saving for the down payment of a house. To this point, the "Income X Full Credit X Mortgage" variables allows us to observe the impact of

[8] It is important to note that since we only observe individuals' pension contribution over the period 2004-2006 housing purchases over the period 2006-2009, we are unable to examine whether individuals return to complying with the mandatory savings requirement after purchasing their home. This is a fruitful area of future study following the publication of the next round of the EPS.

labor income on this forced tradeoff between pension compliance and savings. Note that in specifications (1) and (3), this last variable is significantly less than one, indicating that the earlier observation that individuals are 500–1,000 times more likely to fail to comply with the pension contribution regime when saving for a house is a baseline for the poorest individuals in the sample. This last variable shows that as labor income increases, the impact of the tradeoff between compliance and savings is lessened for wealthier individuals.

4.5 Analysis and Conclusions

In the last section, we showed that individuals who are planning on purchasing a home are less likely to contribute to their pension. Furthermore, this rational non-compliance appears to be more prevalent among lower-income households, because wealthier individuals do not have to choose between complying with the pension savings requirement and saving up a down payment on a house. In Section **??**, our analysis is identified by estimating Cox's proportional hazard model.

However, this model assumes that the time-dependent portion of the hazard rate is constant for all observations. Given this assumption, all variation in the hazard rate is caused by the variables defined explored in Table 4.6. However, we can return to the non-parametric hazard rates estimated above in Subsection 4.3.3 to see what is driving the results above. Specifically, we build on Figure 4.3 by subdividing individuals with full access to credit by whether they are planning on purchasing a home with a mortgage. These hazard rates are presented in Figure 4.5.

Figure 4.5 shows once again that individuals who do not have access to credit are more likely to default on their pension contribution requirements than people with access. Now, however, we subdivide those with access to credit into those who are planning on purchasing a house with a mortgage and those who are not. Here we observe that those who are planning to buy a house have a greater hazard rate.

Therefore, Figure 4.5 enables us to nonparametrically present the key findings from Section 4.4 above.

The nonparametric framework enables us t to explain the unusually large point estimates presented in the pervious section. Since Cox's proportional hazard model holds the shape of the hazard function constant when estimating the marginal impact of each regressor, significant changes in the shape of the hazard function between two subgroups are identified in the hazard ratios of the regressors that distinguish between these groups. Figure 4.5 identifies a major shift in the hazard function between the potential borrowers who plan on financing a house and those who do not. In the Cox framework, this leads a single regressor to represent the entirety of this observed shift, which can lead to outsized estimates. The nonparametric estimation technique, on the other hand, allows for a more subtle parsing of the variation in hazard rates that better reflects the true scale of differences between groups.

Ultimately, the findings presented in this paper focus on the tradeoff between saving for retirement and investing in durable goods. Our most pertinent result is that individuals are constrained by the mandatory pension contribution requirement to favor retirement savings over investment in real estate. This result is softened when households receive housing subsidies or are wealthy enough to choose both options simultaneously. Finally, for individuals who choose to purchase a home despite government policy favoring pension savings, they are often able to default on their contribution requirement and save for housing.

Finally, we must ask who are the people who are faced with this tradeoff? This is to say, who satisfies the following three conditions: 1) Contributing to a pension program, 2) Having access to credit, and 3) Not receiving a housing subsidy. By the requirements of the Chilean pension system, all individuals who are privately employed must contribute regularly to their retirement account every pay period,

so long as they are not self-employed. This requirement binds for all individuals irrespective of their level of income. Therefore all individuals who are not self-employed satisfy the first condition.

As for the second condition, we have shown that access to credit is unsurprisingly correlated with household income. Therefore, our analysis focuses on middle- and high-income households who have access to the private credit market. That said, the definition of the variable used to determine access to credit is based upon an individual's ability to borrow six-month's salary in the event of an emergency. This definition obviously scales with income, however, since wealthier individuals would need to provide an increasing amount of collateral to borrow an increasing nominal amount of money, there is reason to believe that it is altogether easier for high-income individuals to acquire six-months' salary on credit than middle-income individuals.

