Essays in Insurance Economics

by

Cameron McNeill Ellis

(Under the Direction of William B. Vogt)

Abstract

The first chapter of this dissertation considers the welfare effects of increasing access to secondary markets in life insurance. In this chapter, I propose and estimate a life-cycle savings model for the life insurance lapse decision with dynamic, heterogeneous bequest motives. I then perform a counterfactual analysis with competitive secondary markets and find them to be Pareto improving for my sample with an average value increase to consumer's welfare by $1,346 per policy-holder. The second chapter of this dissertation considers how optional two-part tariffs can serve as a signaling device for life insurance contracts. I test for consumer self-selection using detailed, policy-level data within the context of life insurance backdating. I am able to identify, through a control function approach, the information about lapse risk a consumer reveals when they choose to backdate. I find consumers a) who are less likely to lapse self-select into the two-part tariff pricing structure and b) exhibit behavior consistent with sunk cost bias. The final chapter of this dissertation considers how Medicaid expansion can affect private insurance markets. I use policy-level data from the Health Insurance Exchanges to identify and estimate the effects of Medicaid expansion on the private health insurance market premiums. I find that expanding Medicaid reduces average monthly premiums by $32.4, a decrease of 11.86%.

Index words: Life Insurance, Health Insurance, Affordable Care Act, Secondary Markets, Bequest Motives, Life Insurance Lapsing, Asymmetric Information

Essays in Insurance Economics

by

Cameron McNeill Ellis

B.S., Georgia College & State University, 2010

M.B.A., Georgia College & State University, 2011

A Dissertation Submitted to the Graduate Faculty of The University of Georgia in Partial Fulfillment of the Requirements for the Degree

Doctor of Philosophy

Athens, Georgia

2017

Essays in Insurance Economics

by

Cameron McNeill Ellis

Major Professor: William B. Vogt

Committee: Scott E. Atkinson
James M. Carson

Electronic Version Approved:

Suzanne Barbour
Dean of the Graduate School
The University of Georgia
May 2017

Essays in

Insurance Economics

Cameron M$^{\text{c}}$Neill Ellis

April 24, 2017

Dedication

This dissertation is dedicated to my parents, Carole Knight and Charlie Ellis.

Acknowledgments

I would first like to acknowledge the help and guidance of my dissertation chair, William B. Vogt. I would also like to thank my committee members, Jim Carson and Scott Atkinson. In addition, I would like to thank the following individuals who made helpful comments that served to improve my dissertation essays: Meghan Esson, Evan Eastman, Joshua Frederick, Meghan Skira, Sabine Wende, Jon Williams, Lars Powell, Ian Schmutte, Art Snow, Alexander Muermann, Irina Gemmo, David Russell, Rob Hoyt, Kryzysztof Ostaszewski, Josh Kinsler, and Svetlana Paschenko.

Contents

List of Figures

List of Tables

Introduction

The first chapter of this dissertation considers the welfare effects of increasing access to secondary markets in life insurance. Life insurance contracts can be exceptionally long term and are typically written with a level premium structure. Because death risk increases with age, the actuarial value of a life insurance policy increases over time and becomes positive far enough into the policy. Life insurance is also unique in that the payout is valued through a bequest motive. If bequest preferences are dynamic and subject to unexpected shocks, the consumer's value of a life insurance contract can become negative even when the actuarial value is positive. Thus, there are potential gains from trade from a secondary market for life insurance policies. While these markets do exist, they are limited and controversial. In this paper, I propose and estimate a life-cycle savings model for the life insurance lapse decision with dynamic, heterogeneous bequest motives. I then perform a counterfactual analysis with competitive secondary markets and find them to be Pareto improving for my sample with an average value increase to consumer's welfare by $1,346 per policy-holder.

The second chapter of this dissertation considers how optional two-part tariffs can serve as a signaling device for life insurance contracts. I develop a model of insurance pricing under heterogeneous lapsing with asymmetric information about lapse likelihood within the context of an optional two-part tariff as a signaling device for future policyholder behavior. Specifically, I test for consumer self-selection using detailed, policy-level data within the context of life insurance backdating (a common practice that resembles a two-part tariff). I am able to identify, through a control function approach, the information about lapse risk a consumer reveals when they choose to backdate. I provide empirical evidence that consumers a) who are less likely to lapse self-select into the two-part tariff pricing structure and b) exhibit behavior consistent with sunk cost bias.

The final chapter of this dissertation considers how Medicaid expansion can affect private insurance markets. The Affordable Care Act is one of the most debated and dividing pieces of legislation in recent memory. One of the main elements of the ACA is the Optional expansion of Medicaid eligibility from the poverty line to 138% of the poverty line and inclusion of childless adults. Nearly all of the debate has focused on the direct effects of the newly covered, but there are also important other effects to consider. If the newly-eligible portion differs from the general populace then expansion of Medicaid can affect the private market for health insurance. I use policy-level data from the Health Insurance Exchanges to identify and estimate the effects of Medicaid expansion on the private health insurance market premiums. I find that expanding Medicaid reduces average monthly premiums by \$32.4, a decrease of 11.86%.

This dissertation covers, broadly, consumer responses to various quirks in the insurance industry, life and health specifically, and the implications this consumer behavior has on optimal policy. The insurance industry in particular is subject to problems from asymmetric information. The final chapter of this dissertation examines, and finds evidence of, adverse selection in private, individual health insurance and how an expansion of public insurance can help negate these problems. The second chapter covers a different form of selection and asymmetric information: the risk that a consumer will lapse on their life insurance policy. The policy recommendation from this chapter is for life insurance companies to offer a menu of two-part tariff pricing structures and allowing the consumers to self select into the policy designed for their own lapse type. The first chapter examines a different quirk in the life insurance industry – that the benefits of the contract are received through a bequest motive. Life insurance contracts are typically long in nature and bequest motives have a dynamic aspect to them. This can lead to cases where the actuarial value of the life insurance is positive but the consumer does not want it anymore because their bequest preference has changed. The policy recommendation in this chapter is the removal of extant limits on secondary markets for life insurance.

Chapter 1
Dynamic Bequest Motives and Secondary Markets for Life Insurance

1.1 Introduction:

Relative to the majority of insurance policies, whole life insurance contracts are exceptionally long term, often lasting in excess of 30 years. These policies are also typically written where premiums are the same each period and not adjusted to changing risks, called level premiums. Because death risk increases with age, the actuarial value of a life insurance policy increases over time and often becomes positive later in the policy. Life insurance is also different from other forms of insurance in that the payout is not received by the insured, but instead by a beneficiary. This means the consumer's value of the insurance is through a bequest motive. If bequest preferences are dynamic and subject to unexpected shocks, such as through divorce or the aging of a child, the consumer's value of a life insurance contract can become negative even when the actuarial value is positive. Thus potential gains from a secondary market exist. For example, a market where a risk neutral firm pays the policy holder a lump sum up front (and takes over the subsequent premium payments) for the policy holder to name the firm as the beneficiary of the policy. These markets do exist, but they are limited, controversial, and actively lobbied against by life insurance firms (Gottlieb and Smetters, 2014). In this paper, I propose and estimate a life-cycle savings model for the life insurance lapse decision with dynamic, heterogeneous bequest motives. I then perform a counterfactual analysis with competitive secondary markets and find them to be Pareto improving for my sample with an average compensation-value increase to consumers of \$1,346.

There is a secondary market for life insurance in the United States. This market originally arose as "viatical settlements" in the 1980s in response to the AIDS epidemic. Young insureds who contracted HIV/AIDS had drastically reduced life expectancies as well as financial hardships from high medical costs. Individual investors, and later firms, began purchasing the life insurance contracts from these terminally ill patients for lump sums that were greater than the cash values of the policies (Doherty and Singer, 2003b). Increases in medical technology that extended the life expectancy of HIV/AIDS patients brought a temporary end to the viatical settlement market, but the development of life settlement firms in the early 2000s has allowed more individuals to access the market.[1] Though life settlement markets are growing, they are still underdeveloped. Of the $492 billion in face value of life insurance held by consumers over age 65, nearly 20% have an economic value that exceeds the surrender value (Doherty and Singer, 2003a). The total face value of policies held in the secondary market is less than $15 billion (Gatzert et al., 2009). These numbers still understate the potential value of complete secondary markets as they do not consider the increase in the "option value" of the life insurance contract. Even for policies that are not currently saleable on secondary markets, the value of the policy would increase. Due to level premiums, there always exists, in a competitive market, a potential point in the life of the contract where it becomes investment-worthy.[2]

Of course, even in the absence of secondary markets, consumers are not forced to continue paying into a life insurance contract they no longer want. At any point in a life insurance contract a consumer can simply stop paying premiums, referred to as lapsing, and the rates for this are enormous. 88% of whole life insurance policies never pay a death benefit. More surprisingly, 76% of whole life policies sold to people at age 65 fail to pay a death benefit

[1]Though often used interchangeably, viatical settlements involve a terminally ill insured while life settlements do not.

[2]For non-competitive markets this will depend on interest rates.

(Gottlieb and Smetters, 2014).[3] Between 1991 and 2010, $29.7 trillion of new individual life insurance coverage was issued in the United States. During this same time period, $24 trillion of coverage lapsed (Gottlieb and Smetters, 2014).

In order to explore the potential of secondary markets, it is important to first understand the various reasons consumers lapse on policies, the rates of which are enormous. This literature is well developed and has naturally segmented the lapse decision into: preference shocks, income shocks, policy replacement, as well as behavioral explanations.[4]

Preference shocks refer to any number of situations where the consumer's preference for life insurance has changed, typically through shocks to bequest motives (Fang and Kung, 2010; Liebenberg et al., 2012; Fei et al., 2015).[5] Examples include: divorce, death of a spouse or child, children becoming self-sufficient, increase in spousal income, etc. Income shocks refer broadly to consumers experiencing a negative shock to income and thereby having insufficient funds to pay premiums. The effect of an income shock on lapse rates is stronger in whole life insurance due to the dual presence of a surrender value and larger premiums; this subset is referred to in the literature as the "emergency fund hypothesis" (Linton, 1932; Outreville, 1990; Kuo et al., 2003). The aptly named "policy replacement hypothesis" refers to consumers who lapse on one policy because they found a better one (Outreville, 1990; Carson and Forster, 2000). Because insurance premiums are collected, and invested, long before benefits are paid out, expectations about interest rates play an important part in the determination of premium rates. Thus, the "interest rate hypothesis" is a specific case of the policy replacement hypothesis where the driver of the newly available superior policy is a change in expectations of future interest rates (Schott, 1971; Pesando, 1974; Kuo et al.,

[3]Whole life insurance policies are also called cash value policies since they accumulate a "cash value" that can be redeemed at the surrender of the policy. However, even after long periods, this is generally much less than the face value or the actuarial value of the policy.

[4]The vast majority of literature on insurance lapsation is focused on whole life insurance specifically, but much of the theory applies to a broader spectrum of insurance contracts.

[5]This reason for lapsing generally only covers negative shocks to the insurance value of the policy. Positive shocks to preferences are generally subsumed by the policy replacement category.

2003). The final category of research on lapse rates focuses on non-expected utility models of consumer behavior and how these various behavioral assumptions can influence the decision to lapse on a policy (Shefrin, 2002; Mulholland and Finke, 2014; Gottlieb and Smetters, 2014).[6]

Because of the "dynamic shock" nature of the general reasons for lapsing on a life insurance policy, a dynamic, life-cycle model is the natural structure to examine the mechanics of the lapse decision. The theoretical literature on life insurance demand features extensive use of dynamic programming (e.g. Yaari, 1965; Fischer, 1973; Lewis, 1989). However, to the author's knowledge, the empirical literature has not used these methods. This article bridges that gap. The investigation into the welfare implications of increasing consumer access to secondary markets for life insurance is the main contribution of this paper as well as the formal treatment of life insurance demand as a function of shocks to both bequest motives and income. The rest of the article proceeds as follows: Section 1.2 discusses relevant institutional mechanics of the life insurance industry, Section 1.3 establishes the model and makes explicit the implicit assumptions required, Section 1.4 describes in detail the estimation strategy, Section 1.5 describes my data, Section 1.6 presents the parameter-fitting and welfare analysis results, and Section 1.7 offers conclusions.

1.2 Institutional Mechanics:

There are, in general, two types of life insurance: term and whole. Term life insurance provides protection for a given amount of time, where the term is generally somewhere between 1 and 30 years. These policies are useful for situations where protection from an untimely death is needed only for a defined period of time, such as until a child graduates college or until retirement. These policies are generally cheaper than whole life policies.[7]

[6]For an excellent, detailed analysis of these hypotheses the interested reader is referred to Eling and Kochanski (2013).

[7]Cheaper in the sense that an individual can purchase a policy with a greater face value for the same premium amount.

Whole life insurance is a permanent insurance contract that, once purchased, will insure against death for any number of years.[8] Whole life insurance, unlike term policies, is combined with a savings component called the cash value. This cash value grows throughout the policy and is returned to the consumer in the event of a policy surrender.[9] The popular financial advice is to "buy term and invest the difference;" however the academic literature is divided on the issue. Consumers are also generally allowed to borrow against their accumulated cash value at a prescribed (fixed or variable) interest rate. Throughout this paper I am exclusively examining whole life insurance.

Modeling the life insurance lapse decision is, in a way, the mirror image of modeling the demand for life insurance, the theory and estimation of which has been a rich area of research for over 50 years. Yaari (1965) is the original theoretical work, using bequest motives and an uncertain lifetime as motivation for demand. Fischer (1973) expanded on this notion, incorporating a full life-cycle model into the life insurance purchase decision. Campbell (1980) was the first to move away from pure bequest motives and into a Becker (1965) household production model. Cambell's model focuses on life insurance as protection, for the household, against the untimely death of the primary wage earner. Somewhat independently, a stream of work has arisen in the finance literature on the asset value of a life insurance contract, most of which take a portfolio theory approach (e.g. Karni and Zilcha, 1986; Mayers and Smith Jr, 1983; Doherty, 1984; Lin and Grace, 2007).[10]

Secondary markets, originally called viatical settlements, arose in the 1980s in response to the AIDS epidemic. Young insureds who contracted HIV/AIDS had drastically reduced

[8]This is not strictly true, as life insurers will generally set an age at which all policyholders are assumed to "die" and benefits are paid out. This is generally 100 or (more recently) 120 years, but can differ based on the insurer and the policy.

[9]Variable life insurance offers policyholders the chance to choose from certain portfolios of securities where to invest their cash value, with the potential to earn a higher return.

[10]This literature also examines one of the "annuity puzzles" where individuals simultaneously hold a life insurance and an annuity contract. See Gottlieb (2012a) for an overview of this literature and a possible solution using prospect theory.

life expectancies as well as financial hardships from the high medical costs. Individual investors, and later firms, began purchasing the life insurance contracts from these terminally ill patients for lump sums that were greater than the cash values of the policies (Doherty and Singer, 2003b). Increases in medical technology that extended the life expectancy of HIV/AIDS patients brought a temporary end to the viatical settlement market. However, the development of life settlement firms in the early 2000s has allowed more individuals to access the market.

Gatzert et al. (2009) suggest that a secondary market for life insurance has the potential to reduce life insurer profits, as it provides an alternative to lapsing for individuals. Individuals who would have lapsed to obtain the cash value instead sell their policy to a life settlement firm. The life settlement firm will never lapse, thus the policy is guaranteed to eventually pay out. Because life insurance companies underwrite based on expected lapse rates, their profits depend on the unrequited income from lapsed policies to cover the losses from policies that pay out. Gatzert et al. (2009) finds that insurer profits would be decreased in the presence of a secondary market and can even be negative in the presence of adverse selection (i.e., individuals with good health surrender but individuals with poor health sell to the secondary market). Even if, in equilibrium, the insurance market captures some of the increased value of the policies, insurers will almost certainly lose money in the short term. This is because current policies have been underwritten under assumption of lapse rates consistent with limited secondary markets. If complete secondary markets suddenly appeared, insurers would not be able to change the premiums on these policies, many of which may last upwards of thirty years, and would experience losses on these policies.

In the counterfactual section of this paper, I implicitly assume perfect competition in the life insurance market. While there have been studies that have looked at competition in other insurance markets, in particular for property-liability (Joskow, 1973) and health (Dafny, 2010) insurance, little work has been done on competition in the life insurance industry. One exception is Brown and Goolsbee (2002), who analyze whether competition for term life insurance increased since the Internet became widely available. They find that by reducing search costs with insurance price comparison websites, competition increased in the life insurance industry.

1.3 Model and Assumptions:

My model is drawn from the classic life-cycle savings models which have been often used to study life insurance demand (e.g. Yaari, 1965; Fischer, 1973; Lewis, 1989), savings (e.g. De Nardi et al., 2010), annuity demand (e.g. Pashchenko, 2013), and more. The key extension of my model from the rest of the literature is the allowance for both observed and unobserved heterogeneity in the bequest motive.

Within each period, consumers receive utility from both current consumption as well as their expected bequests. Each period, consumer flow utility is given by

$$u(c_t, B_t, \Gamma_t) = (1 - \Gamma_t)\frac{c_t^{1-\rho}}{1-\rho} + \Gamma_t * p(age_t) * \frac{(B_t + \theta)^{1-\rho}}{1-\rho} \tag{1.1}$$

where c_t is the current period consumption, B_t is the amount the consumer would leave as a bequest should they die at the end of period t, and $p(age_t)$ is the probability the consumer dies at the end of period t given their age. Γ_t is the relative weight factor for bequest motives and is heterogeneous across consumers and across time. Each consumer's potential bequest amount B_t consists of their choice of savings A_{t+1} as well as the face value $face$ of their life insurance contract, if they have one. The choice set for consumers consists of whether or not to lapse on their life insurance as well as their current period consumption.

The timing of events within period t is:

1 : Consumer realizes current value for income I_t and bequest intensity Γ_t.

2 : Consumer jointly decides optimal consumption and whether or not to lapse on life insurance, ie $(c_t, L_t) \in S^*_{t+1}$ are chosen.

3 : Utility from consumption $\left((1-\Gamma_t)\frac{c_t^{1-\rho}}{1-\rho}\right)$ is gained.

4 : Consumer dies with probability $p(age_t)$ and receives bequest utility $\left(\Gamma_t \cdot \frac{(B_t+\theta)^{1-\rho}}{1-\rho}\right)$.

5 : Consumer does not die with probability $1-p(age_t)$ and continues to period $t+1$.