Finally, we have already stated that there are multiple kinds of housing subsidies provided by the Chilean government. While many of the subsidies are aimed at encouraging developers to build low-income housing, our interest lies in traditional direct subsidies to potential homeowners. For low-income households, these subsidies are akin to receiving a low-interest loan from the government to purchase a home with little money down. However for the middle- and higher-income individuals who satisfy the first two conditions listed here, government subsidies take the form of a government provided mortgage loaned on par with private-sector interest rates. Furthermore, households must demonstrate that they have accumulated a sizeable down payment for a house prior to receiving the subsidy, and even then must contend with the limited availability of government funds. Hence many middle-income individuals who qualify for a government subsidy fail to receive one.

Those who satisfy all three conditions are most susceptible to the government-mandated tradeoff highlighted in this paper. Otherwise law-abiding individuals who are looking to purchase a new — or in many cases, first — home are compelled to

choose between satisfying the pension contribution requirement or accumulating a sufficient down payment on a home to receive a private-sector mortgage. Thus we see that how government policies designed to spur long-term retirement savings and investment in housing can work at cross purposes.

Individuals are faced with economic tradeoffs every day, and a key role of public policy is to incentivize the socially beneficial outcomes of these decisions. They main lesson for pension administrators and policymakers proffered in this paper is to recognize that governments are frequently unwieldy organizations that seek to encourage a variety of social aims through various institutions. However, these institutions will often be competing for individuals' scarce resources, which can undermine the aims of policymakers and reduce the efficiency of the institutions they manage.

Defined contribution pension programs place the burden of saving for retirement on the individual while maintaining a social safety net for those who do not save enough. It is only by encouraging regular pension contributions and credibly punishing those who fail to contribute that policymakers can hope to limit the number of free-riders who will turn to unfunded pay-as-you-go pension programs when they have not saved enough. However, in addition to punishing those who fail to contribute regularly to their pensions, policymakers must recognize that many people choose not to contribute to their illiquid pensions in order to take advantage of other politically-encouraged economic opportunities. Thus, policymakers have the opportunity to reorganize the institutions they oversee to mitigate the effects of institutional competition and encourage responsible savings through a variety of paths, including defined-contribution pension programs and durable goods such as housing.

Table 4.1: Homeowners vs. Renters

	Homeowners	Renters	T-Test
Age	39.50	37.32	6.61
HH Size	5.20	5.28	0.99
HH Income per capita	753,493	703,761	1.11
Total HH Savings	151,513	72,001	2.12
Full Credit (%)	35.7	27.7	5.46
Obs	3,440	1,212	
Weighted Proportion	72.0%	28.0%	

Table 4.2: Future Homebuyers' Access to Full Credit by Home Financing

Financing	Percent with Full Credit	Weighted Proportion	Obs
Cash	24.9	19.7	45
Mortgage	40.1	44.1	105
Other Credit	9.4	6.3	14
Inherited	4.2	3.3	9
Other	16.5	26.6	51
Total	29.2	–	224

Table 4.3: Future Homeowners by Subsidy

	No Subsidy	Subsidy	T-Test
Age	35.40	33.77	1.04
HH Size	5.05	5.84	2.05
HH Income per capita	915,504	699,556	0.78
Total HH Savings	384,044	223,784	0.83
Full Credit (%)	48.6	23.1	3.07
Obs	105	119	
Weighted Proportion	45.9%	54.1%	