The consumer's problem is therefore to maximize utility V given their current state S_t. Which can be described by the Bellman equation

$$V(S_t) = \max_{c_t, L_t} \left\{ (1-\Gamma_t)\frac{c_t^{1-\rho}}{1-\rho} + \Gamma_t * \cdot p(age_t) \cdot \frac{(B_t+\theta)^{1-\rho}}{1-\rho} + (1-p(age_t)) \cdot \beta E[V(S_{t+1})] \right\} \tag{1.2}$$

$$S^*_{t+1} = \underset{c_t, L_t}{\operatorname{argmax}} \left\{ (1-\Gamma_t)\frac{c_t^{1-\rho}}{1-\rho} + \Gamma_t \cdot p(age_t) \cdot \frac{(B_t+\theta)^{1-\rho}}{1-\rho} + (1-p(age_t)) \cdot \beta E[V(S^*_{t+1})] \right\} \tag{1.3}$$

Subject to

$$S_t = \{L_{t-1}, A_t, I_t, \epsilon_t, age_t, face, cash_t, prem, married, kids\} \tag{1.4}$$

$$B_t = A_{t+1} + (1-L_t) \cdot face \tag{1.5}$$

$$A_{t+1} = A_t + I_t - c_t - prem \cdot (1-L_t) + cash_t \cdot L_t \cdot (1-L_{t-1}) \tag{1.6}$$

$$L_{t+1} \geq L_t \in (0,1) \tag{1.7}$$

$$cash_{t+1} = cash_t \cdot (1 + r) \tag{1.8}$$

$$\Gamma_t = \frac{e^{X\gamma+\epsilon_t}}{1 + e^{X\gamma+\epsilon_t}} \tag{1.9}$$

$$I_{t+1} = F(I_{t+1}|S_t) + e_t^I \tag{1.10}$$

$$\epsilon_{t+1} = \epsilon_t + e_t^\epsilon \tag{1.11}$$

$$e_t^\epsilon \sim N(0, \sigma_\epsilon^2) \tag{1.12}$$

$$e_t^I \sim N(0, \sigma_I^2) \tag{1.13}$$

Where:

Decision Variables:	
c_t	consumption in time t
L_t	lapse state in time t
Life Insurance Variables:	
$cash_t$	the cash value of the insurance plan in time t
$face$	the face value of the life insurance policy
$prem$	the annual premium amounts
Other State Variables:	
X	list of covariates for bequest motives
	(inclusive of an intercept term)
I_t	income (inclusive of capital gains) in time t
A_t	assets in time t
$p(age_t)$	death probability for a given consumer age
ϵ_t	unobserved heterogeneity in bequest motive share Γ_t
$Married$	binary indicator for if the consumer is married (in X)
$Kids$	number of children in the household (in X)
Estimated Parameters:	
Γ_t	flow utility share of the bequest motive in time t
ρ	CRRA parameters of consumption and bequests
σ_ϵ^2	variance of the unobserved part of Γ_t
θ	shift factor for bequests
Drawn Parameters:	
β	discount rate
σ_I^2	variance of Income process

Equation (4) presents the state space, Equation (5) defines the raw amount left as a bequest if the consumer dies at the end of period t, Equation (6) is the budget constraint, Equation (7) prevents "un-lapsing," Equation (8) describes the exogenous updating process of the cash value of the policy, Equation (9) defines the relative weight for bequests as a function of observable covariates, and Equations (10) and (11) describe the update processes for income and the unobservable heterogeneity in bequest motive. $F(I_{t+1}|S_t)$ is the consumer's expectation of next period's income given their current state. The process for estimating F is described in the next section.

S^*_{t+1} defines the optimal policy choice given S_t. The choice set for the consumer is the decision to lapse on their life insurance L_t and the choice of consumption this period c_t.[11] The decision to include *age* as a state variable is important and warrants discussion. Traditionally, differential death probabilities related to age would be incorporated through defining *age* as t and solving a finite - horizon problem with period - specific discounting. Including age as a state variable instead allows me to write the Bellman Equation (1) as time - independent. This increases the number of state variables, but allows for a single functional estimation of the value function (with *age* as an input) as opposed to a different value function for each $t \leq T$.[12]

As with all economic models I am making a number of implicit assumptions that are necessary to find a solution. The main assumption is that, in addition to ignorance of the initial purchase decision, once a policy is purchased, the set of other policies has no effect on the lapse decision. This assumption is made because I have no way of observing the potential purchase set for each consumer. With reference to the prior literature on lapsing, I cannot identify the decision set for consumers who are replacing one policy with another. Instead, I treat them as continuing coverage. Combining the purchase and lapse decision under one

[11]In practice, I estimate the choice of A_{t+1} by substituting for c_t using the budget constraint.

[12]I am not the first to use this modeling approach. De Nardi et al. (2010) use a similar method.

framework is an excellent avenue for future research. The other main component missing from my model is term insurance. In my data I observe if consumers have term insurance and the face amount, but I do not observe the premiums they pay or the length of the term. Additionally, I assume that consumers have perfect foresight with regards to their future marital status as well as the number of children in the household.[13]

1.4 Methods:

The theoretical literature on life insurance demand leans heavily on dynamic, life-cycle savings models. The empirical literature relies instead on reduced form models. The reason for the dichotomy between the is simply that fitting data to dynamic programming models is challenging and becomes exponentially more difficult the more structure is added to the model – structure that is necessary to properly examine life insurance. The traditional method of solving dynamic programming problems involves a discretization of the potential state space and then solving, simultaneously, for all of the possible combinations of states. This solution method returns an exact solution but, because the number of possible combinations of states increases exponentially with the number of state variables, becomes unfeasible as the state space grows large.

Conventional models deal with this challenge by assuming away some of the state space either through, often dramatic, discretization of continuous states or through defining one state as a function of another, losing valuable information in the process. Rather than sacrificing structure and getting an exact solution to the wrong problem, I investigate alternative solution methods. Specifically, I make use of approximate dynamic programming methods. Approximate dynamic programming consists of a broad class of "practical" solution methods including (but not limited to): Q-Learning, post-decision state variable estimation, value function approximation, Monte Carlo policy approximation, parameterized policy function

[13] Assuming, instead, that they have zero foresight does not change the resulting estimated value function in a meaningful way.

approximation, and many more.[14] In this article, I use a combination of value function approximation in the general solution algorithm and policy function approximation in calculating moments for the outer generalized method of moments (GMM) procedure.

Value function approximation, started initially by Bellman and Dreyfus (1959), has a long history, across several fields, as a solution to the "curse of dimensionality" commonly associated with the traditional method of numerically solving dynamic programming problems. The basic idea is simple – rather than solving the value function individually for every possible state, sample the state space and use a functional approximation instead. So long as the approximation structure is flexible enough, the standard recursive iterative solution will converge to a close approximation of the correct value function.[15] The most important decision when using value function approximation is what functional form the approximate value function will take. Value functions are notoriously misbehaving, so the functional form chosen must be flexible; but the whole purpose is to reduce computation time, so the functional form must also require a relatively small number of observations to fit.[16] Given these two competing requisites, artificial neural networks are a natural choice. So much so that a sub-discipline has arisen called "neuro-dynamic programming" about the use of artificial neural networks in approximate dynamic programming.

1.4.1 Approximate Dynamic Programming:

Dynamic programming began in earnest with Bellman's (1957) seminal text. The direct descendant of Bellman's framework are problems currently called Markov decision processes.[17] Bellman and Dreyfus (1959) was the first to consider the use of approximate value functions but, due largely to the lack of available computational power, little progress was

[14] Bertsekas and Tsitsiklis (1996), Judd (1998), and Powell (2007) provide excellent overviews with an operations, economics, and general emphasis respectively.

[15] Detailed discussion of convergence results is in Section 1.4.

[16] Judd (1998) provides a discussion of potential methods.

[17] Puterman (2014) is a modern reference for the interested reader.

made. Independently, computer scientists working in control theory, and using Hamilton-Jacobi equations in place of Bellman, developed their own branch of approximate dynamic programming called heuristic dynamic programming (Werbos, 1974). Tsitsiklis (1994) first made the general connection between approximate dynamic programming, reinforcement learning, and stochastic approximation.[18] Rust (1997) and Judd (1998) initially brought the techniques to economics.

The general solution algorithm for approximate value function iteration follows. For a general Bellman equation of the form

$$V(S_t) = \max_{S_{t+1}} \{u(S_t, S_{t+1})) + \beta E[V(S_{t+1})]\} \tag{1.14}$$

$$S^*_{t+1} = \operatorname*{argmax}_{S_{t+1}} \{u(S_t, S_{t+1})) + \beta E[V(S_{t+1})]\} \tag{1.15}$$

define $\hat{V}^j(S_t; \nu^j)$ to be the jth iteration of the parameterized (by vector ν) approximation of V. For simplicity (and relation to traditional recursive solution methods) let $\hat{V}^0(S_t; \nu^0) = 0$.[19] The iterative process for $\hat{V}^j(S_t; \nu^j)$ is

$$\hat{V}^{j+1}(S_t; \nu^{j+1}) \approx \max_{S_{t+1}} \Big\{u(S_t, S_{t+1})) + \beta E[\hat{V}^j(S_{t+1}; \nu^j)]\Big\} \tag{1.16}$$

where $\approx$ is an approximation operator.[20,21] From that, define

$$\hat{S}^j_{t+1} = \operatorname*{argmax}_{S_{t+1}} \Big\{u(S_t, S_{t+1})) + \beta E[\hat{V}^j(S_{t+1}; \nu^j)]\Big\} \tag{1.17}$$

[18] Tsitsiklis and Van Roy (1997) and Bertsekas and Tsitsiklis (1996) are further discussions of the idea.

[19] It does not actually matter what the initial guess is.

[20] This iterative process is often called the "Bellman Operator" in the Operations literature.

[21] I use a neural network structure with hidden layers of (5,5) neurons fitted via resilient backpropagation as my approximation operator.

If the approximation operator $\approx$ covers V in that there exists a finite parameter vector ν such that for any compact set $K \subset R^n$ and $\epsilon > 0$

$$\sup_{x \in K} |V(x) - \hat{V}(x;\nu)| \leq \epsilon \tag{1.18}$$

then, if the limit exists,

$$\lim_{j \to \infty} \hat{V}^j(S_t;\nu^j) \leq \approx V(S_t) \tag{1.19}$$

That is, if the approximating function and policy choice converge, then they converge to within the best approximation one could get (under the same functional form) if one knew the actual value function.[22] The traditional numeric solution method for dynamic programming problems is actually a limiting case of this algorithm where the approximation architecture $\approx$ is either a step function or a linear spline, depending on if linear interpolation is used. See Bertsekas and Tsitsiklis (1996) for a proof of the bounded convergence result.

Artificial neural networks are a natural choice for the general architecture of the functional approximation.[23] Neural networks were originally inspired by the biological structure of the central nervous system (McCulloch and Pitts, 1943). Neural networks consist of several layers of interconnected groups of nodes (or neurons) which each node taking input from all nodes in the previous layer. Hornik (1991) shows that any multi-layer, feed-forward structured neural network is a universal approximator.[24] That is, Equation (1.18) holds for every continuous function on compact subsets of $\mathbb{R}^n$.

[22] This is only true for problems with a discount rate $\beta < 1$. For problems with $\beta = 1$, called stochastic shortest path problems, a convergence result exists but is much weaker. For such problems, the limit will be within the error in approximation multiplied by the number of iterations until convergence. Bertsekas and Tsitsiklis (1996) discusses this result further.

[23] While neural networks are my preferred functional form, they are not necessary. Running the same general procedure using a random forest architecture and a polynomial architecture produce very similar results. However, both methods take significantly longer to converge. An OLS approximation architecture did not converge for the estimated parameters.

[24] Prior, Cybenko (1989) showed the same result specifically for sigmoid activation curves.

While rarely used in the decades following their initial development, neural networks saw a resurgence in the 1980s due to the increased availability of computational power and the development of the backpropagation algorithm which allowed for much quicker training (Werbos, 1974). Large scale neural networks are currently being used to approximate problems ranging from suggesting YouTube videos to recognizing handwriting to solving the traveling salesman problem for international shipping. The use of neural networks as an approximation architecture for value functions in dynamic programming is so popular it has developed its own sub-discipline called "neuro-dynamic programming," as discussed in Bertsekas and Tsitsiklis (1996).

1.4.2 GMM Procedure:

Because of the complexity of the value function, fitting the parameters is a non-trivial problem. Solving for the unobserved heterogeneity (e_t^ϵ) for a maximum likelihood procedure is unreasonable, so I use a moments-based method instead. For moments, I use the expected value of lapsing, the expected value of asset choice, the covariance of lapsing and each exogenous, observed state variable, and the covariance of asset choice and each exogenous, observed state variable. Given a solution $\hat{V}$, the natural method to calculate the sample moments $g(S_{i,t})$ for an observation $S_{i,t}$ is

$$g(S_{i,t}) = g\left(\int_{\epsilon_t} \underset{S_{i,t+1}}{\operatorname{argmax}} \left\{ u(S_{i,t}, S_{i,t+1})) + \beta \int_{\epsilon_{t+1}} \int_{I_{t+1}} \hat{V}(S_{i,t+1}) \partial \epsilon_{t+1} \partial I_{t+1} \right\} \partial \epsilon_t \right) \tag{1.20}$$

which is computationally difficult. However, the computational burden of calculating the moments can be dramatically reduced by approximating the policy function. For this, I fit a logit specification for the lapse choice, if it exists, and two OLS specifications for asset choice: one for each lapse choice. Formally

$$\bar{L}_{i,t} = \begin{cases} 0, & \text{if } L_{i,t-1} = 0 \ \& \ \dfrac{exp(X_i\beta_1)}{1+exp(X_i\beta_1)} < .5 \\ 1, & \text{if } L_{i,t-1} = 0 \ \& \ \dfrac{exp(X_i\beta_1)}{1+exp(X_i\beta_1)} >= .5 \\ 1, & \text{if } L_{i,t-1} = 1 \end{cases} \tag{1.21}$$

$$\bar{A}_{i,t+1} = \begin{cases} X_i\beta_2, & \text{if } \bar{L}_{i,t} = 0 \\ X_i\beta_3, & \text{if } \bar{L}_{i,t} = 1 \end{cases} \tag{1.22}$$

where X_i is the vector of states, S_t, interactions between states, and squares of states for individual i. $\beta_1, \beta_2, \beta_3$ are fit from the final $\hat{S}_{t+1}$ calculated from the sample used to fit the value function. The sample moments are then calculated from

$$g(\bar{S}_{i,t+1}) = g\left(\int_{\epsilon_t} \left\{\bar{L}_{i,t}, \bar{A}_{i,t+1}\right\} \partial\epsilon_{i,t}\right) \tag{1.23}$$

To address the possibility of local minima in the moment function, I use a global - global - local, three-stage GMM procedure. In the first stage, a genetic search algorithm iterates over a bounded parameter space looking for a global optimum to the unweighted moment function; this solution is then used to calculate an efficient weighting matrix. Using this weighting matrix, the genetic search algorithm again searches for a global optimum. Once the global search procedure is complete, the solution is used as a starting point for the third stage gradient-descent algorithm.[25]

Genetic search algorithms are a global search methodology based on evolutionary biology and commonly used in search problems with many parameters and over rough surfaces.[26] For minimization, the basic idea is to evaluate the function for numerous randomly selected

[25]The specific gradient-descent algorithm is the unbounded, limited-memory Broyden-Fletcher-Goldfarb-Shanno (LM-BFGS) method.

[26]See Golberg (1989) for an in depth discussion and convergence results.

points, drop the highest points, and then generate new points to test by combining (via weighted averaging) the lowest points.[27] Aside from the obvious benefit of avoiding local minima, genetic search algorithms are also of the class of "embarrassingly parallel" methods: this method drastically speeds up the entire search process by allowing me to parallelize the outermost algorithm.[28]

1.4.3 Counterfactual Exercise:

In order to estimate the welfare effects of fully-functioning, competitive secondary markets for life insurance, I perform a counterfactual exercise. To approximate the state of the world under perfectly competitive secondary markets, I adjust two variables for each of my observed consumers who own life insurance: $Premiums$ and the $Cash\,Value$. An increase in the prevalence of secondary markets would lead to a reduction in lapsation rates for policies. Because life insurance premiums are, in large part, underwritten based on assumed lapse rates, a reduction in lapse rates would lead to an increase in premiums. To my knowledge there is no work that estimates the effect on premiums secondary markets alone would have. However, Gottlieb and Smetters (2014) estimates premiums would increase by around 12% if there was no lapse-supported pricing. I therefore increase premiums by 12% for everyone in my sample. Because some consumers would still lapse on a policy even under perfect secondary markets, due to the policy not being actuarially a good deal yet, the actual premium increase will be lower than 12%. Thus, my welfare analysis can be viewed as a lower bound on the effect secondary markets would actually have.

The other state variable that I adjust is the cash value of the policy. Because I want to capture both the current and the option value increase, this adjustment is more difficult.[29] In order to capture the option value, I give each individual their own updating process for

[27] Additionally, the very best points are also kept. This process is called "elitism" in the literature.
[28] For the final gradient-descent section, the inner-most optimization step is parallelized but the neural network fitting is not.
[29] See Appendix A for a description of the solution method for the counterfactual.

the cash value of their policy. Where before the update process was

$$cash_{i,t+1} = (1 + r)cash_{i,t}$$

the new process is

$$cash_{i,t+1} = (cash_{i,t} - max\{0, AV_{i,t}\}) * (1 + r) + max\{0, AV_{i,t+1}\}$$

Where $AV_{i,t}$ is the discounted expected payout of the policy (assuming no lapsing) minus the discounted expected premium. In other words, the cash value in each period increases by the actuarial value (i.e., the sale price on a secondary market).[30] I calculate the welfare change under this new regime for everyone in my sample who owns life insurance and then calculate an individual compensation equivalent amount.

1.4.4 Expectation of Future Income:

In order to estimate an individual's expectation of future income, I make use of random forest regression techniques. Popular in machine learning, random forests are an ensemble learning method most commonly used for classification problems, but easily adjustable for regression as well. Developed to correct for the over-fitting inherent in single decision trees, random forests operate by constructing a multitude of decision trees via bootstrapping and then using the mean prediction.[31] Using a data-driven model selection method alleviates me from having to specify the "correct" income process.

The gap between econometrics and machine learning is largely one of predictive power vs. interpretability. The most popular methods in machine learning, such as random forests and neural networks, are often referred to as "black box" algorithms. This moniker references the

[30] In calculating the change this way I'm assuming the life settlement firm withdraws any accrued cash value immediately.

[31] See Breiman (2001) for a discussion of random forest predictors.

difficulty in interpretation of the estimated parameters relative to the linear (in some way) methods pervasive in economics.[32] However, the cost of opacity typically buys an increase in predictive power which, when leveraged properly, is extremely useful in economics. For instance, Belloni et al. (2012) uses the LASSO estimator to allow for non-linearity in the estimation of treatment effects in sparse models; Wager and Athey (2015) discusses the use of random forests to estimate models with heterogeneous treatment effects.

A good general rule, and one this paper follows, is to use classic econometric techniques when interpretability is required and the underlying data process is well-behaved, such as in the estimation of a policy functions, and to use machine learning techniques where prediction is the dominant requisite, such as in value function estimation or in the estimation of a consumer's expectation of future income.

1.5 Data:

The data used for analysis come from the 2007-2009 panel version of the Survey of Consumer Finances (SCF). Typically, the SCF is conducted as a three year cross-section. The 2007 portion of the 2007-2009 panel was conducted as part of the normal triennial cross-section. After the 2007 SCF was largely completed, the Federal Reserve Board designed and implemented the 2009 follow-up survey to provide a fuller picture of the effects of the intervening recession on households' finances than was available from aggregate data. The 2009 SCF's structure largely mirrored that of the 2007 survey so it is possible to construct parallel estimates of wealth, and specifically life insurance status, for households in both years. The 2007 - 2009 panel study is one of two sources of panel data in the SCF with the other occurring as a two year panel in 1983 - 1989 (Bucks et al., 2013).[33]

[32]Many machine learning algorithms are also fully non-parametric.

[33]The choice to collect the 2009 data did not occur until the final stages of the 2007 survey so no extra steps were taken as part of the 2007 survey to improve the odds of re-contacting respondents. Despite this, 97% of 2007 respondents were located in 2009 and 90% of those found agreed to take part again.

The SCF is designed to over sample wealthier households. In addition to myriad questions regarding household demographics, wealth, income, and pensions, the survey also contains a series of questions related to life insurance holdings. This panel structure allows for the identification of life insurance lapse decisions. The SCF has a long history of use in examining life insurance questions. Liebenberg et al. (2010) and Liebenberg et al. (2012) exploit the 1983 - 1989 panel to examine demand for policy loans and life insurance, respectively. Lin and Grace (2007) and Gutter and Hatcher (2008) use the cross-section version of the SCF to examine life insurance demand.[34] Summary statistics as well as subsetted between people who do and do not hold life insurance are presented in Table 1.1. Approximately 28% of my sample currently has whole life insurance and, of those that currently hold, I observe around 29% lapse. This may seem large, but my sample is over a two-year span and during the financial crisis. As shown in previous literature, consumers who hold whole life insurance are older, richer, and more likely to be married.

Variable	Full Sample		Have Insurance		No Insurance	
	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
Have Whole	0.28	0.45				
Lapse			0.29	0.45		
Assets ($000s)	1,126.98	2,694.54	1,731.78	2,347.43	884.46	2,347.43
Income ($000s)	127.15	254.27	173.11	342.92	108.72	205.63
Age	50.88	15.36	55.22	14.28	49.14	15.44
Face Value ($000s)			238.60	532.48		
Cash Value ($000s)			36.53	106.38		
Premium			1352.90	2,156.28		
Married	0.60	0.49	0.70	0.46	0.56	0.50
Kids	0.88	1.18	0.81	1.11	0.91	1.21
N	2,823		808		2,015	

Note: This table shows summary statistics for my sample separated for consumers who do and not not currently hold whole life insurance.

Table 1.1: Summary Stats

[34]Kennickell and Starr-Mccluer (1997) provides discussion on the benefits of using the panel version of the SCF as opposed to the repeated cross-section.

The full vector of parameters, not including the income and death processes, is

$$\{\rho, \theta, \gamma_0, \gamma_1, \gamma_2, \sigma_\epsilon, \beta, r\}$$

$$V(S_t) = \max_{c_t, L_t} \left\{ (1-\Gamma_t)\frac{c_t^{1-\rho}}{1-\rho} + \Gamma_t * \cdot p(age_t) \cdot \frac{(B_t+\theta)^{1-\rho}}{1-\rho} + (1-p(age_t)) \cdot \beta E[V(S_{t+1})] \right\}$$
$$\Gamma_t = \frac{e^{X\gamma+\epsilon_t}}{1+e^{X\gamma+\epsilon_t}}$$
$$X = (1, married, kids)$$

The first 6, $\{\rho, \theta, \gamma_0, \gamma_1, \gamma_2, \sigma_\epsilon\}$, are determined via the GMM procedure described in the preceding section. The last three are set exogenously. β is assumed to be .97 and r (the updating of cash values) is calculated from the data to be .05535.[35] The standard deviation of the error in consumer's income expectation is calculated to be \$23,105.[36]

1.6 Results:

1.6.1 Model Estimation:

Table 1.2 shows the estimated parameter values. The estimated value for the constant relative risk aversion parameter ρ of 4.986 falls within the range of values commonly found in the literature. For instance, De Nardi et al. (2010) find a value of 3.66. As would be expected, I do find a much smaller bequest shift parameter (\$101434 vs. \$273000, in De Nardi et al. (2010)). This divergence can be attributed to my inclusion of life insurance as a form of a bequest that De Nardi et al. (2010) does not pick up.

To my knowledge, I am the first to estimate a dynamic life-cycle model with both observed and unobserved heterogeneity in bequest motives, even without the inclusion of life insurance.

[35]This is not directly a return on investment as often a portion of the premium is added to the cash value.
[36]I treat capital gains as a part of income so this estimate is more realistic than it seems at first glance.

	Estimate	Interpretation
Estimated:		
ρ	4.986	Constant Relative Risk Aversion
	(.443)	
θ	101.434	Bequest Shift (in \$000s)
	(5.811)	
σ_e	1.116	SD of Bequest Heterogeneity
	(0.133)	
γ_0	-1.315	Intercept of Bequest Heterogeneity
	(1.793)	
γ_1	0.414	Coefficient for Marriage
	(0.171)	
γ_2	0.985	Coefficient for Kids
	(0.704)	
Exogenous:		
β	0.97	Discount Rate
r	0.053	Increase in Cash Value
σ_I	23.105	SD Income Process (in \$000s)

Note: This table shows the estimated and exogenously set parameters. Std. errors are in parentheses.

Table 1.2: Parameter Estimates

Thus, the rest of my parameter estimates are not directly comparable to previous results in the literature, however they do still have a direct interpretation. The point estimates are positive for the coefficients γ_1, γ_2 on both of the observable covariates (marriage and number of kids) for bequest motives; however, the coefficient (γ_2) on kids is not statistically significant. This implies, as predicted in prior literature, that married households place a stronger weight on leaving a bequest than single households.

Figure 1.1 shows the estimated policy function for my sample for the A_{t+1} decision at various levels of I_t and A_t as well as the actual choice made by my sample (smoothed via spline). As expected, a consumer's current net worth is a pretty good predictor of their future net worth. My model appears to do very well at predicting future assets for all levels of income.

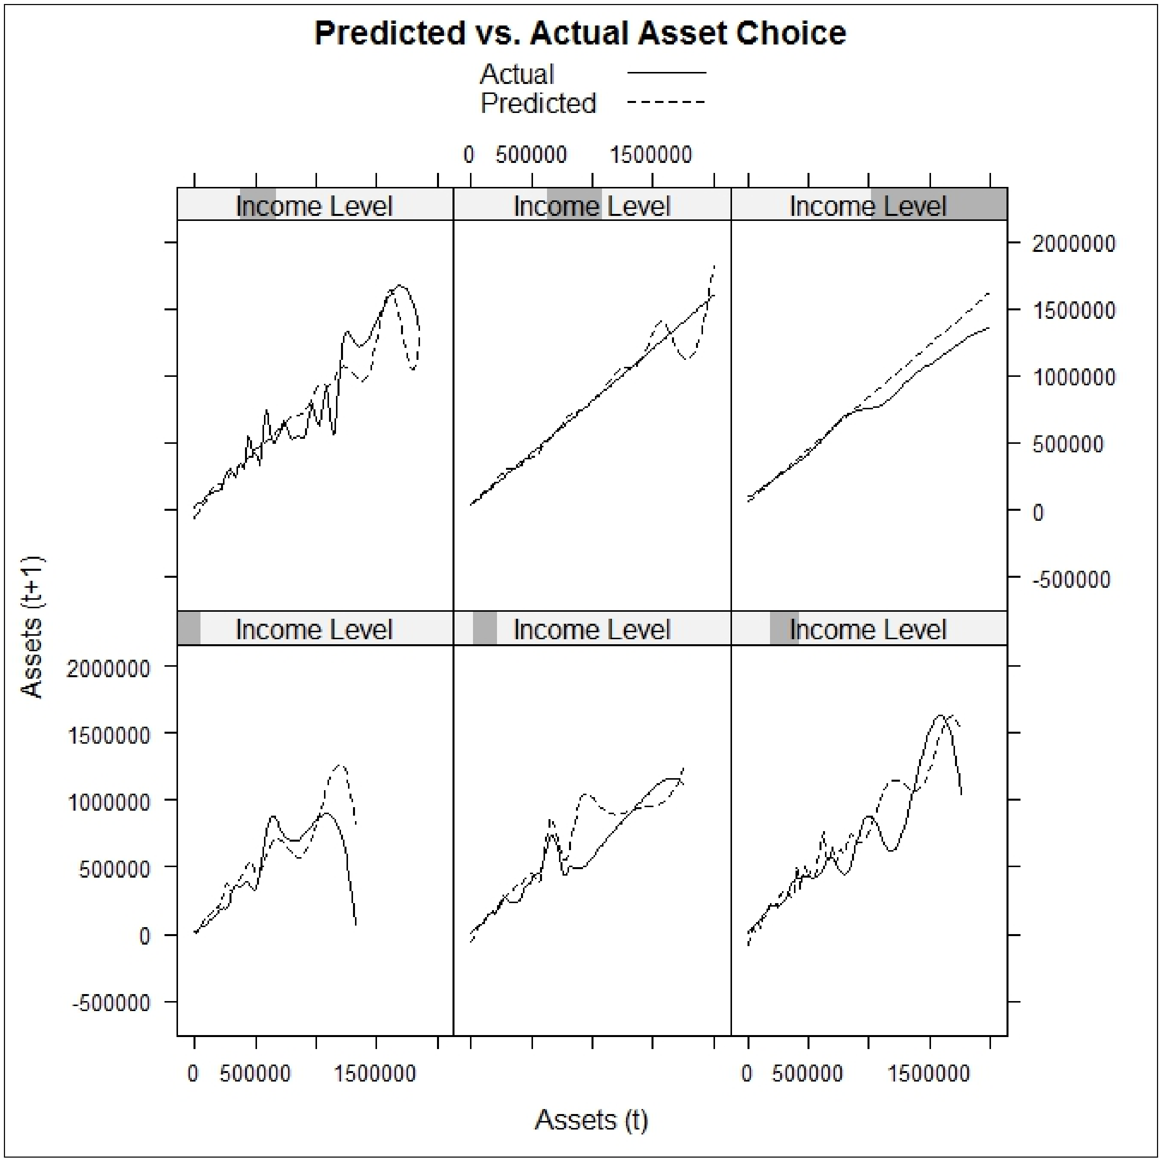

Note: This figure shows actual asset choice vs. the one predicted by my model for various levels of A_t and I_t. The upper breaks used for income are (in \$000s): 25.772, 39.191, 56.443, 76.465, 110.757, and 198.716. There are around 400 observations represented in each panel. There is a 10% overlap across each panel barrier. People with $A_t > 2,000,000$ and $I_t > 200,000$ are not graphed for clarity sake.

Figure 1.1: Actual vs. Predicted Asset Choice by Assets and Income Level

Figure 1.2 shows the predicted lapse choice for the portion of my sample who had life insurance at various levels of I_t and A_t as well as the actual choice made by my sample

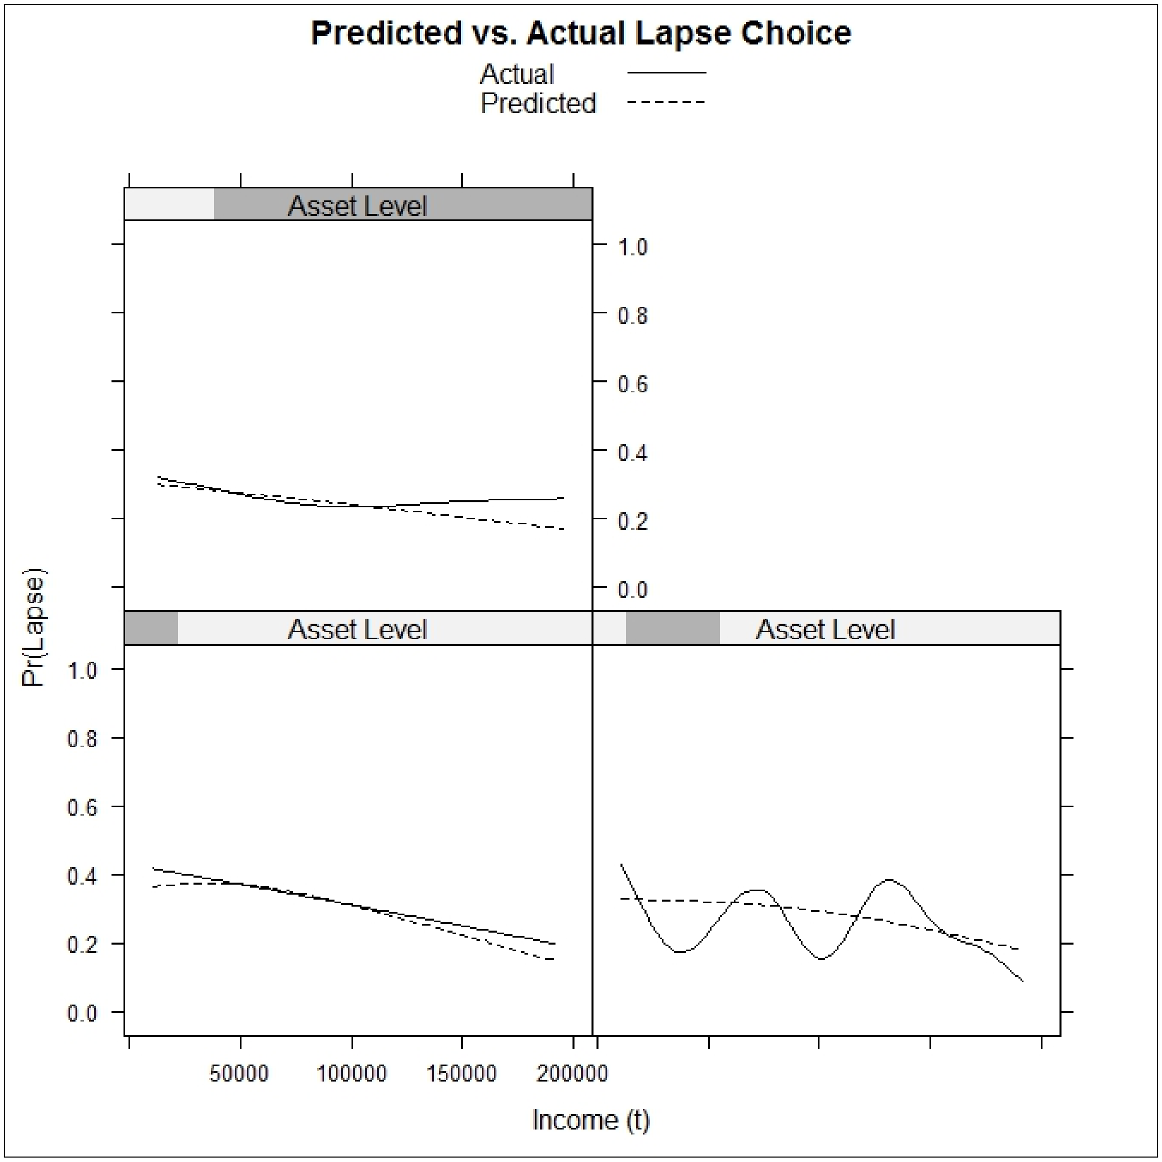

Note: This figure shows actual lapse choice vs. the one predicted by my model for various levels of A_t and I_t. The upper breaks used for assets are (in \$000s): 224.669 and 541.265. There are around 240 observations represented in each panel. There is a 25% overlap across each panel barrier. People with $A_t > 2,000,000$ and $I_t > 200,000$ are not graphed for clarity sake.

Figure 1.2: Actual vs. Predicted Lapse Choice by Assets and Income Level

(smoothed via spline). As predicted by the income-shock hypothesis, consumers with lower levels of assets and income, i.e., those more likely to be subject to liquidity constraints, are more likely to lapse on their life insurance policies. My model does a good job predicting lapse decisions for consumers with low net worth as well as high net worth, but appears unable to explain much of the variance in the lapse decision for consumers with net worth between $225,000 and $550,000. This is due to the income-effect dominating for the poorer consumers and the bequest-shock effect dominating for the rich consumers, each of which my model seems to pick up well. However, for middle to upper middle class consumers, the effects interact and my model seems to do a poor job picking that up.

Figure 1.3 shows the estimated policy function for my sample who has life insurance for the A_{t+1} decision at various levels of I_t and A_t and the $Face\ Value$ of their policy as well as the actual choice made by my sample (smoothed via spline). My model continues to perform well in prediction of asset choice; however, I am unable to predict much of the variation in asset choice for consumers with low to medium income, lower asset levels, and very large life insurance policies. Life insurance is intended to be insurance against the loss of income, so consumers who are "over-insuring" are potentially acting based on private information that my data, and therefore my model, are failing to account for.

Figure 1.4 shows the predicted lapse choice for the portion of my sample who had life insurance at various levels of I_t, A_t, and a binary indicator for the presence of children in the household as well as the actual choice made by my sample (smoothed via spline). The difference in how actual lapse rates change with respect to income for households that have children versus households that do not is interesting. Households with children for low and medium levels of net worth have a strikingly linear decrease in propensity to lapse as their income increases. A similar relation exists for consumers with children but with high assets; however, as would be expected, increasing levels of income have very little effect on lapse propensity, likely due to the absence of liquidity constraints.