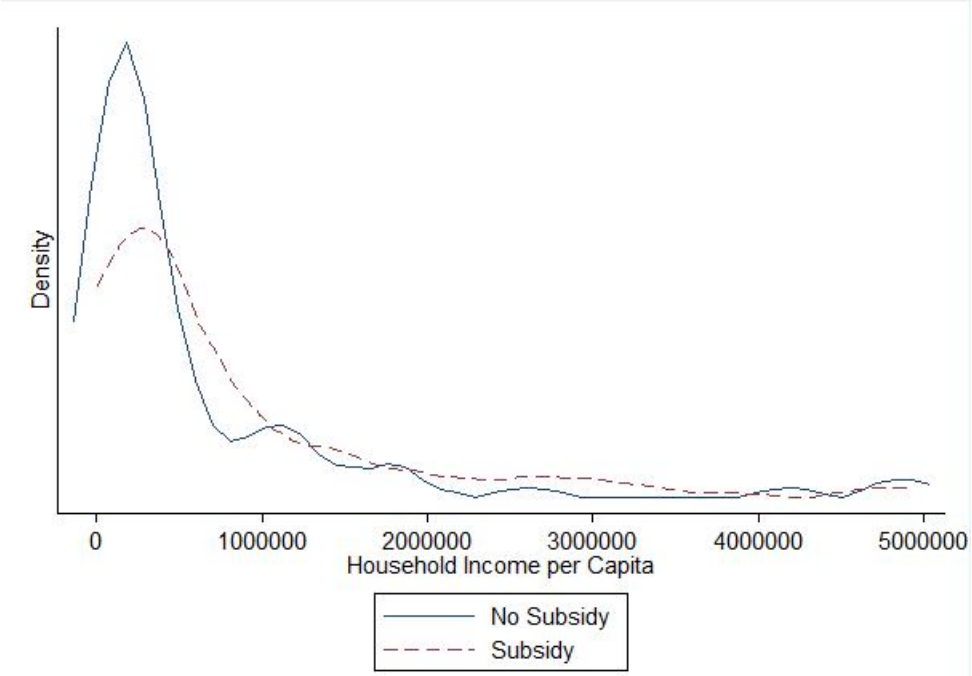

FIGURE 4.1: Empirical Density of New Homeowners' Total Household Income per Capita

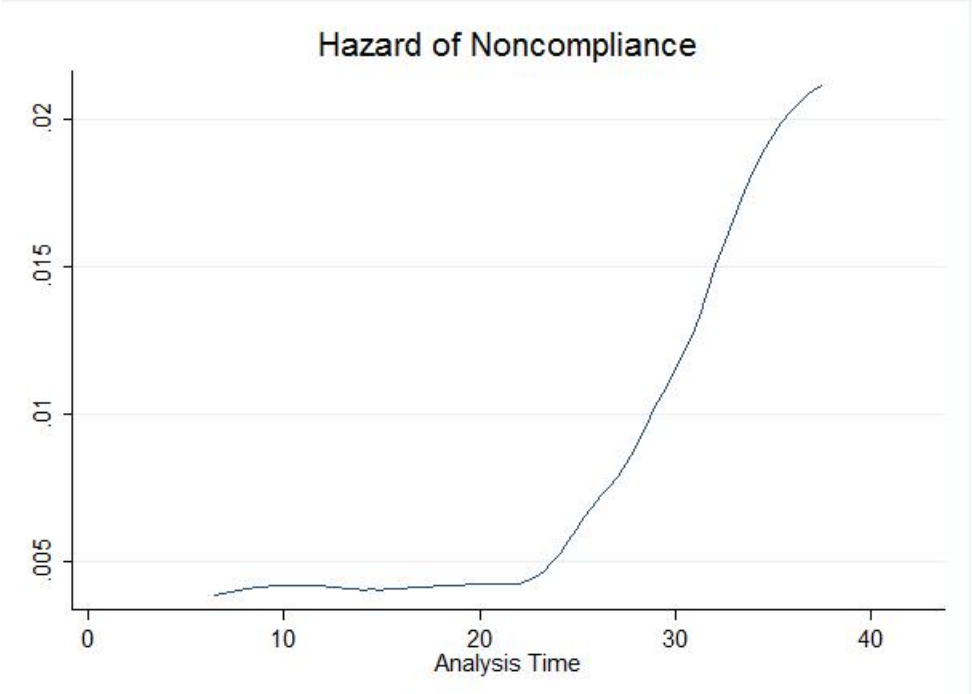

FIGURE 4.2: Hazard Rate of Failure to Contribute to Pension

Table 4.4: Failure to Contribute by Full Access to Credit

Access to Credit	Weighted Share of Population	Obs	Cox Relative Hazard
Yes	24.7%	1,267	0.403%
No	75.2%	3,385	1.463%