Figure 1.5 shows the predicted lapse choice for the portion of my sample who had life insurance at various levels of I_t, A_t, and age as well as the actual choice made by my sample

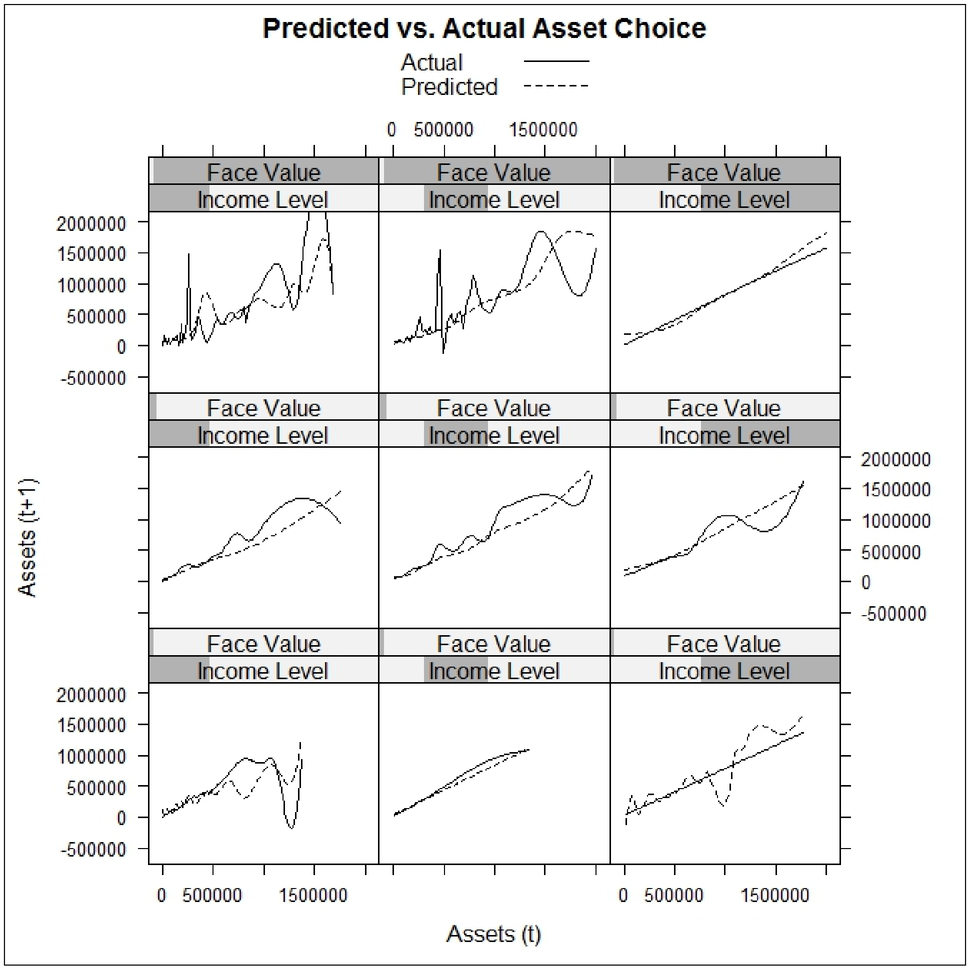

Note: This figure shows actual asset choice vs. the one predicted by my model for various levels of A_t, I_t, and Face Value for people with life insurance. The upper breaks used for income are (in $000s): 58.573, and 97.125. The upper breaks for Face Value are (in $000s): 38.210 and 100.010. There are around 230 observations represented in each panel. There is a 25% overlap across each panel barrier. People with $A_t > 2,000,000$ and $I_t > 200,000$ are not graphed for clarity sake.

Figure 1.3: Actual vs. Predicted Asset Choice by Assets, Income, and Face Value

(smoothed via spline). Age appears to be negatively correlated with lapse rates and my model does a good job picking this up with the exception of the richest rich.

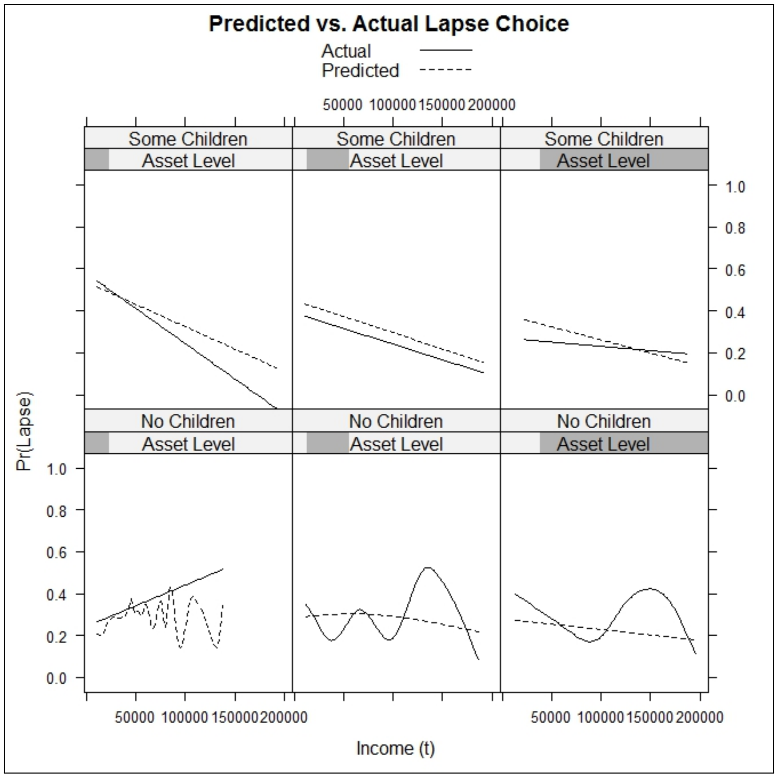

Note: This figure shows actual lapse choice vs. the one predicted by my model for various levels of A_t, I_t, and the number of children in the household for people with life insurance. The upper breaks used for assets are (in \$000s): 224.699, and 541.265. There are between 190 and 270 observations represented in each panel. There is a 25% overlap across each panel barrier. People with $A_t > 2,000,000$ and $I_t > 200,000$ are not graphed for clarity sake.

Figure 1.4: Actual vs. Predicted Lapse Choice by Assets, Income, and Number of Children

1.6.2 Counterfactual Analysis:

Using the estimated parameters from Table 1.2, I re-calculate each individual's value under a scenario where premiums increase by 12% and the cash value in time period t increases by $max\{0, Act.Value_t\}$.[37] Interestingly, everyone in my sample is better off. I find that the average utility, of policy-holders, increases, in one-time compensation units, by

[37] The premium increase estimate is derived from Gottlieb and Smetters's (2014) estimate of premium increases if there was no lapsing. Because some lapsing would still occur under perfect secondary markets, this is a conservative estimate. See Appendix A for description of the solution method used to solve for the value function in the counterfactual estimation.

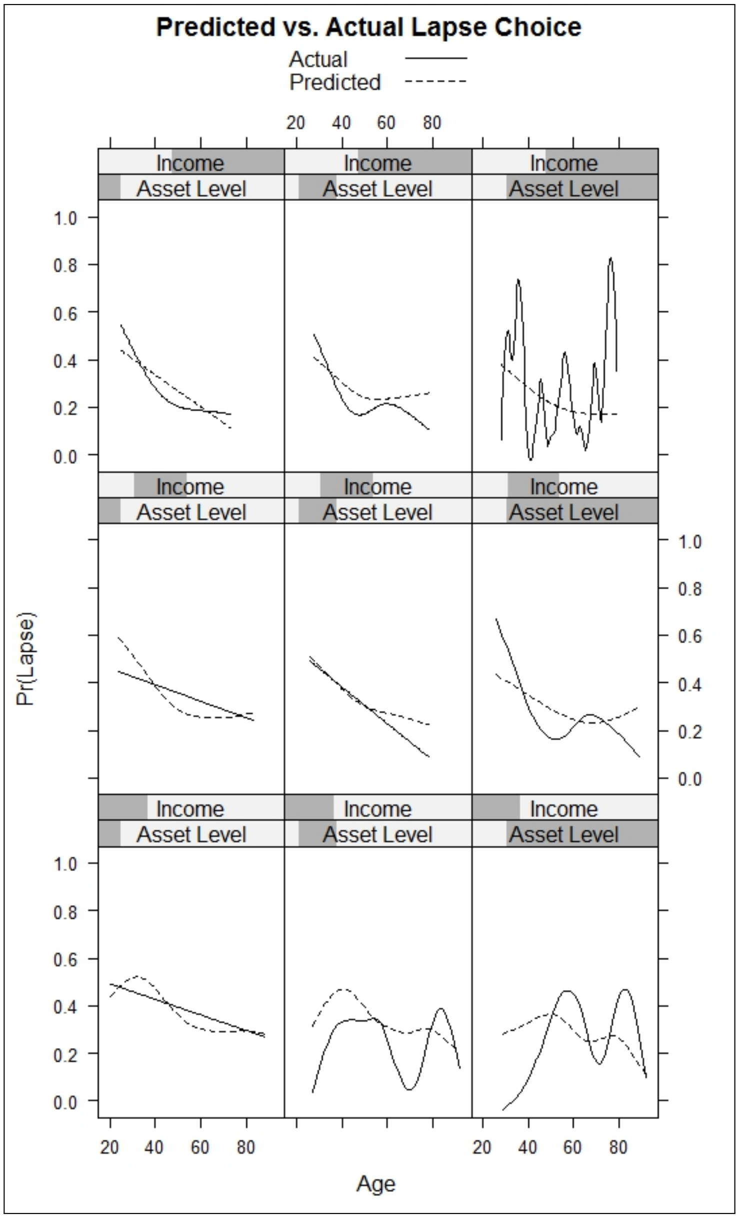

Note: This figure shows actual lapse choice vs. the one predicted by my model for various levels of A_t, I_t, and age for people with life insurance. The upper breaks used for assets are (in \$000s): 224.699, and 541.265. The upper breaks used for income are (in \$000s): 58.573, and 97.125. There are between 190 and 270 observations represented in each panel. There is a 25% overlap across each panel barrier. People with $A_t > 2,000,000$ and $I_t > 200,000$ are not graphed for clarity sake.

Figure 1.5: Actual vs. Predicted Lapse Choice by Assets, Income, and Age

$1,346. This means that the increase in the value of life insurance policies, even considering the increased premiums, is, on average, $1,346. This welfare increase is a lower bound for two reasons: I estimate the change in premiums based on the stoppage of all lapsing, while some would still occur under secondary markets, and I do not consider the, strictly positive, welfare benefits on the extensive margin. That is, I do not account for benefits gained by people who do not purchase a life insurance policy under the status quo, but would do so if secondary markets were more prevalent and accessible.

Figure 1.6 shows the single-period compensation equivalent plotted against the actuarial value for my sample at various levels of A_t, I_t, and age. While the compensation value moves partially congruently with the estimated actuarial value of the each individual's policy. While Figure 1.6 appears to show a positive correlation with compensation value and net worth, this relation is illusional and due to a positive correlation of the face value of policies and the consumer's net worth. To account for this omitted variable I fit an OLS specification to the compensation-equivalent welfare increase. Table 1.3 shows this result.

Interestingly, poorer and younger consumers benefit the most from the changed policy. Young consumers are relatively better off because the increased option value simply has more expected time to play out. That is, younger consumers are going to be subject to more shocks over the rest of their life than older consumers will. Because both income and bequest motive shocks are persistent, this increases the variance of the future insurance value of the contract, which necessarily increases the option value. Poorer consumers are relatively better of than richer (in both assets and income) consumers because they are more likely to lapse on their policies due to income constraints.

Figure 1.7 shows the difference in predicted lapse rates under the status quo versus the counterfactual. As would be expected, "lapse" rates increase. This is due to my definition of lapsing being inclusive of selling a policy. In reality, there are likely people who would be worse off under the counterfactual, due to the premium increase, and may lapse because of

	Dependent variable:
	Income Differential
Assets (\$000s)	-3.061^{***}
	(1.303)
Income (\$000s)	-1.440^{***}
	(0.352)
Age	-10.269^{***}
	(2.464)
Face Value (\$000s)	0.143
	(4.029)
Cash Value (\$000s)	2.716^{***}
	(0.429)
Annual Premium	-0.855^{***}
	(0.011)
Married	31.248
	(89.611)
Kids	28.736
	(70.805)
Constant	602.011^{***}
	(27.498)
Observations	8082

Note: *p<0.1; **p<0.05; ***p<0.01. This table shows the compensation-equivalent change in utility for the counterfactual fitted via OLS. Std. errors are in parentheses.

Table 1.3: Counterfactual

it. However, in my sample, everyone is strictly better off from the policy and therefore the increased lapse rates are entirely due to consumers being able to sell their policies.

Figure 1.8 shows the compensation equivalent welfare increase split between the increase due to consumers selling their current policies vs. the value increase due to the option value of being able to sell the policy later. The initial increase dominates for younger consumers with high value policies. This may not seem intuitive at first, surely younger consumers have a lower actuarial value for their policies? While true, younger consumers, in particular those with higher levels of assets, simply don't value the insurance as much and thus, as shown in Figure 1.5, are much more likely to lapse. The option value dominates the initial value for consumers who have smaller policies.

1.7 Conclusions:

Relative to the majority of insurance policies, life insurance contracts can be exceptionally long term, often in excess of 30 years. These policies are also typically written with a level premium structure. Because death risk increases with age, the actuarial value of a life insurance policy increases over time and often becomes positive towards the end of the policy. Life insurance is also different from other forms of insurance in that the payout is not received by the insured, but instead by a beneficiary. This means that value of the insurance to the insured must be discounted through a bequest motive. If bequest preferences are dynamic and subject to unexpected shocks, such as through divorce or the aging of a child, the value to the insured of a life insurance contract can become negative later in the policy even when the actuarial value is positive. Thus potential gains from trade could be acquired from a secondary market for life insurance policies. A market where the firm pays the policy holder a lump sum up front (and all of the subsequent premiums) for the policy holder to name the firm as the beneficiary of the policy. While these markets do exist, they are limited and controversial. In this paper, I proposed and estimated a structural model for the life insurance lapse decision. I then performed a counterfactual analysis with competitive secondary markets and found them to be Pareto improving for my sample with an average value increase to consumers of $1,346. Younger and poorer consumers benefited the most.

This paper makes three contributions: The primary contribution of this paper is to bring together the theoretical and empirical strands of life insurance lapsation estimation under a structural framework, allowing for counterfactual policy analysis. Dynamic life-cycle models are the natural method of examining decisions about the long-term, bequest motivated, decisions involving the purchase and lapsing of life insurance policies. While the theoretical literature has embraced this framework, the empirical literature still relies almost exclusively on a reduced form estimation approach. This has greatly limited researcher's ability to conduct counterfactual exercises to examine various policy proposals regarding life insurance including, but not limited to, increasing the prevalence and accessibility of secondary markets.

The second contribution is the addition to the academic discussion of the welfare benefits of secondary markets in life insurance, finding them to be largely welfare increasing. However, the defining feature of whole life insurance that makes secondary markets viable, the long-term nature of the contracts, is also a major barrier to implementation. Because insurance firms have underwritten existing contracts under the assumption that lapse rates would remain stable, the introduction of functioning secondary markets would lead to losses on many current policies. This has led to many insurance firms banding together to lobby against the widespread opening of secondary markets.

The third contribution is to exemplify, to the economics community, the potential of neural networks as an approximation architecture in solving complicated dynamic programming problems through the method of approximate dynamic programming. This paper closed the schism in the complexity of modeling framework between the theoretical and empirical literatures examining life insurance lapse behavior. The reason this schism existed is because of the computational difficulty involved in solving the complicated dynamic programming problem necessary to structurally model consumer behavior toward life insurance. Rather than sacrificing necessary structure, I drew on the machine learning and operations literature for an alternative solution method. I do not advance the methods used any further, however I am, to my knowledge, the first in the economics literature to use neuro-dynamic programming methods, which potentially have broad applications beyond life insurance lapse decisions.

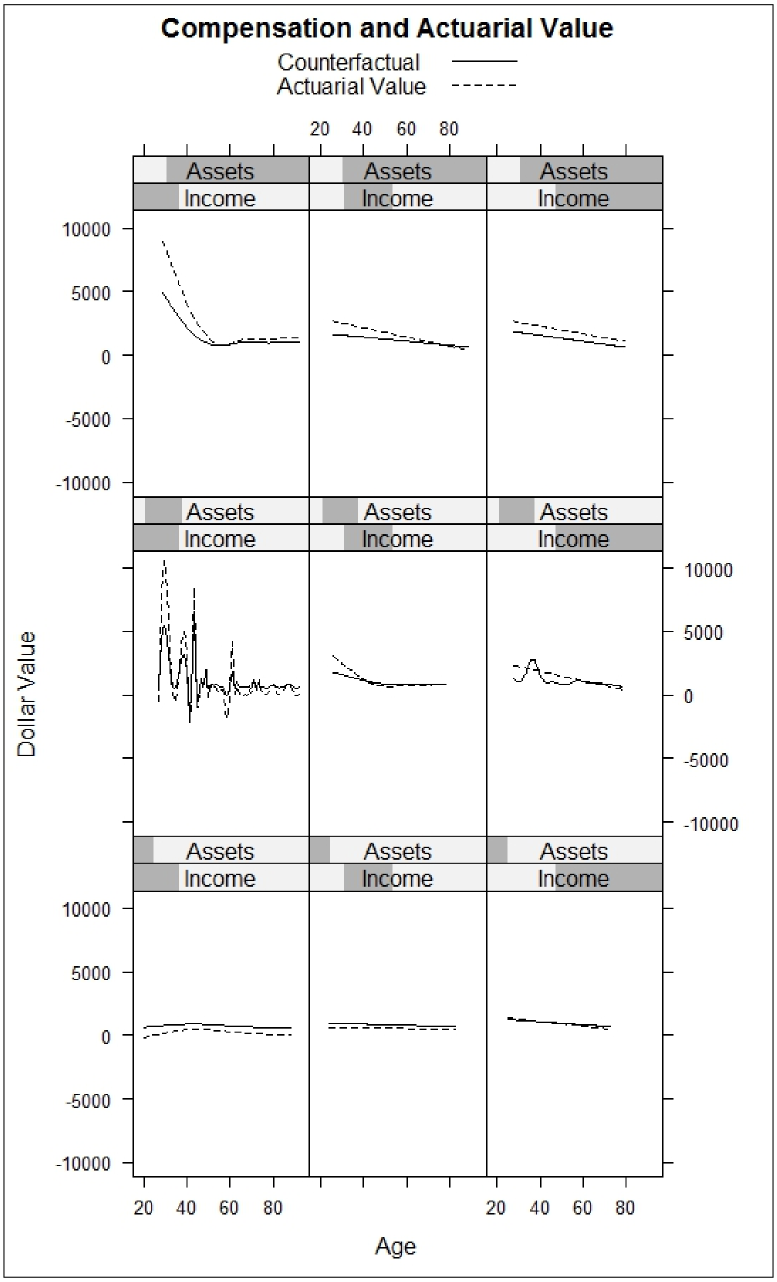

Note: This figure shows the compensation equivalent plotted against the actuarial value for my sample at various levels of A_t, I_t, and age. The upper breaks used for assets are (in \$000s): 224.699, and 541.265. The upper breaks used for income are (in \$000s): 58.573, and 97.125. There are between 190 and 270 observations represented in each panel. There is a 25% overlap across each panel barrier. People with $A_t > 2,000,000$ and $I_t > 200,000$ are not graphed for clarity sake.

Figure 1.6: Compensation Equivalent by Actuarial Value

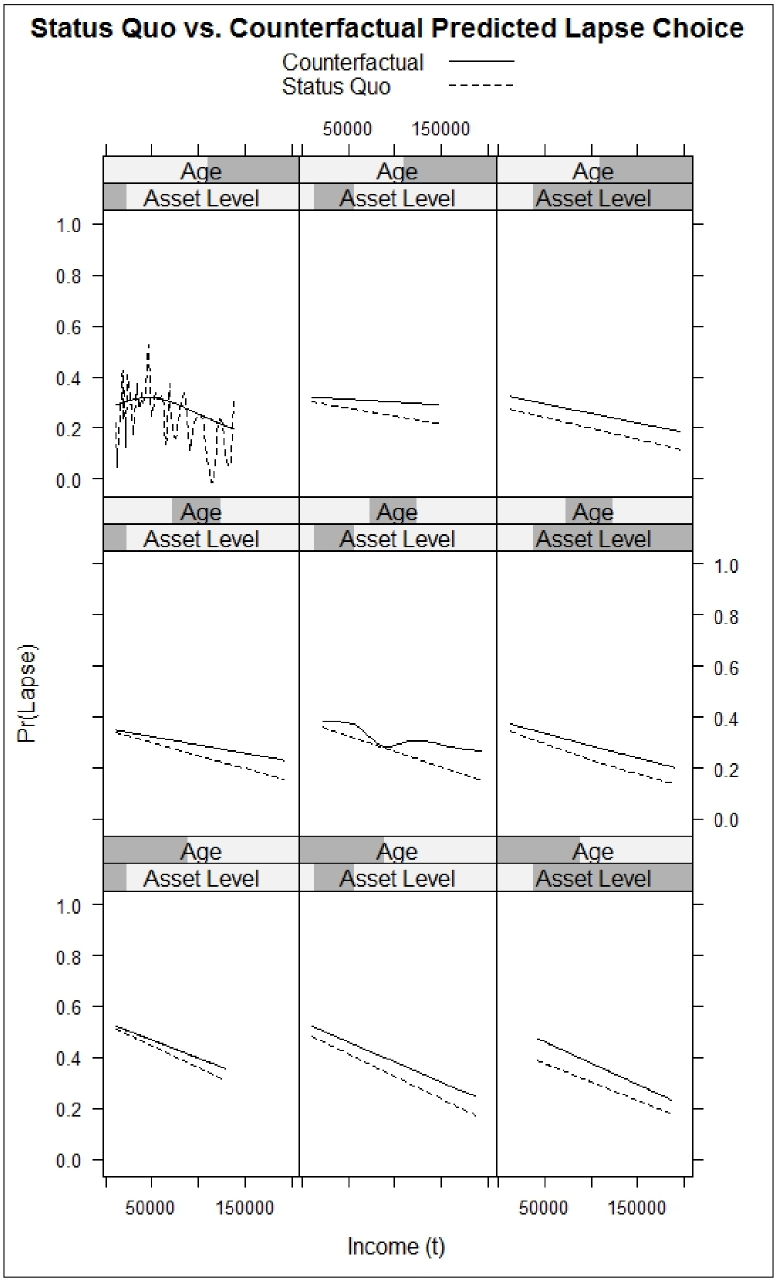

Note: This figure shows predicted lapse choice under the status quo vs. the counterfactual for various levels of A_t, I_t, and age. The upper breaks used for assets are (in \$000s): 224.699, and 541.265. The upper breaks for age are 50, 62, and 95. There are between 190 and 270 observations represented in each panel. There is a 25% overlap across each panel barrier. People with $A_t > 2,000,000$ and $I_t > 200,000$ are not graphed for clarity sake.

Figure 1.7: Predicted Lapse Choice under Status Quo and Counterfactual

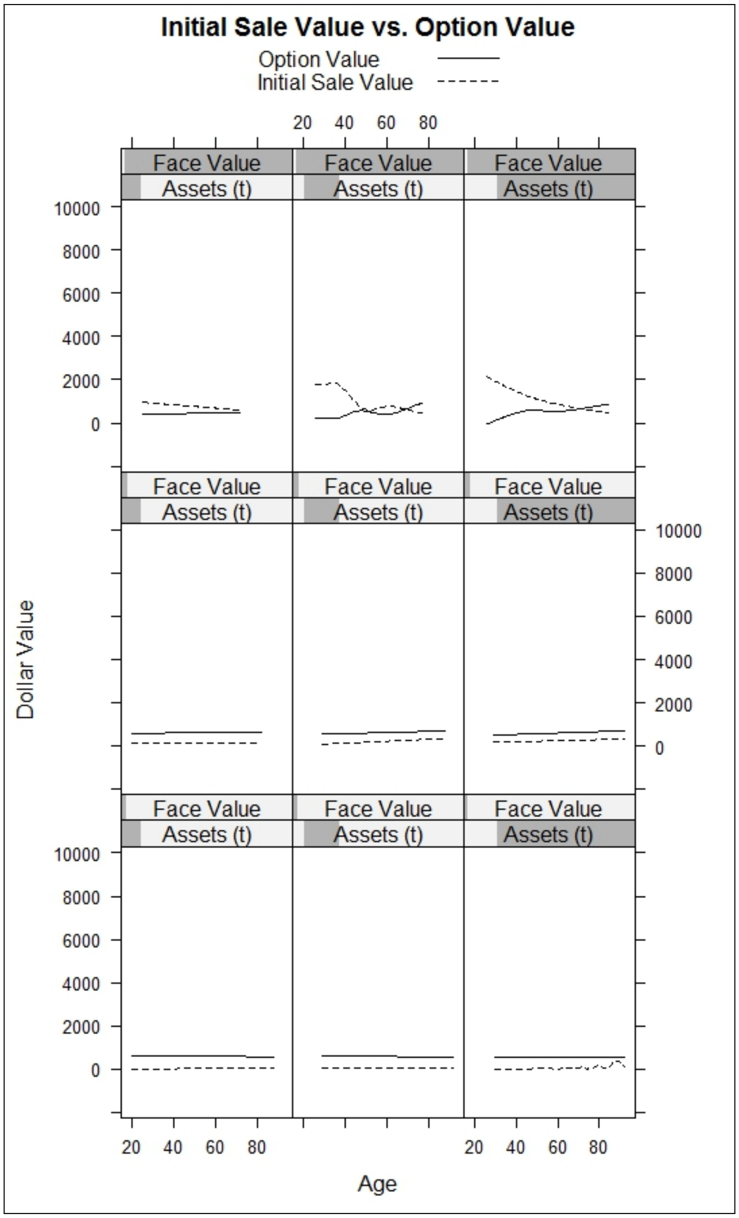

Note: This figure shows the compensation equivalent welfare increase split between the increase due to consumers selling their current policies vs. the value increase due to the option value of being able to sell the policy later for various levels of A_t, I_t, and age. The upper breaks used for assets are (in \$000s): 224.699, and 541.265. There are between 190 and 270 observations represented in each panel. The upper breaks for age are 50, 62, and 95. There is a 25% overlap across each panel barrier. People with $A_t > 2,000,000$ and $I_t > 200,000$ are not graphed for clarity sake.

Figure 1.8: Compensation Equivalent: Option Value vs. Current Value

Appendices:

A Counterfactual Estimation Procedure:

In order to estimate the welfare effects of fully functioning, competitive secondary markets for life insurance, I perform a counterfactual exercise. To approximate the state of the world under perfectly competitive secondary markets, I adjust two state variables for each of my observed consumers who own life insurance: Premiums and the Cash Value. An increase in the prevalence of secondary markets would lead to a reduction in lapsation rates for policies. Because life insurance premiums are, in large part, underwritten based on assumed lapse rates, a reduction in lapse rates would lead to an increase in premiums. To my knowledge there is no work that estimates the effect on premiums secondary markets alone would have. However, Gottlieb and Smetters (2014) estimates premiums would increase by around 12% if there was no lapse-supported pricing. I therefore increase premiums by 12% for everyone in my sample. Because some consumers would still lapse on a policy even under perfect secondary markets, due to the policy not being actuarially a good deal yet, the actual premium increase will be lower than 12%; thus, my welfare analysis can be viewed as a lower bound on the effect secondary markets would actually have.

The other state variable that I adjust is the cash value of the policy. Because I want to capture both the current and the option value increase, this adjustment is more difficult. In order to capture the option value, I give each individual their own updating process for the cash value of their policy. Where before the update process was

$$cash_{i,t+1} = (1 + r)cash_{i,t}$$

the new process is

$$cash_{i,t+1} = (1+r) \cdot (cash_{i,t} - max\{0, AV_{i,t}\}) + max\{0, AV_{i,t+1}\}$$

Where $AV_{i,t}$ is the discounted expected payout of the policy (assuming no lapsing) minus the discounted expected premium. In other words, the cash value in each period increases by the actuarial value (i.e., the sale price on a secondary market).[38]

Because the updating of the cash value is no longer time consistent, solving for the change in welfare becomes more difficult. I address this by including a policy function approximation approach in my solution algorithm. The idea is similar to the value function approximation method described above with a few slight modifications: Define

$$\hat{S}_{t+1}(S_t; \hat{\phi}) \approx S^*_{t+1}(S_t|\hat{V}) = \underset{S_{t+1}}{\text{argmax}} \Big\{ u(S_t, S_{t+1})) + \beta E[\hat{V}(S_{t+1})] \Big\} \tag{1.24}$$

$$\hat{S}_{t+j}(S_t; \hat{\phi}) \equiv \hat{S}_{t+j}(\hat{S}_{t+j-1}(...\hat{S}_{t+2}(\hat{S}_{t+1}(S_t)))) \tag{1.25}$$

That is, $\hat{S}_{t+1}(S_t; \hat{\phi})$ is an approximation of the actual policy choice given an approximate value function $\hat{V}$. Equation (24) defines the estimation of future states given the current estimated policy function and current state. The derivation of $S^*_{t+1}(S_t|\hat{V})$ is the same as in the ADP solution method above. The updating procedure of the approximate value function changes however. Before

$$\hat{V}^{j+1}(S_t; \nu^{j+1}) \approx \max_{S_{t+1}} \Big\{ u(S_t, S_{t+1})) + \beta E[\hat{V}^j(S_{t+1}; \nu^j)] \Big\}$$

[38] In calculating the change this way I'm assuming the life settlement firm withdraws the cash value immediately.

Whereas under the new method:

$$\hat{V}^{j+1}(S_t;\nu^{j+1}) \approx \sum_{i=0}^{\infty}\left\{\beta^i E\left[u(\hat{S}_{t+i},\hat{S}_{t+i+1}))\right]\right\}$$

That is, the previously estimated value function is not used in the calculation of the value of the estimated policy derived from the previously estimated value function. For the policy function approximation architecture I use the same form described in the GMM procedure:

$$\bar{L}_{i,t} = \begin{cases} 0, & \text{if } L_{i,t-1} = 0 \ \& \ \dfrac{exp(X_i\beta_1)}{1+exp(X_i\beta_1)} < .5 \\ 1, & \text{if } L_{i,t-1} = 0 \ \& \ \dfrac{exp(X_i\beta_1)}{1+exp(X_i\beta_1)} >= .5 \\ 1, & \text{if } L_{i,t-1} = 1 \end{cases} \tag{1.26}$$

$$\bar{A}_{i,t+1} = \begin{cases} X_i\beta_2, & \text{if } \bar{L}_{i,t} = 0 \\ X_i\beta_3, & \text{if } \bar{L}_{i,t} = 1 \end{cases} \tag{1.27}$$

where X_i is the vector of states, S_t, interactions between states, and squares of states for individual i.

This method and the method of only approximating the value function converge to essentially the same values.[39] However, this new method allows me to calculate the value of a policy with individual, time-inconsistent updating of the cash values.

[39]The combined method actually converges faster as well. I use the slower version as my preferred specification because of the greater ease of exposition.

Chapter 2

Sunk Costs and Signaling: Two-Part Tariffs in Life Insurance[1]

2.1 Introduction:

There are large, upfront, fixed costs to writing a life insurance policy. Both agent commission and direct underwriting costs (e.g. fees for physicals and blood tests) are fully paid a few years into contracts that can last 10-30 years. Because of these upfront costs, insurers can actually lose money on policies when the consumer lapses early into the contract, even if no death benefit is ever paid out. Thus, to properly price contracts, insurers must estimate lapse risks. However, because consumers often have more knowledge about their lapse likelihood than the insurer, asymmetric information arises and room for a signaling mechanism exists. In this article, we develop a model of life insurance pricing under heterogeneous lapse behavior with asymmetric information about lapse likelihood. We establish the existence of a separating equilibrium under a menu of contracts containing an optional two-part tariff. We then show the consumer's choice serves as a signaling device for private information on lapse likelihood. Using detailed, policy-level data on the practice of life insurance backdating as an example of our proposed optional two-part tariff, we empirically test our model's prediction of consumer self-selection.[2] We use a control function approach to separately identify selection effects from potential sunk cost fallacy. We find that consumers who choose to take

[1]This chapter is joint work with Jim M. Carson, Robert E. Hoyt, and Krzysztof Ostaszewski.

[2]Backdating occurs when the contract's start date is earlier then its application date. Most often, policies are backdated to align financial documents or set a specific premium payment date. However a smaller number (6% of all policies in our sample) are backdated to save age, which is described in detail later. Throughout this paper we are only considering policies backdated to save age.

part in the two-part tariff by backdating their policies are less likely to lapse, due to both self-selection and sunk costs.

Lapse rates in life insurance are enormous. Between 1991 and 2010, $29.7 trillion of new individual life insurance coverage was issued in the United States. During this same time period, $24 trillion of coverage lapsed. 85% of term life insurance policies never pay a death benefit. More surprisingly, 74% of term life policies sold to people at age 65 fail to pay a death benefit (Gottlieb and Smetters, 2014).[3] For every policy that ends in death or term maturation, 36 policies lapse due to nonpayment of premiums (Purushotham, 2006). We focus our analysis on term life insurance and, though the vast majority of literature on insurance lapsing is focused on whole life insurance, much of the theory applies to term insurance as well. The literature on lapse behavior is well-developed and can be condensed into: preference shocks, income shocks, policy replacement, and non-expected utility explanations.

Preference shocks refer to any number of situations where the private value of the life insurance contract has changed (Fang and Kung, 2010; Liebenberg et al., 2012; Fei et al., 2015).[4] Examples include: divorce, death of a spouse or child, children becoming self-sufficient, increase in spousal income, etc. Income shocks refer to consumers experiencing a negative shock to income and thereby having insufficient funds to pay premiums. The effect of an income shock on lapse rates is stronger in whole life insurance due to the presence of a surrender value. This is referred to in the literature as the emergency fund hypothesis (Linton, 1932; Outreville, 1990; Kuo et al., 2003). The aptly named policy replacement hypothesis refers to consumers who lapse on one policy because they found a better one (Outreville, 1990; Carson and Forster, 2000). The interest rate hypothesis is a specific case of the policy replacement hypothesis where the driver of better available policies is a change

[3] Whole life insurance, which doesn't expire, has similar overall lapse rates and lower per-year lapse rates.

[4] This reason for lapsing generally only covers negative shocks to the insurance value of the policy. Positive shocks to preferences are typically subsumed by the policy replacement category.

in expectations of future interest rates (Schott, 1971; Pesando, 1974; Kuo et al., 2003).[5] The final category of research on lapse rates focuses on non-expected utility models of consumer behavior and how these various behavioral assumptions can influence the decision to lapse on a policy (Shefrin, 2002; Mulholland and Finke, 2014; Gottlieb and Smetters, 2014).[6] Our theory requries that consumers potentially have some private knowledge about their lapse proclivity. Of these four main determinants, consumers most likely have private knowledge about income shocks and preference shocks.

Though asymmetric information on lapse risk initially appears different than the canonical study of asymmetric information on loss probability, the same intuition applies. In our case, we observe that life insurance incurs a front-loading of underwriting-related costs both through direct underwriting costs and agent commission. It is expected that the costs of underwriting will be recovered over the long life of the contract. However, when a consumer lapses on a policy early into the contract (e.g., in the first several years), the insurer is unable to recoup the entire underwriting cost. Thus, consumers with a higher likelihood of lapsing have a higher (expected) average cost per year because the fixed cost of underwriting is spread over a smaller time span. If consumers have private knowledge about this cost, the same adverse selection issues arise.

Much has been written on possible asymmetric information issues in life insurance, though the extant literature focuses solely on loss as opposed to lapse. Cawley and Philipson (1999) show the existence of a negative correlation between risk and coverage in life insurance that runs counter to the testable implications set forth by Chiappori et al. (2006). Cutler et al. (2008) explore reasons for this apparent "advantageous selection" citing a negative correlation for risk aversion and risky behaviors. Gottlieb (2012b) suggests it is unlikely

[5]Because insurance premiums are collected, and invested, long before benefits are paid out, expectations about interest rates play an important part in the determination of premium rates.