Pearson's $\chi^2(1) = 232.55$

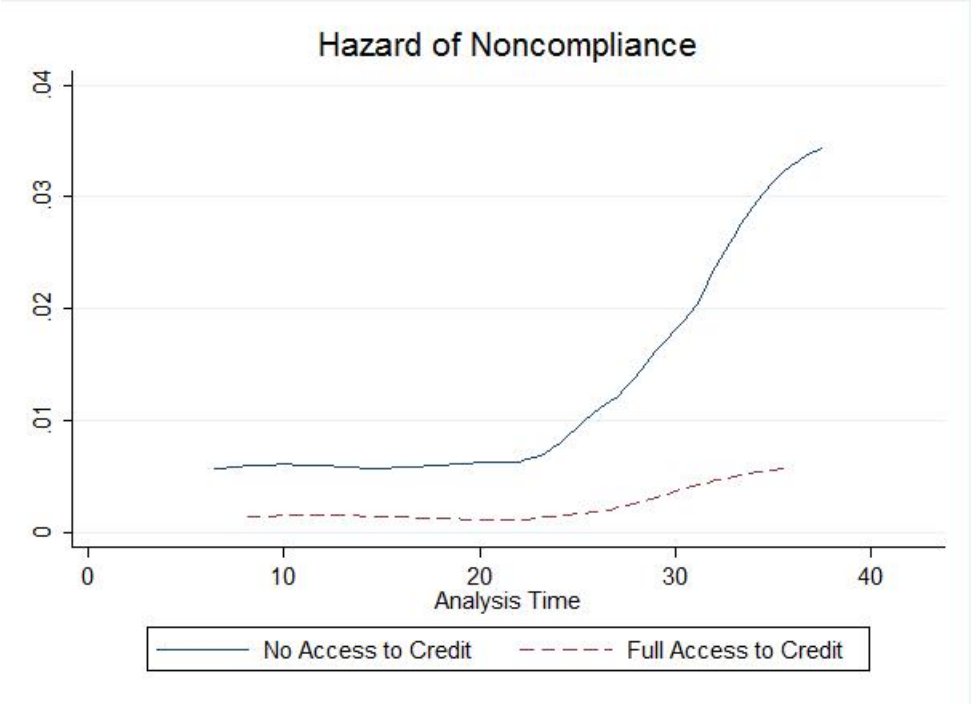

FIGURE 4.3: Hazard Rate by Access to Credit

Table 4.5: Failure to Contribute by Purchase Date of House

Housing Variable	Weighted Share of Population	Obs	Cox Relative Hazard
New House in 3 Years	6.9%	336	0.675%
New Mortgage in 3 Years	1.9%	108	0.291%