[6]For an excellent, detailed analysis of these hypotheses the interested reader is referred to Eling and Kochanski (2013).

there is an adverse selection problem in life insurance, as it is priced close to actuarially fair levels. However, the possible presence of simultaneous advantageous selection, which produces opposite observable correlation as adverse selection, clouds this analysis (De Meza and Webb, 2001). Under advantageous selection, individuals who have lower risk also tend to be more risk averse. This produces a situation where less risky individuals are more likely to purchase insurance due to their higher risk aversion even though they have a lower probability of experiencing a loss.[7]

The main difference in lapse risk versus loss risk is that the common methods of addressing asymmetric information do not readily apply: deductibles and copays will not cause consumers to self-select into their proper lapse-risk group. Instead we show insurers can offer a menu of contracts where the variety lies in combinations of an initial fee and recurring premiums, i.e., a menu of two-part tariffs. We show that consumers who choose to pay the higher up-front fee and lower recurring premiums are signaling their intention to persist and not to lapse early in the contract period.

Two-part tariffs, originally examined by Gabor (1955) and Bowman (1957), are central to the literature concerning price discrimination. The simplest form of second-degree price discrimination, two-part tariffs are commonly used across numerous industries. The first formal analysis of this form of price discrimination was Oi (1971) which examines the optimal pricing structure for amusement parks (Disneyland in particular); Schmalense (1981) expands on this work. Blackstone (1975) looks into potential two-part tariffs for Electrofax copying machines. Schmalense (2015) looks at the "razor-and-blades" pricing strategy.[8] The extreme version of a two-part tariff is "buffet pricing," where the per-period price is set to zero and the entire payment is subsumed by the tariff (Nahata et al., 1999). Prior work examined on why life insurance contracts tend to be actuarially front-loaded, but the extant literature has

[7]For a broad discussion of adverse and advantageous selection, see Einav and Finkelstein (2011).

[8]For a more in-depth discussion of the extant literature on two-part tariffs, the interested reader is referred to Stole (2007); for a discussion of the theory the reader is referred to Tirole (1988).

failed to consider that optional front-loading acts as a signaling device (Hendel and Lizzeri, 2003; Hofmann and Browne, 2013).[9]

A version of an optional two-part tariff pricing method exists in the life insurance industry, called backdating. A life insurance contract is considered backdated when the insurance contract bears a start date that is prior to the application date.[10] The consumer chooses to pay for coverage for the time prior to their application and, because the consumer is still alive at the time of application, no direct benefit will be paid for this prior coverage. Rather, consumers do this to "save age." Because life insurance policies feature level premiums that are based largely on the age of the applicant at the beginning of the contract, backdating to save age lowers the per-period premium paid throughout the life of the contract.[11] Throughout this paper we are only considering policies backdated to save age. The choice to backdate is, in essence, the choice to pay an initial upfront payment to have lower per-period payments, which looks a lot like an optional two-part tariff. The extant literature on backdating is limited. Carson (1994) shows that the net present value of backdating can be positive after relatively short periods of time depending on the length of the backdate period, the discount rate, and age. Carson and Ostaszewski (2004) further shows that the actuarial present value of backdating is generally positive.Carson et al. (2012) examine the incentives and welfare economics of life insurance backdating.

Prior work on life insurance backdating has failed to address if and why decreased lapse rates occur. The obvious avenue through which greater policy persistence might occur is via a selection mechanism – those consumers with asymmetric knowledge of their low lapse risk are more likely to pay the initial tariff because the future stream of decreased premiums

[9]There is a somewhat similar vein of work on reclassification risk in health insurance (Herring and Pauly, 2006; Pashchenko and Porapakkarm, 2015; Handel et al., 2015).

[10]This process is regulated by states with typical maximum backdating of one year.

[11]We note also that other forms of two-part tariffs could be offered by insurers (in addition to backdating), such as the option to pay for underwriting /policy issue expenses separately from mortality charges, for example.

is longer and, therefore, more valuable to them. Behavioral economics offers a different mechanism: sunk cost fallacy (Arkes and Blumer, 1985).[12] Sunk cost fallacy would lead to lower lapse rates because consumers may have an aversion to "wasting" the high upfront tariff of backdating. If, for instance, consumers are randomly assigned to the tariff pricing structure, we would not expect those who pay the tariff will exhibit lower lapse rates, unless those consumers are exhibiting sunk cost fallacy if properly controlling for the differences in per-period premiums. Through our empirical structure we are able to separately identify the self-selection and behavioral effects of backdating on lapsing behavior. Our work builds on the growing field evidence of behavior consistent with sunk cost fallacy (Ho et al., 2017).

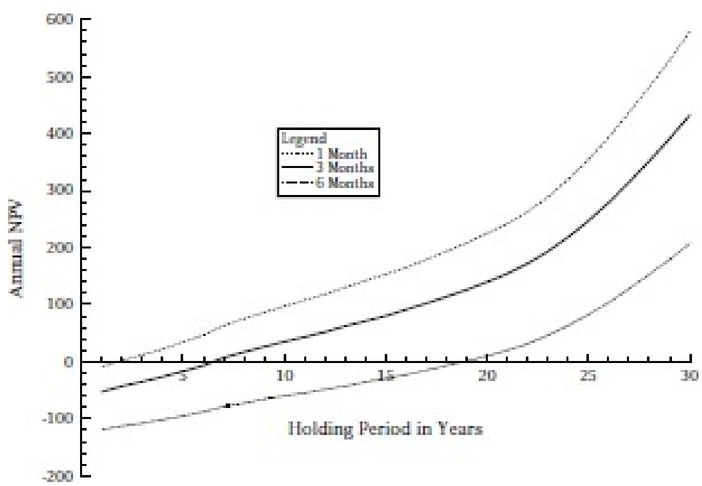

Note: Annual net present value (NPV) for a 35 year old male for a $250,000, 30 year term policy. Source: Carson and Ostaszewski (2004)

Figure 2.1: NPV of Backdated Policy Relative to Normal Policy

Because of the discrete nature of life insurance underwriting, specifically the use of integer ages, the relative distance from a consumer's application date to their birthday will change the size of the initial tariff while having no effect on the size of the premium reduction.

[12] We acknowledge the existence of potential rational explanations for consumers exhibiting what appears to be sunk cost fallacy (McAfee et al., 2010). We are only documenting the existence of the behavior, not necessarily providing evidence of irrational behavior.

Drawn from Carson and Ostaszewski (2004), Figure 1 shows how the actuarial value to a 35-year-old male consumer of backdating a policy evolves over time compared to a traditionally purchased policy for different lengths of backdating. Consumers who apply for life insurance far away from their birthdays must pay a larger upfront cost to get the same reduction in premium. Thus, it takes longer for the premium reduction to outweigh the initial tariff. We exploit this randomness in initial tariff size to separately identify the selection and sunk cost effects of backdating on lapse proclivity. We find each avenue to be a significant contributor to a consumer's decreased propensity to lapse.

Our contribution to the literature is twofold: we are the first to consider life insurance lapsing as a form of adverse selection; we also explore, both theoretically and empirically, the role of optional two-part tariffs as a signaling mechanism using life insurance backdating as our primary example. The findings shed light on why insurers continue to engage in backdating when alternative pricing mechanisms exist – backdating is used by insurers as a screening device to deal with hidden knowledge of lapse intentions that engender a form of adverse selection. As part of our results, we also show evidence of consumer behavior consistent with sunk cost fallacy.

The rest of this article proceeds as follows: the next section describes a simple model of life insurance pricing under heterogeneous lapse behavior with perfect asymmetric information about lapse likelihood and then examines the role of optional two-part tariffs as a signaling device. The third section describes our data and empirical methods, and the fourth section presents our results. The final section offers conclusions.

2.2 Theory:

We begin with a simple expression of how insurers can use optional two-part tariffs to signal a consumer's lapse risk.[13] Let there be two periods. Consider a term insurance contract

[13]In this section we are ignoring any irrational behavior, specifically sunk cost fallacy.

offered by a competitive life insurer lasting both periods with level premiums. Let there be two types of consumers: those who lapse after the first period and those who continue into the second. λ is the share of consumers who lapse on their policies after the first period and $(1-\lambda)$ is the share who do not lapse. Both types of consumers have an equal chance ρ_1 of loss L in the first period and ρ_2 of loss L in the second period. Define $\rho_2 > \rho_1$ to exemplify that death risk increases with age. Consumers are indistinguishable to the insurer who must incur a fixed underwriting cost $F > 0$. The profit Π for a representative insurer, separated by period, offering full coverage at level premium P is

$$\Pi = \Pi_1 + \Pi_2$$

$$\Pi_1 = P - \rho_1 L - F$$

$$\Pi_2 = (1-\lambda)P - (1-\lambda)\rho_2 L$$

Where Π_t is the profit in period t.

With a competitive insurer, we use the zero profit condition and solve for equilibrium price

$$P^* = \frac{F + \rho_1 L + (1-\lambda)\rho_2 L}{(2-\lambda)}.$$

The first derivative of premium with respect to lapse rate is then

$$\frac{\partial P^*}{\partial \lambda} = \frac{F + (\rho_1 - \rho_2)L}{(2-\lambda)^2}$$

The sign of this derivative, which determines whether or not lapsing is costly to the insurer, is determined by the relative size of the underwriting cost F and the difference in expected costs between periods $(\rho_1 - \rho_2)L$. In situations such as shorter term life insurance written for younger consumers, it is likely that $F > (\rho_1 - \rho_2)L$ and thus $\frac{\partial P^*}{\partial l} > 0$ meaning lapsing is costly to insurers. In other situations, such as longer term insurance written to the elderly, $(\rho_1 - \rho_2)L$ may be quite large and lapsing may actually reduce the costs of the insurer. In either case, the optimal price for the two types differs and room for a signaling mechanism exists.

If insurers instead offer an optional two-part tariff, the equilibrium changes. A pooling equilibrium is impossible under the standard Rothschild and Stiglitz (1976) argument so we instead search for a separating equilibrium. This equilibrium is defined by two separate premium offers (P_1, P_2) where $P_1 > P_2$ and an initial tariff T that a consumer must pay to access P_2.[14] This equilibrium space is defined by a solution to the following:

$$P_1 = \rho_1 L + F \tag{2.1}$$

$$2P_2 + T = (\rho_1 + \rho_2)L + F \tag{2.2}$$

$$2P_2 + T \leq 2P_1 \tag{2.3}$$

$$P_2 + T \geq P_1 \tag{2.4}$$

Equations (1) and (2) represent the zero profit conditions; Equations (3) and (4) represent the incentive compatibility constraints. A space of solutions always exists, with an intuitive example being to set the tariff T as slightly higher than the underwriting cost $F + \epsilon$ and P_2 as slightly lower than the expected, first period, loss $\rho_1 L - \epsilon$, where ϵ is small.[15]

2.3 Data and Methods:

A Data:

The data for the analysis come from a medium-sized and geographically diverse mutual life insurer in the U.S. and spans the years 2006 through 2014.[16] The full sample includes

[14] Because our model is only two periods, a two-part tariff for those consumers who lapse after one period is indistinguishable from a single premium.

[15] Unlike Rothschild and Stiglitz (1976), a solution always exists in our model.

[16] We use these years to simplify the issue of a term policy expiring vs. a consumer lapsing.

data for 97,522 term life insurance contracts. In order to avoid the complication of non-monotonicity in lapse cost, we focus exclusively on term insurance between 10-20 years.[17] Because death claims are so rarely paid in these term policies, the potential for consumers holding their policies "too long" appears remote. That is, with reference to our model, it is likely that $F > (\rho_1 - \rho_2)L$, thus making improved persistency strictly a benefit for the insurer.[18] This allows us to view higher likelihood of lapse strictly as an adverse form of selection and adhere more to the predicted selection results of our model. To avoid the complication of classifying deaths in our hazard model, we remove the 727 policies that end in death.[19]

Variable	Full Sample		Backdate		Don't Backdate	
	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
Backdate	0.06	0.29	1	-	0	-
Backdate Opp.	181.26	89.18	60.29	32.03	189.08	85.99
Face Amount($000s)	213.54	288.72	311.13	424.58	207.23	276.49
Issue Age	46.20	11.17	46.46	10.60	42.65	11.17
Annual Premium	578.47	922.11	937.61	1420.76	555.26	875.10
Lapse	0.31	0.46	0.23	0.42	0.32	0.47
Days in Force	1,612.72	1,115.63	1853.40	1087.45	1597.17	1115.66
N	46,199		2,804		43,395	

Note: This table shows summary statistics separated by consumers who did and did not backdate.

Table 2.1: Summary Stats

Life insurance policies are often backdated for reasons other than saving age (such as to align financial documents). To address this, we define a policy as backdated only if the effective date is within one week prior to the consumer's birthday. Our identification strategy relies on randomness in applying for life insurance with regard to a consumer's

[17]This focus also eliminates the yearly renewable term contracts. As these policies are spot contracts and do not have level premiums, a two-part tariff has no meaning. 10-year terms are the smallest length term policies in our data after 1-year renewable term.

[18]In our full sample of 97,522 contracts, we observe 727 deaths; in our selected sample, we observe 494.

[19]Inclusion of policies that end in death and classifying them as non-lapsed policies does not affect our results. Similarly, classifying the policies as lapsing at death does not affect our results.

birthday. However, there is evidence that consumers do alter their behavior during the short time frame around their birthday (Dai et al., 2014). To address this we drop those consumers who apply within one month (before or after) of their birthday.[20] After these modifications, and the removal of missing data, our sample consists of 46,199 active or lapsed policies. Summary statistics for these policies can be found in Table 1. Because they have to pay a lower tariff for the same reduction in premiums, consumers who apply for life insurance closer to their birthdays are more likely to backdate their policies. Additionally, consumers with backdated policies both hold their policies longer and have lower lapse rates. Interestingly, consumers who backdate tend to have larger policies.

B Base Specification:

To first check our proposed optional two-part tariff selection mechanism, we examine if backdated policies exhibit lower lapse rates. Our initial specification for the hazard of lapsing is defined as follows:[21]

$$\begin{aligned}
\lambda(Lapse_{it}|X) &= \lambda_0(t)exp(X_i\beta')u_i \qquad (2.5)\\
X_i &= (\mathbf{Backdate_i}, IssueAge_i, Male_i, FaceAmt_i, AnnPrem_i,\\
&\quad UWClass_i, TermLength_i, IssueYear_i, State_i)\\
u_i &\sim exp(N(0,\sigma^2)).
\end{aligned}$$

[20]If we include these consumers our results are slightly strengthened.

[21]Contrary to most hazard models, our data are not truly panel and thus we do not represent them as such. Each policy represents one row in our data and we derive the hazard specifications from a combination of our time-invariant covariates: length of time policy was/is enforced and a binary indicator for if the policy lapsed. Since our data's time frame for policy issuance is smaller than the minimum term length used in our data (10 years) we do not have to account for issues with a policy expiring without lapse or policy-holder death.

We fit Equation (5) using the partial likelihood method of Cox (1972). $Lapse_{it}$ is equal to 1 in time t if the policy has lapsed at any time equal or prior to t and 0 otherwise. For the independent variables in the model, $Backdate_i$ is 1 if the consumer was found to be saving age and 0 otherwise. $IssueAge_i$ is the age of the insured at time of application. $FaceAmt_i$ is the face value of the life insurance policy. $AnnPrem_i$ is the annual premium for the policy. $UWClass_i$ is a series of dummy variables for the underwriting class of the insured (with four dummy variables for five different categories), with $UWClass1_i$ being the healthiest. $Male_i$ is a binary variable equal to 1 if the insured is male, and 0 otherwise. $Year_i$ is a series of dummy variables for the year of policy issuance. $TermLength_i$ is a series of dummy variables for length of term (10, 15, or 20 years). $State_i$ is a series of dummy variables for state of issuance. Because we are going to use a control function in the next step, we additionally specify the structure of the error term u_i as being log-normally distributed with mean 0 and variance σ^2.

C Control Function Approach:

We wish to examine if information about a consumer's lapse risk is contained within consumer's decision to backdate. However, our above model is not able to separately identify an effect from the decision to backdate or an effect from simply having backdated. In other words, we are unable to tell with this initial specification if differential lapse behavior is due to selection based on prior knowledge of lapse risk or if it is based on the differential pricing structure caused by backdating. Since we control for differences in premiums, any non-selection effect must be due to the initial tariff. This tariff is only paid at the start of the policy, thus consumers who act differently based only on having paid the tariff (not selection effects) must be falling prey to a classic sunk cost fallacy.

We are concerned with the selection effect and sunk cost effect separately, rather than the net effect. Thus our identification problem is classic selection-bias which we address through instrumental variables. We instrument for $\mathbf{Backdate_i}$ using the difference between the approval date of the policy and the consumer's birthday ($\mathbf{BackdateOpp_i}$), i.e., how large of an initial tariff the consumer must pay in order to acquire a lower premium. The exclusion restriction is satisfied assuming that consumers do not consider their birthdays when applying for life insurance. Because consumers are more likely to purchase life insurance close to their birthdays, we drop those who apply to purchase life insurance within a month of their birthdays (both before and after). Figure 2.2 shows the distribution of days required to save age along with a fitted quadratic curve. The slope is statistically significantly negative, however this effect is exceptionally small: a consumer is .02% more likely to apply 30 days after their birthday vs. 335 days.[22]

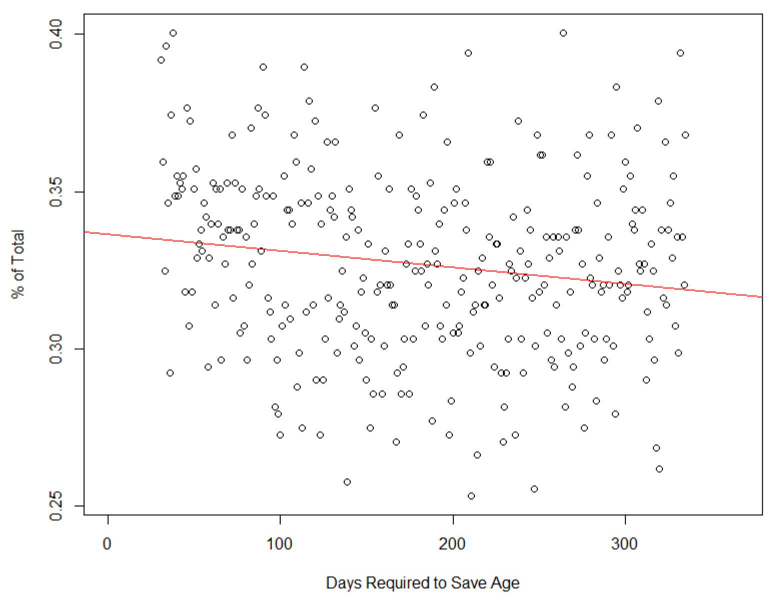

Note: This figure shows the distribution of application days relative to birthdays, i.e. how many days of premiums consumers must pay to backdate. Consumers who purchase within 30 days (before or after) of their birthday are not included.