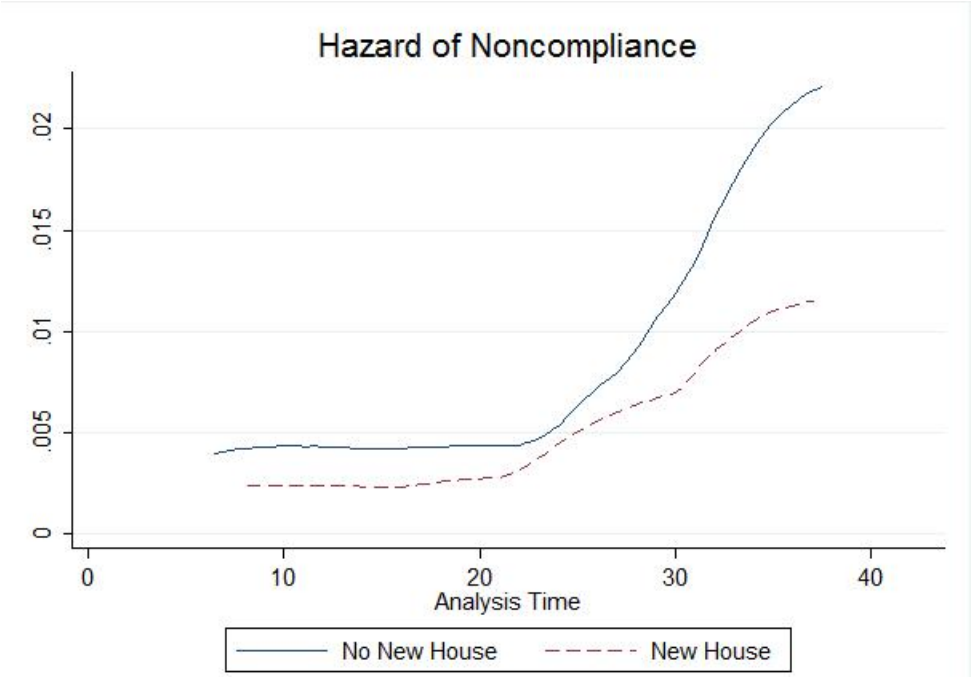

FIGURE 4.4: Hazard Rate by Future Homeownership

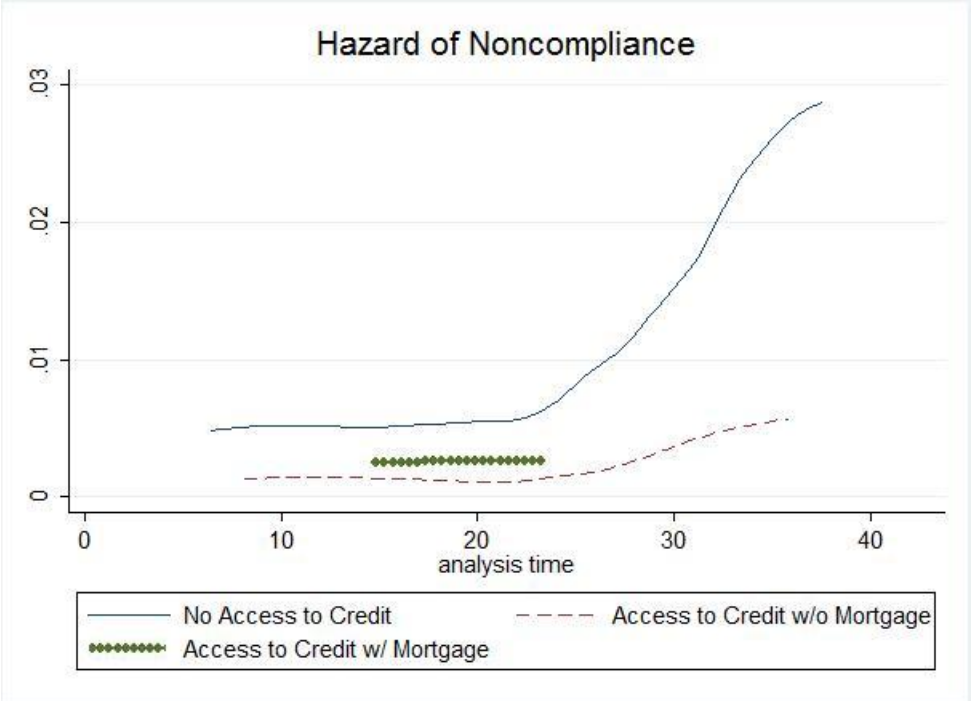

FIGURE 4.5: Hazard Rate by Access to Credit and Home Mortgage Financing

Table 4.6: Proportional Hazard Model on Time to Noncompliance: Unsubsidized Housing

VARIABLES	(1) Full Credit	(2) Partial Credit	(3) Full Credit	(4) Partial Credit
Age	0.847***	0.849***	0.844***	0.846***
	(0.0415)	(0.0417)	(0.0414)	(0.0417)
Age^2	1.002***	1.002***	1.002***	1.002***
	(0.000628)	(0.000630)	(0.000630)	(0.000631)
Years of Education	1.004	1.005	1.005	1.006
	(0.0340)	(0.0341)	(0.0341)	(0.0343)
Labor Income	0.999***	0.999***	0.999***	0.999***
	(7.29e-08)	(7.41e-08)	(7.26e-08)	(7.37e-08)
New House	0.584	0.584		
	(0.218)	(0.218)		
New Mortgage			0.296*	0.168*
			(0.218)	(0.154)
Full Credit Access	0.326***		0.321***	
	(0.0884)		(0.0873)	
Partial Credit Access		0.325***		0.319***
		(0.0839)		(0.0826)
Full Credit X Mortgage	509.2***		1,027***	
	(954.2)		(2,040)	
Partial Credit X Mortgage		58.57*		319.1**
		(125.0)		(856.4)
Income X Full Credit X Mortgage	0.999***		0.999***	
	(4.70e-07)		(4.69e-07)	
Income X Partial Credit X Mortgage		0.999		0.999
		(4.27e-07)		(4.67e-07)
Region-Specific Dummies	Yes	Yes	Yes	Yes
Observations	3,083	3,005	3,083	3,005

s.e. in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Bibliography

Acuña, R. R. and Iglesias, A. P. (2001), *Chile's Pension Reform after 20 Years*, World Bank, Social Protection discussion paper no. 0129, Washington, DC.

Afzal-ur-Rahman (1980), *Economic Doctrines of Islam*, vol. 1, Islamic Publication, Lahore, second edn.

Alvayay, J. R. and Schwartz, Jr., A. L. (1997), "Housing and Mortgage Market Policies in Chile," *Journal of Real Estate Literature*, 5, 47–55.

Arenas de Mesa, A., Behrman, J. R., and Bravo, D. (2004), "Characteristics of and Determinants of the Density of Contributions in a Private Social Security System," University of Michigan Retirement Research Center Working Paper.

Arenas de Mesa, A., Bravo, D., Behrman, J. R., Mitchell, O. S., and Todd, P. E. (2008), *The Chilean Pension Reform Turns 25: Lessons from the Social Protection Survey*, Oxford University Press, Oxford.

Barr, N. A. (2002), "Reforming Pensions: Myths, Truths, and Policy Choices," *International Social Security Review*, 55, 3–36.

Barrientos, A. (1996), "Pension Reform and Pension Coverage in Chile: Lessons for Other Countries," *Bulletin of Latin American Research*, 15, 309–322.

Barrientos, A. (2002), "Health Policy in Chile: The Return of the Public Sector?" *Bulletin of Latin American Research*, 21, 442–459.

Barrientos, A. and Lloyd-Sherlock, P. (2000), "Reforming Health Insurance in Argentina and Chile," *Health Policy and Planning*, 14, 417–423.

Bayındır, A. (1986), *İslâm Muhakeme Hukuku (Osmanlı Devri Uygulaması)*, İslâmî İlimler Araştırma Vakfı, Istanbul.

Becker, C. M., Marchenko, G. A., Khakimzhanov, S., Seitenova, A. S., and Ivliev, V. (2009), *Social Security Reform in Transition Economies: Lessons from Kazakhstan*, Palgrave Macmillan, New York, NY.

Bhattacharya, U., Galpin, N., and Haslem, B. (2007), "The Home Court Advantage in International Corporate Litigation," *Journal of Law and Economics*, 50, 625–659.

Bodenhausen, G. V. (1955), "Stereotypic Biases in Social Decision Making and Memory: Testing Process Models of Stereotype Use," *Journal of Personality and Social Psychology*, 88, 726–737.

Bruce, N. (2000), "The Chilean Health Care Reforms: Model or Myth?" *Journal of Public and International Affairs*, 11, 69–86.

Chapra, M. U. (1992), *Islam and the Economic Challenge*, Islamic Foundation, Leicester.

Clermont, K. M. and Eisenburg, T. (1996), "Xenophilia in American Courts," *Harvard Law Review*, 109, 1120–1143.

Coşgel, M., Miceli, T., and Ahmed, R. (2009), "Law, State Power, and Taxation in Islamic History," *Journal of Economic Behavior and Organization*, 71, 704–717.

Cox, D. R. (1972), "Regression Models and Life-Tables," *Journal of the Royal Statistical Society*, 34, 187–220.

Cummings, J. and DiPasquale, D. (2002), "The Spatial Implications of Housing Policy in Chile," in *Chile: Political Economy of Urban Development*, eds. E. L. Glaeser and J. R. Meyer, Harvard University Press, Cambridge, MA.

Dam, K. W. (ed.) (2006), *The Law-Growth Nexus: The Rule of Law and Economic Development*, Brookings Institution Press, Washington, DC.

DiIanni, I. (2010), "The Role of Competition in the Market for Adjudication," forthcoming in Supreme Court Economic Review.

Ekinci, E. B. (ed.) (2004), *Tanzimat ve Sonrası Osmanlı Mahkemeleri*, Arı Sanat, Istanbul.

El-Gamal, M. A. (ed.) (2006), *Islamic Finance: Law, Economics, and Practice*, Cambridge University Press, New York.

Faroqhi, S. (1997), "Sidjill, Ottoman administrative usage," in *Encyclopedia of Islam*, vol. 9, pp. 539–545, Brill, Leiden, second edn.

Feld, L. P. and Voigt, S. (2003), "Economic Growth and Judicial Independence: Cross-Country Evidence Using a New Set of Indicators," *European Journal of Political Economy*, 19, 497–527.

Ferguson, B. (2004), "Housing Finance Options for Low and Medium Income Families: Analysis of the Latin America Experience," *Housing Finance International*, pp. 11–14.

Gaudefroy-Demombynes, M. (ed.) (1950), *Muslim Institutions*, George Allen and Unwin, London, Translated by John P. MacGregor.