Figure 2.2: Plot of Days Required to Save Age vs. Percent of Total Sample.

[22]This can be attributed to *Benford's Law* mod *365*.

To instrument here we redefine the error term

$$u_i = exp(\alpha v_i + e_i)$$
$$e_i \sim N(0, \sigma^2)$$
$$Backdate_i = f(Z_i) + v_i$$
$$Z_i = (\mathbf{BackdateOpp_i}, IssueAge_i, Male_i, FaceAmt_i, AnnPrem_i,$$
$$UWClass_i, TermLength_i, IssueYear_i, State_i)$$

Where $f(Z_i)$ is fitted via a Probit specification.

This structure allows us to use a control function approach.[23] Our identified model is

Full Model:

$$\lambda(Lapse_{it}|X) = \lambda_0(t)exp(X_i\beta')u_i \tag{2.6}$$
$$X_i = (\mathbf{Backdate_i}, IssueAge_i, Male_i, FaceAmt_i, AnnPrem_i,$$
$$UWClass_i, TermLength_i, IssueYear_i, State_i)$$
$$u_i = exp(\alpha v_i + e_i) \tag{2.7}$$
$$e_i \sim N(0, \sigma^2) \tag{2.8}$$

First Stage:

$$\mathbf{Backdate_i} = f(Z_i) + v_i \tag{2.9}$$
$$Z_i = (\mathbf{BackdateOpp_i}, IssueAge_i, Male_i, FaceAmt_i, AnnPrem_i,$$
$$UWClass_i, TermLength_i, IssueYear_i, State_i)$$

Second Stage:

$$\lambda(Lapse_{it}|\hat{X}) = \lambda_0(t)exp(\hat{X}_i\beta')e_i \tag{2.10}$$

[23] Also commonly referred to as two-stage residual inclusion. This method originated with Heckman and Hotz (1989). For an excellent review with a health focus, the reader is referred to Terza et al. (2008).

$$\hat{X}_i = (\hat{\mathbf{v}}_\mathbf{i}, \mathbf{Backdate_i}, IssueAge_i, Male_i, FaceAmt_i, AnnPrem_i,$$
$$UWClass_i, TermLength_i, IssueYear_i, State_i).$$

Here the second stage (10) is the same as (5) with the inclusion of the residuals from the first stage (9).[24] Following Basu and Coe (2015) we use Anscombe residuals (Anscombe and Tukey, 1963) to account for the relative rarity of backdating.[25] Equation (10) allows us to separately identify the selection effect (coefficient on $\hat{\mathbf{v}}_\mathbf{i}$) from the sunk cost effect (coefficient on $\mathbf{Backdate_i}$).

The interpretation of the coefficient on $\hat{\mathbf{v}}_\mathbf{i}$ is similar in intuition to the coefficient on the inverse Mills ratio in the classic Heckman selection model. That is, the significance of the coefficient shows whether some form of selection is occurring. The interpretation of the coefficient on $\mathbf{Backdate_i}$ as only the effect of paying the initial tariff on future decisions depends vitally on the implicit assumption that premium changes are being effectively controlled for via the additively linear term. There is no reason to believe this is true. To account for arbitrary non-linearity in the control variables while still preserving the ability to instrument we turn to a Lasso technique.

D Lasso:

The Lasso (least absolute shrinkage and selection operator) is a model selection technique originally developed by Tibshirani (1996) as an improvement on step-wise regression and adapted to Cox hazard models by Tibshirani et al. (1997). The technique is currently popular

[24]It is important to note that $-1 < v_i < 1$ and thus is v_i is not normally distributed. However, $E[v_i] = 0$ since the predictions of a Probit specification are unbiased. In Equation (10), v_i is transformed into approximately normal $\hat{\mathbf{v}}_\mathbf{i}$ via an Anscombe transformation. See Basu and Coe (2015) for further discussion on 2SRI with a binary first-stage as well as Anscombe residuals.

[25]Anscombe residuals are a transformation of the standard residuals into an approximately standard normal distribution. Standard errors are bootstrapped.

in the machine learning literature and was introduced to the econometrics literature by Belloni et al. (2012).[26] The Lasso is in the class of l_1-penalized methods of model selection.[27]

The Lasso is beneficial here in two specific ways. First, the model selection allows us to account for (nearly) arbitrary non-linearity in our control variables via polynomial approximation. Rather than including only linear representations of our control variables, we allow linear, squared, and cubic terms as well as all possible two-variable (inclusive of squared and cubed variables) interaction terms and then allow the Lasso to select the important terms. The second benefit of this method of approximating non-linearity in control variables is the preservation of the linear nature of the treatment variables (here $\mathbf{Backdate_i}$ and the selection effect $\hat{\mathbf{v}}_\mathbf{i}$) allowing for control function instrumentation. Formally, our final model is

Final Model:

$$\lambda(Lapse_{it}|X) = \lambda_0(t)exp(X_i\beta')u_i \tag{2.11}$$

$$X_i = (\mathbf{Backdate_i}, IssueAge_i, Male_i, FaceAmt_i, AnnPrem_i,$$
$$UWClass_i, TermLength_i, IssueYear_i, State_i)$$

$$u_i = exp(\alpha v_i + e_i) \tag{2.12}$$

$$e_i \sim N(0, \sigma^2) \tag{2.13}$$

First Stage:

$$\mathbf{Backdate_i} = f(Z_i) + v_i \tag{2.14}$$

$$Z_i = (\mathbf{BackdateOpp_i}, IssueAge_i, Male_i, FaceAmt_i, AnnPrem_i,$$
$$UWClass_i, TermLength_i, IssueYear_i, State_i)$$

Second Stage:

[26] See also Belloni et al. (2014b), Belloni et al. (2014a), Belloni et al. (2016), and Chernozhukov et al. (2015).

[27] That is, they generally take the form: $\beta_{Lasso} = \underset{\beta}{\text{argmin}} \left\{ \sum (y - X\beta)^2 \right\}$ subject to $\sum |\beta| \leq l_1$.

$$\lambda(Lapse_{it}|\hat{L}) = \lambda_0(t) exp(\hat{L}_i \beta'_{Lasso}) e_i$$
$$\text{where } \beta'_{Lasso} = \underset{\beta}{\text{argmax}} \{l_{Cox}(\beta L)\} \text{ subject to } |\beta| \leq l_1 \tag{2.15}$$
$$\hat{L}_i = S(\hat{\mathbf{v}}_\mathbf{i}, \mathbf{Backdate_i}; \hat{\mathbf{X}}_\mathbf{i}, \hat{\mathbf{X}}_\mathbf{i}^\mathbf{2}, \hat{\mathbf{X}}_\mathbf{i}^\mathbf{3}, \hat{\mathbf{X}}_\mathbf{i} : \hat{\mathbf{X}}_\mathbf{i})$$

Third Stage:

$$\lambda(Lapse_{it}|\hat{L}^p) = \lambda_0(t) exp(\hat{L}_i^p \beta') e_i \tag{2.16}$$
$$\hat{L}_i^p = (\hat{\mathbf{v}}_\mathbf{i}, \mathbf{Backdate_i}; \hat{L}_i \text{ such that } \beta_{Lasso,i} \neq 0)..$$

Where the first stage is the same as above. $\hat{L}_i = S(\hat{\mathbf{v}}_\mathbf{i}, \mathbf{Backdate_i}; \hat{\mathbf{X}}_\mathbf{i}, \hat{\mathbf{X}}_\mathbf{i}^\mathbf{2}, \hat{\mathbf{X}}_\mathbf{i}^\mathbf{3}, \hat{\mathbf{X}}_\mathbf{i} : \hat{\mathbf{X}}_\mathbf{i})$ is the standardized collection of all control variables, squared terms, cubed terms, and two-variable interactions. Equation (15) describes the Lasso procedure where $l_{Cox}(\beta L)$ is the likelihood function for the Cox proportional hazard model.[28] Following Belloni et al. (2016) in Equation (16) we estimate the unpenalized Cox proportional hazard using the two treatment variables and all of the control variables whose coefficients were non-zero in the second stage. The included variables and combinations of variables in $\hat{L}_i^p$ can be interpreted as the optimal polynomial form of the control variables that can be represented in a limited (via the choice of l_1) number of terms.

[28] In Equation (15), the coefficients on $\hat{\mathbf{v}}_\mathbf{i}$ and $\mathbf{Backdate_i}$ are not penalized. l_1 is determined via cross-validation techniques (Goeman, 2010). Additionally, the linear versions of all variables in $\hat{\mathbf{X}}_i$ are included in $\hat{L}_i^P$ regardless of their coefficient in the Lasso stage.

2.4 Results:

	Dependent variable:
	Hazard(Lapse)
Variable of Interest:	
Backdate	-0.244^{***}
	(0.041)
Controls:	
Issue Age	-0.029^{***}
	(0.001)
Male	0.045^{***}
	(0.017)
Face Amount($000s)	-0.001^{***}
	(0.000)
Annual Premium	0.000^{***}
	(0.000)
LengthofTerm15	-0.498^{***}
	(0.033)
LengthofTerm20	-0.286^{***}
	(0.018)
UWRank2	0.174^{***}
	(0.044)
UWRank3	0.504^{***}
	(0.076)
UWRank4	1.194^{***}
	(0.044)
UWRank5	0.717^{***}
	(0.040)
State Effects?	Yes
Year Effects?	Yes
Observations	46,199
R^2	0.101
Log Likelihood	-146,033.940
Wald Test	$4,834.540^{***}$
LR Test	$4,913.589^{***}$
Score (Logrank) Test	$5,087.000^{***}$

Note: *p<0.1; **p<0.05; ***p<0.01. This table shows Equation (5) fitted via the Cox proportional hazards technique. Std. errors are in parentheses.

Table 2.2: Cox Proportional Hazard Model

The results from our initial specification (Equation (5)) are presented in Table 2. The initial regression results are congruent with with our theoretical predictions. Policy-owners who choose to backdate their life insurance contracts (effectively paying a two-part tariff) signal their lower likelihood for lapsing by their willingness to pay for time that already has elapsed that only results in net saving if the policy is held for a relatively long period of time. However, this coefficient does not fully identify the selection effect that we seek. Potentially, consumers who backdate may have no additional knowledge of their propensity for lapsing and are instead exhibiting sunk cost bias.

	Dependent variable:
	Backdate
Variable of Interest:	
Backdate Opp.	−0.019***
	(0.000)
Controls:	
Issue Age	0.020***
	(0.001)
Male	0.150***
	(0.026)
Face Amount($000s)	0.001***
	(0.000)
Annual Premium	0.000***
	(0.000)
LengthofTerm15	0.151***
	(0.041)
LengthofTerm20	−0.008
	(0.029)
UWRank2	−0.024
	(0.046)
UWRank3	−0.152
	(0.103)
UWRank4	−0.228***
	(0.059)
UWRank5	−0.109**
	(0.043)
Constant	−0.728***
	(0.099)
State Effects?	Yes
Year Effects?	Yes
Observations	46,199
Log Likelihood	-6,231.460
Akaike Inf. Crit.	12,530.920

Note: *p<0.1; **p<0.05; ***p<0.01. This table shows fitted values for the first stage in our control function procedure (Equation (9)). Std. errors are in parentheses.

Table 2.3: First Stage Regression

To identify these separate effects, we exploit inherent randomness in the time of year, relative to the consumer's birthday, that the consumer applies for life insurance. This allows us to perform a pseudo-random experiment exploiting variation in the initial tariff consumers have to pay while holding constant the reduction in future premiums. The results from the first stage in our control function approach (Equation (9)) are presented in Table 3. Our instrument is strong and loads in the predicted manner – people who have to pay a higher initial tariff, *ceteris paribus*, are less likely to choose to backdate.

	Dependent variable:	
	Hazard(Lapse)	
	Initial	Instrumented
	(1)	(2)
Variables of Interest:		
Backdate	−0.244***	−0.105*
	(0.041)	(0.058)
Stage 1 Residuals		−0.099***
		(0.028)
Controls:		
Issue Age	−0.029***	−0.030***
	(0.001)	(0.001)
Male	0.045***	0.043**
	(0.017)	(0.018)
Face Amount($000s)	−0.001***	−0.001***
	(0.000)	(0.000)
Annual Premium	0.000***	0.000***
	(0.000)	(0.000)
LengthofTerm15	−0.498***	−0.501***
	(0.033)	(0.033)
LengthofTerm20	−0.286***	−0.286***
	(0.018)	(0.018)
UWRank2	0.174***	0.174***
	(0.044)	(0.042)
UWRank3	0.504***	0.509***
	(0.076)	(0.075)
UWRank4	1.194***	1.197***
	(0.044)	(0.043)
UWRank5	0.717***	0.718***
	(0.040)	(0.038)
State Effects?	Yes	Yes
Year Effects?	Yes	Yes
Observations	46,199	46,199
R^2	0.101	0.101
Log Likelihood	-146,033.940	-146,026.961
Wald Test	4,834.540***	4,848.330***
LR Test	4,913.589***	4,927.546***
Score (Logrank) Test	5,087.000***	5,098.664***

Note: *p<0.1; **p<0.05; ***p<0.01. This table presents again for comparison our initial model from Table 2 in the first column and shows fitted values for the second stage in our control function approach in the second column. Bootstrapped std. errors are in parentheses.

Table 2.4: Two-Stage Residual Inclusion Estimation

We then re-estimate the hazard specification from Table 2 (Equation (5)), this time including the transformed residuals from the first stage. The results from our instrumented model (Equation (10)) are presented in the second column of Table 4. Our original estimated effect was indeed a combination of both selection and sunk cost bias. The selection effect (-0.099) is statistically significant, and it also appears that a significant portion (-0.105) of the reduction in lapse likelihood is due to sunk cost bias.

The third column of table 5 shows the results of the Lasso procedure. In addition to the 12 independent variables shown in the table, the model includes state effects, year effects, and 82 (out of over 800 potential) other forms of the control variables (either squared, cubed, or interaction) that the Lasso procedure selected as important.[29] The inclusion of these variables does not change the significance of our results, though the point estimates are slightly diminished.[30] The selection effect does not have an interpretation beyond directional comparative statics, however the sunk cost effect does. Exponentiating the coefficient on $\mathbf{Backdate_i}$ results in a marginal relative hazard rate of 92.5%. That is, the sunk cost effect of backdating reduces the per-period hazard rate of lapsing by 7.5%.

[29]The extra control variables are suppressed in the table for space, a full table reporting the coefficients on all Lasso control variables can be found on the author's website.

[30]Interestingly, the significance of the coefficients for the linear terms of $Male$ and $FaceAmount$ goes away. This is likely due to the inclusion, via the Lasso, of many significant interaction terms containing those variables.

	Dependent variable:		
	Hazard(Lapse)		
	Initial	Instrumented	Lasso Controlled
	(1)	(2)	(3)
Variables of Interest:			
Backdate	-0.244^{***}	-0.105^{*}	-0.078^{*}
	(0.041)	(0.058)	(0.045)
Stage 1 Residuals		-0.099^{***}	-0.072^{***}
		(0.028)	(0.020)
Controls:			
Issue Age	-0.029^{***}	-0.030^{***}	-0.020^{***}
	(0.001)	(0.001)	(0.002)
Male	0.045^{***}	0.043^{**}	-0.015
	(0.017)	(0.018)	(0.034)
Face Amount($000s)	-0.001^{***}	-0.001^{***}	-0.000
	(0.000)	(0.000)	(0.000)
Annual Premium	0.000^{***}	0.000^{***}	-0.000^{***}
	(0.000)	(0.000)	(0.000)
Length of Term	-0.498^{***}	-0.501^{***}	-0.340^{***}
	(0.033)	(0.033)	(0.099)
UW Rank	-0.286^{***}	-0.286^{***}	-0.192^{*}
	(0.018)	(0.018)	(0.106)
UWRank2	0.174^{***}	0.174^{***}	0.352^{***}
	(0.044)	(0.042)	(0.074)
UWRank3	0.504^{***}	0.509^{***}	0.848^{***}
	(0.076)	(0.075)	(0.181)
UWRank4	1.194^{***}	1.197^{***}	1.087^{***}
	(0.044)	(0.043)	(0.090)
UWRank5	0.717^{***}	0.718^{***}	0.630^{***}
	(0.040)	(0.038)	(0.126)
State Effects?	Yes	Yes	Yes
Year Effects?	Yes	Yes	Yes
Lasso Controls?	No	No	Yes (82)
Observations	46,199	46,199	46,199
R^2	0.101	0.101	0.122
Max. Possible R^2	0.998	0.998	0.998
Log Likelihood	-146,033.940	-146,026.961	-145,477.643
Wald Test	$4,834.540^{***}$	$4,848.330^{***}$	$5,437.540^{***}$
LR Test	$4,913.589^{***}$	$4,927.546^{***}$	$6,026.184^{***}$
Score (Logrank) Test	$5,087.000^{***}$	$5,098.664^{***}$	$6,619.428^{***}$

Note: *p<0.1; **p<0.05; ***p<0.01. This table presents again for comparison both our initial model from Table 2 in the first column and the fitted coefficients for the second stage in our control function approach in the second column. The third column represents the results of the Lasso procedure. A full representation and list of the Lasso coefficients can be found on the author's website. Bootstrapped std. errors are in parentheses.

Table 2.5: Lasso Estimation

2.5 Conclusions:

Our results indicate that asymmetric information about lapse risk can be reduced through a firm offering a menu of two-part tariff contracts and allowing consumers to self-select into the contract designed for their lapse type. These optional two-part tariffs serve as a signaling device for insurers on consumers' likelihood of lapsing. Consumers who choose to pay the two-part tariff (e.g., backdating their life insurance contracts) signal their lower likelihood for lapsing by their willingness to pay for time that already has elapsed to have lower premiums that only results in net savings if the policy is held for a long time. If consumers terminate the policy early, they do not reap the benefit of the lower premium level. Such a willingness translates into a significantly lower hazard of lapsing, thus aligning the interests of the consumer with the insurer. Our research provides key insight into why insurers do not use continuous (with regards to age) pricing despite the computational ease of doing so – the value of the signal provided by offering the optional two-part tariff outweighs any actuarial downside caused by discreteness in years.