Gazal-Ayal, O. and Sulitzeanu-Kenan, R. (2010), "Let My People Go: Ethnic In-Group Bias in Judicial Decision—Evidence from a Randomized Natural Experiment," *Journal of Empirical Legal Studies*, 7, 403–428.

Gill, I. S., Packard, T., and Yermo, J. (2005), *Keeping the Promise of Social Security in Latin America*, Stanford University Press, Palo Alto, CA.

Heckman, J. J., Nelson, R. L., and Cabatingan, L. (eds.) (2004), *Global Perspectives on the Rule of Law*, Routledge, New York.

Helmke, G. (2002), "The Logic of Strategic Defection: Court-Executive Relations in Argentina under Dictatorship and Democracy," *American Political Science Review*, 96, 291–303.

Helmke, G. (ed.) (2005), *Courts under Constraints: Judges, Generals and Presidents in Argentina*, Cambridge University Press, New York.

Holst, J., Laaser, U., and Hohmann, J. (2004), "Chilean Health Insurance System: A Source of Inequity and Selective Social Insecurity," *Journal of Public Health*, 12, 271–282.

Imber, C. (ed.) (2002), *The Ottoman Empire, 1300–1650: The Structure of Power*, Palgrave Macmilan, New York.

Kessler, D., Meites, T., and Miller, G. (1996), "Explaining the Deviations from the Fifty-Percent Rule: A Multimodal Approach to the Selection of Cases for Litigation," *Journal of Legal Studies*, 25, 233–259.

Klerman, D. (2007), "Jurisdictional Competition and the Evolution of the Common Law," *University of Chicago Law Review*, 74, 1179–1226.

Kotlikoff, L. (1999), "The World Bank's Approach and the Right Approach to Pension Reform," Boston: Boston University Department of Economics, unpublished manuscript.

Kuran, T. (2004), "The Economic Ascent of the Middle East's Religious Minorities: The Role of Islamic Legal Pluralism," *Journal of Legal Studies*, 33, 475–515.

Kuran, T. (ed.) (2010–2011b), *Mahkeme Kayıtları Işığında 17. Yüzyıl İstanbul'unda Sosyo-Ekonomik Yaşam / Social and Economic Life in Seventeenth-Century Istanbul: Glimpses from Court Records*, İş Bankası Kütür Yayınları, Istanbul, 10 vols.

Kuran, T. (2011a), *The Long Divergence: How Islamic Law Held Back the Middle East*, Princeton University Press, Princeton.

Landes, W. M. and Posner, R. A. (1979), "Adjudication as a Private Good," *Journal of Legal Studies*, 8, 235–284.

Lira, R. (1994), "Pension Funds and Housing Finance in Chile: A Question of Social Efficacy," *Journal of Housing Research*, 5, 229–245.

Mandaville, J. E. (1966), "The Ottoman Court records of Syria and Jordan," *Journal of the American Oriental Society*, 86, 311–319.

Mantran, R. (1962), *Istanbul dans la Seconde Moitié du XVII Siècle*, Librarie Adrien Maisonneuve, Paris.

Masters, B. (2001), *Christians and Jews in the Ottoman Arab World: The Roots of Sectarianism*, Cambridge University Press, New York.

Mitchell, O. S., Todd, P., and Bravo, D. (2007), "Learning from the Chilean Experience: The Determinants of Pension Switching," Pension Research Council Working Paper.

Moore, K. A. (2003), "Xenophobia in American Courts," *Northwestern University Law Review*, 97, 1497–1550.

Naqvi, S. N. H. (2003), *Perspectives on Morality and Human Well-Being: A Contribution to Islamic Economics*, Islamic Foundation, Leicester.

Orszag, P. R. and Stiglitz, J. E. (2001), "Rethinking Pension Reform: Ten Myths About Social Security Systems," in *New Ideas About Old Age Security: Toward Sustainable Pension Systems in the 21st Century*, eds. R. Holzmann and J. E. Stiglitz, pp. 17–56, The World Bank, Washington, DC.

Ortaylı, I. (1994), *Hukuk ve İdare Adamı Olarak Osmanlı Devletinde Kadı*, Turhan Kitabevi, Ankara.

Özmucur, S. and Şevket Pamuk (2002), "Real Wages and Standards of Living in the Ottoman Empire, 1489–1914," *Journal of Economic History*, 62, 293–321.

Pardo, C. A. (1999), "Housing Finance in Chile: Primary and Secondary Mortgage Financing," *Housing Finance International*, pp. 15–25.

Pardo, C. A. (2001), "The Chilean 'Tripartite' Approach: Loans, Family Savings, and State Subsidies," *Housing Finance International*, pp. 32–41.

Piñera, J. (2004), "Empowering Workers: The Privatization of Social Security in Chile," in *Social Security and its Discontents*, ed. M. B. Tanner, Cato Institute, Washington, DC.

Porter, J. (1771), *Observations on the Religion, Law, Government, and Manners of the Turks*, J. Nourse, London, second edn.

Priest, G. L. and Klein, B. (1984), "The Selection of Disputes for Litigation," *Journal of Legal Studies*, 13, 1–55.

Ramseyer, J. M. and Rasmusen, E. B. (2003), *Measuring Judicial Independence: The Political Economy of Judging in Japan*, University of Chicago Press, Chicago.

Ruprah, I. J. and Marcano, L. T. (2005), *A Meta-Impact Evaluation of Social Housing Programs: The Chilean Case*, Inter-American Development Bank's Office of Evaluation and Oversight, Washington, DC.

Saeed, A. and Saeed, H. (2004), *Freedom of Religion, Apostasy and Islam*, Ashgate, Aldershot, UK.

Sapelli, C. (1999), "Mandatory Health Insurance in Chile: Issues of Design in Theory and in Practice," Paper prepared for the IHEA 2nd World Conference.

Sapelli, C. and Torche, A. (2001), "The Mandatory Health Insurance System in Chile: Explaining the Choice between Public and Private Insurance," *International Journal of Health Care Finance and Economics*, 1, 97–110.

Shapiro, M. (1981), *Courts: A Comparative and Political Analysis*, University of Chicago Press, Chicago.

Shavell, S. (1996), "Any Frequency of Plaintiff Victory at Trial Is Possible," *Journal of Legal Studies*, 25, 493–501.

Shayo, M. and Zussman, A. (2010), "Judicial Ingroup Bias in the Shadow of Terrorism," forthcoming in Quarterly Journal of Economics.

Siegelman, P. and Donohue, III, J. J. (1995), "The Selection of Employment Discrimination Disputes for Litigation: Using Business Cycle Effects to Test the Priest-Klein Hypothesis," *Journal of Legal Studies*, 24, 427–462.

Siegelman, P. and Waldfogel, J. (1999), "Toward a Taxonomy of Disputes: New Evidence through the Prism of the Priest/Klein Model," *Journal of Legal Studies*, 29, 101–130.

Solimano, A. and Torche, A. (2008), "La Distribución del Ingreso en Chile 1987–2006: Análisis y Consideraciones de Política," Central Bank of Chile: Working Paper.

Srinivas, P. S. and Yermo, J. (1999), *Do Investment Regulations Compromise Pension Fund Performance?*, The International Bank for Reconstruction and Development, Washington, DC.

The World Bank (1994), *Averting the Old Age Crisis: Policies to Protect the Old and Promote Growth*, Oxford University Press, Oxford.

Uzunçarşılı, I. H. (1965), *Osmanlı Devletinin İlmiye Teşkilâtı*, Türk Tarih Kurumu, Ankara.

van Tulder, F. P. and van Velthoven, B. C. J. (2003), "Econom(etr)ics of Crime and Litigation," *Statistica Neerlandica*, 57, 321–346.

Biography

Scott Jordan Lustig was born in New York City on Wednesday, February 22nd 1984 at 6 lbs. 4 ozs. After graduating from Farmingdale High School in 2002, he studied International Relations and Quantitative Economics at Tufts University, with guest appearances at the London School of Economics and the Center for Russian Language and Culture in Saint Petersburg, Russia.

Following graduation from Tufts University in 2006, he enrolled in a Ph.D. program in the Duke University Department of Economics. There he received the Distinguished Economics Department Graduate Fellowship and the Robert K. Steel Family Summer Fellowship while working with Professors Frank Sloan, Shakeeb Khan, Timur Kuran, and Charles Becker.

After graduating from Duke University in May 2011, Scott will begin working at Citicards in New York City. Once while away at summer camp, Scott earned the Golden Hammer for excellence in wood shop! Has also has an amateur radio license, played guitar in Farmingdale's own "Seven Day Weekend," and once went 60 days without shaving to win the Great Beard-off of 2007.