We additionally find strong evidence that life insurance consumers exhibit behavior consistent with sunk cost fallacy in their lapsing behavior, even when controlling for arbitrary non-linearity in premium effects on lapse proclivity. This interesting finding implies a larger degree of reverse causality with lapse rates and premium structures (including initial tariffs) than is currently being discussed in the literature. Development of a simultaneous model of lapse rates and premium structures would be an excellent direction for further research.

Though our investigation is limited to backdating in life insurance, our theoretical model and the associated signaling power of an optional two-part tariff generalizes to any form of insurance with long-term relationships and large upfront underwriting costs, especially some commercial lines.[31] However, only the selection effects would generalize as firms purchasing

[31] Many forms of property & casualty insurance fall under this veil. For instance: D&O insurance, catastrophe insurance, etc.

insurance likely do not exhibit the same irrational behavioral tendencies that individual consumers do.

Chapter 3

Getting Crowded: Private Market Effects of Medicaid Expansion Refusal[1]

3.1 Introduction:

On March 23, 2010, President Obama signed into effect one of the largest overhauls of the U.S. healthcare system: The Patient Protection and Affordable Care Act (ACA). In addition to numerous other clauses, the ACA mandated for the expansion of Medicaid eligibility requirements to allow many low-income Americans access to the program. This expansion raised the income ceiling of Medicaid eligibility from 100% to 138% of the federal poverty line and allowed childless adults with income below this line access to Medicaid; potentially extending coverage to more than 20 million Americans (Holahan, 2012). To ensure compliance, the ACA permitted the Secretary of the Department of Health and Human Services to withdraw existing federal Medicaid funds to states failing to adopt the new federal eligibility requirements. In response to this drastic overhaul of the U.S. healthcare system, 14 states brought suit questioning the constitutionality of the act. Shortly after, an additional 13 states either joined existing litigation or filed separately. These cases merged together into *National Federation of Independent Business v. Sebelius* which was heard before the U.S. Supreme Court in March, 2012. While the majority of the ACA held up to judicial scrutiny, the Supreme Court ruled that mandatory expansion of Medicaid was unconstitutional. This ruling allowed states the option of "opting out" of the Medicaid expansion.

[1]This chapter is joint work with Meghan I. Esson and Joshua D. Frederick.

The debate over the Medicaid expansion focused on the direct costs and benefits of such an expansion; lawmakers weighed the direct benefits of increasing insurance coverage versus the accounting cost of doing so (Harrington, 2010b). We examine another avenue in which the Medicaid expansion may impact the U.S. healthcare market: the potential effects that a state's expansion of Medicaid may have on premiums in the private market for health insurance. Using prices in the Health Insurance Exchanges as a proxy for the private health insurance market, we find that expanding Medicaid reduces average monthly premiums by \$32.4; a decrease of 11.86%. This finding is robust across different identification strategies.

Two other policies enacted by the ACA, and upheld by the Supreme Court, are especially important to our story: limits on medical underwriting and the individual mandate. Historically, health insurers have been allowed to screen plan enrollees to identify their potential risk factors, a process known as medical underwriting. The ACA restricts this practice. Now, insurers may only vary premiums based on age, tobacco usage, self vs. family coverage, and geographical area. Importantly, health insurers are not allowed to underwrite based on the income of the individual. Because Medicaid eligibility is largely determined by income, the limits on medical underwriting prevent insurers from identifying consumers who would have been covered by the Medicaid expansion. Thus, states' refusal will impact the price of health insurance for all individuals given two assumptions are met. First, there must be a sufficient amount of potentially Medicaid eligible individuals in non-expansion states that buy private insurance instead of going uninsured. Second, those individuals who would have been eligible for Medicaid and purchase health insurance on the private market must have different expected costs than the general, Medicaid ineligible populace.

The first assumption describes health insurance "crowd out." Coined by Cutler and Gruber (1996), crowding out is the phenomenon where individuals forgo private insurance in favor of public insurance. This assumption is supported by the individual mandate portion of the ACA.[2] The mandate requires that individuals must have "minimum essential coverage" or pay a tax penalty.[3]

The extant literature on the existence of crowding out effects of public insurance is conflicting and the overall effect on the private insurance market, both premiums and quality, remains an open question. Cutler and Gruber (1996) argue that as public insurance eligibility increases, individuals may drop their private coverage. In testing this hypothesis, Cutler and Gruber (1996) use the Medicaid expansion for children and pregnant women from 1987-1992 and find evidence that a significant portion of the increase in Medicaid coverage was successively followed by a reduction in private insurance coverage. However, Selden et al. (1998) find no evidence of crowding out. Selden et al. (1998) find that even with free eligibility, the rates of uninsurance amongst children remained high following previous Medicaid expansion periods. As the children themselves are not able to decide on their insurance, it is obvious that their parents are not selecting into the coverage. Selden et al. (1998) further argue that this lack of coverage could be due chiefly to incomplete information.

Further illustrating the split on the existence of a crowding out effect, Ham and Shore-Sheppard (2005) criticize Cutler and Gruber (1996) for failing to consider insurance coverage as a family decision. When accounting for this, Ham and Shore-Sheppard (2005) find that, though overall take up increases the crowding out effect was not significant. This indicates that the new insureds were not previously private market participants.[4]

The second assumption that must be met for Medicaid expansion to impact the price of health insurance for all individuals is there must be a discrepancy between the expected costs of the Medicaid eligible and the rest of the consumers who purchase private insurance. This

[2]Traditionally crowd out is viewed as people who have private insurance suddenly becoming eligible for public insurance. Our situation is a little different because our variation is geographical rather than temporal.

[3]Minimum essential coverage includes all government and job based insurance and most private insurance. Short term plans, fixed benefit plans and vision/dental only plans alone do not qualify for the minimum requirement.

[4]For a detailed literature review of the crowding out literature, please refer to Gruber and Simon (2008).

discrepancy in expected costs occurs through two main channels: The first channel is the gradient of health and wealth, which states that income and health are, generally, negatively correlated (Deaton and Paxson, 1998; Deaton, 2002).[5] The second channel is that low income consumers tend to be more price sensitive than individuals with more disposable income, increasing the potential for adverse selection.

The link between low income and poor health is well documented in the literature and stems from a number of factors. Poor people are more likely to have insufficient nutrition and lower quality diets which can lead to health problems, a relationship that often begins in childhood (Case et al., 2002). Malnutrition weakens the immune system, increasing the likelihood of health problems. Additionally, many simply cannot afford the services (both preventative and ex post) or medicine they require (Deaton and Paxson, 2001).

In addition to the lack of discretionary funds and malnutrition, lower income individuals also tend to live in rural areas, and are therefore far removed from standard health amenities. This distance reduces access to healthcare and health services, and is yet another contributing factor to the negative correlation between low income and poor health. The link between income and education is another avenue; the poor are typically less educated, and subsequently less likely to be aware of health services offered. The compounding effect of each of these relationships leads to a strong negative correlation between income and health.

The other primary avenue is the increased price sensitivity of the poor (Cutler and Reber, 1998). Even if subsidized, the difference in premiums between stringy and generous plans represents a much larger portion of these consumer's income.[6] If a poor individual is willing

[5]We do not require, or assume, a direction of causality. That is, poor health could lead to low income, or low income could lead to poor health outcomes (Bloom and Canning, 2000). We do not take a stance on the direction of the causality of health and income.

[6]Since the decision of the Supreme Court was unexpected at the time the law was written, there exists a subsidy gap in states that opt out of Medicaid. People in those states ineligible for Medicaid with incomes below 100% of the poverty line receive no subsidies with large subsidies beginning at 100% of the poverty line.

to pay the higher premium of a high-coverage plan, they are likely doing this because of asymmetric information about their health – not because they value the risk reduction.

These two channels imply that as more low income, and potentially less healthy, individuals enter the health insurance market, the risk, and therefore cost, of the overall health insurance pool increases. However, due to new underwriting restrictions, health insurance companies cannot charge higher prices based on income and thus must raise prices for everyone.[7] Crowding out is usually implied as a problem for those partaking in the private market, but this case is different. Here the population that is being potentially crowded out by the Medicaid expansion has higher expected costs than the general population. Since insurance companies can not underwrite based on potential Medicaid eligibility, taking these people out of the private market lowers prices for everyone. We thus refer to a state's refusal to partake in the Medicaid expansion as "crowding in" the private market.

Given the ACA's recency and the lag in data availability, there exist few empirical studies on the effects of the law. The projected impact of the ACA is discussed in detail by Harrington (2010a). Sommers et al. (2013) shows the dependent coverage clause, which allows young adults coverage under the private plans of their parents until age 26, substantially increased the coverage rate of young adults aged 19-25 . Depew and Bailey (2015) investigate the impact of this mandate and find it led to an increase in premiums for plans of individuals with children. Pauly et al. (2015) investigates welfare effects on the non-poor who were uninsured prior to the law finding that they are typically worse off due to the mandate. Hilliard et al. (2013) examine the market impact of the supreme court's decision finding evidence of a negative stock market response to health insurers.

Our paper is closely related to Sen and Deleire (2016) and can be viewed as further exploration, expansion, and confirmation of their work. Since individuals are mandated to

[7] Employer-based insurance has never been able to do this due to anti-discrimination laws. The plans offered through the exchanges are allowed more freedom to underwrite (such as age discrimination) but are not allowed to individually underwrite based on income.

have health insurance, in the states where the Medicaid expansion was refused there exists a coverage gap that the private market will, at least partially, pick up. This gap, and the portion picked up, are not random and systematically exhibit higher expected costs. This raises prices for everyone taking part in the health insurance exchanges. In order to fully assess the welfare implications of Medicaid refusal, the "crowding in" effects on the private markets must be considered.

This paper proceeds as follows: Section 3.2 provides institutional details about the ACA and Medicaid; Section 3.3 describes the data; Section 3.4 provides the methods; Section 3.5 estimates the differential impact of refusing the Medicaid expansion on premiums in the health insurance exchanges; Section 3.6 offers conclusions.

3.2 Institutional Details:

A Affordable Care Act:

In addition to the Medicaid expansion (and numerous other more minor reforms), the ACA restricts what health insurance firms are able to utilize for medical underwriting, restricts the variations in premiums over time, includes an individual mandate, which requires individuals to have insurance coverage or pay a fine, coverage is mandatory, and establishes online marketplaces in every state.

We use these marketplaces as a proxy for the private market. The marketplaces play a key role in the goal of expanding coverage. Federal subsidies are only available to those who purchase policies from the Health Insurance Exchange (HIX). States had three separate paths to the development of their HIXs: design and manage their own, let the federal government design and manage it, or some hybrid approach.[8] Regardless of the path chosen, all of the HIXs run in the same general manner; they are designed such that there are tiers of plans

[8] 16 states and DC selected the first option, 27 selected the second, and 7 pursued the third.

from which individuals can select and purchase insurance. In the HIXs there are five possible tiers of plans. Catastrophic, high-deductible, plans are the lowest rung, with the other four tiers, identified by different metals, based on the expected share of healthcare spending the plan covers: bronze (60%), silver (70%), gold (80%), and platinum (90%).

Outside of the additional standardization of covering essential health benefits and a maximum out-of-pocket expenditure ($6,350 for individuals, $12,700 for families), insurers are able to tailor their policies in nearly any way, so long as the insurers are within 2% of the targeted actuarial-value and the mandated benefits are covered. Premiums are required to be community-rated and can vary only across ratings areas, age, tobacco use, and family composition. Aside from these restrictions, insurers are free to set their own initial premiums, but there exists regulation on increasing rates; any rise in premiums greater than 10% must be subject to approval by a state board. There also exists a minimum, plan-level, medical loss ratio; all medical loss ratios must exceed 80% for individual and small-group markets and 85% for large-groups. If the medical loss ratios are not met, the insurers must issue refunds to their customers.[9]

B Medicaid:

We are primarily concerned with the Medicaid expansion provision of the ACA. The Medicaid program began in the late 1960s and was tied to state welfare programs. Medicaid was slowly separated into its own program and by the late 1990s covered most low-income (<100% of the poverty line) children. The State Children's Health Insurance Program (SCHIP), created in 1997, expanded coverage to cover uninsured children in families with modest incomes that were too high to qualify for Medicaid. Prior to the passage of the ACA, there was no federal regulation mandating coverage for low-income, non-disabled adults with no dependents. Though there were a number of provisions to the Medicaid expansion, we focus on

[9] See Abraham and Karaca-Mandic (2011) for an in-depth discussion of focusing regulation on the medical loss ratio.

the raising of the income ceiling for eligibility to 138% of the poverty line and the inclusion of non-disabled, childless adults. It is projected that the Medicaid expansion provision of the ACA would increase overall Medicaid enrollment by more than 20 million.[10] If fully implemented, the Medicaid expansion could reduce the number of uninsured by 48% compared to non implementation (Holahan, 2012). The increase in Medicaid coverage is not free (Harrington, 2010a,b), and many of the states who opposed the expansion did so because of uncertainty in where the burden of the increased cost would eventually fall.

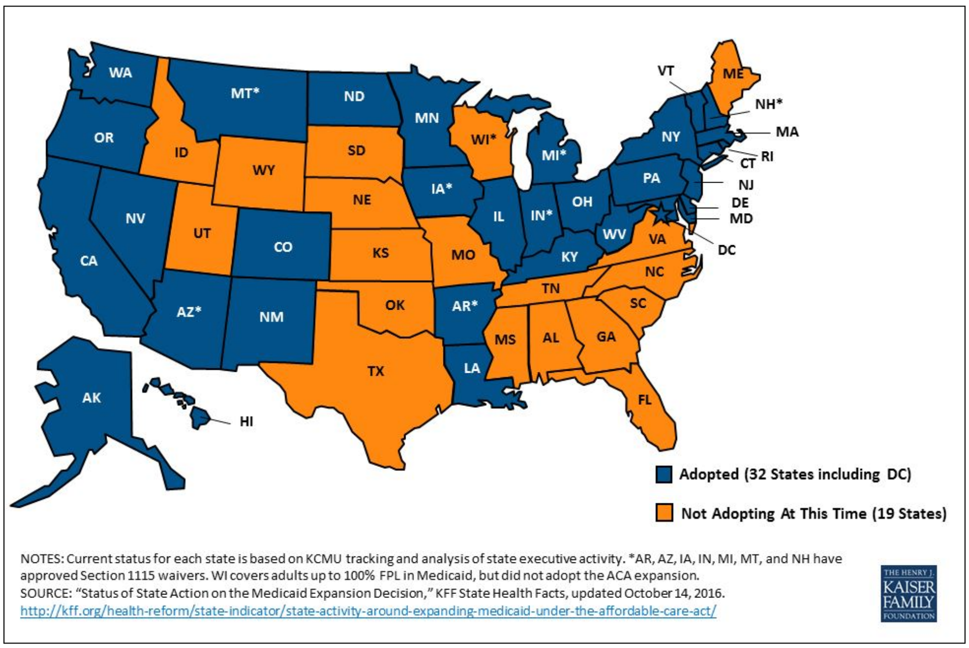

Figure 3.1: Current Status of Medicaid Expansion

As of 2016, 19 states have actively opted out of the Medicaid expansion of the ACA, citing budgetary constraints or costs as their main point of opposition.[11] Many of the opt out states are Southern and Midwestern states, with only a few exceptions. All of the opt out states have Republican governors, speaking to the highly politicized nature of the Medicaid expansion and/or the ACA in general. Figure 3.1 details the current status of each state's

[10]Nearly 21.3 million individuals could gain Medicaid coverage by 2022, of which 14.3 million would only be eligible if all states opted into the expansion.

[11]A few of states are still debating on whether or not to expand Medicaid. For our analysis they are still considered "opt out" states.

decisions to opt in or out of the ACA's Medicaid expansion provision, Figures 3.2, 3.3, and 3.4 show the status at January 1 of 2014, 2015, and 2016 respectively.[12]

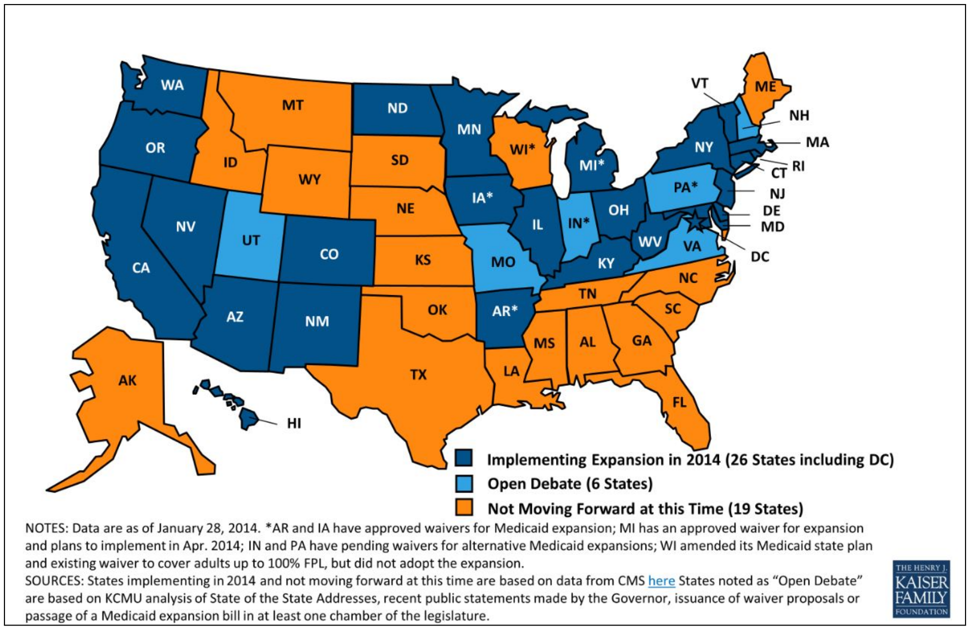

Figure 3.2: January 2014 Medicaid Expansion Status

[12]A few states received waivers for the proposed expansion by providing their own plan for how to increase coverage. Arkansas is an important one for our analysis, so we will further describe their policy here. Rather than including those newly eligible in Medicaid, Arkansas uses the federal funds to subsidize purchases on the Arkansas Health Insurance Exchange. For our analysis, we treat Arkansas as an opt in state. While the lower general health population is partaking in the private market, the subsidy gap does not exist and thus we would expect greatly reduced adverse selection issues. Indiana adopted a similar plan in 2015.

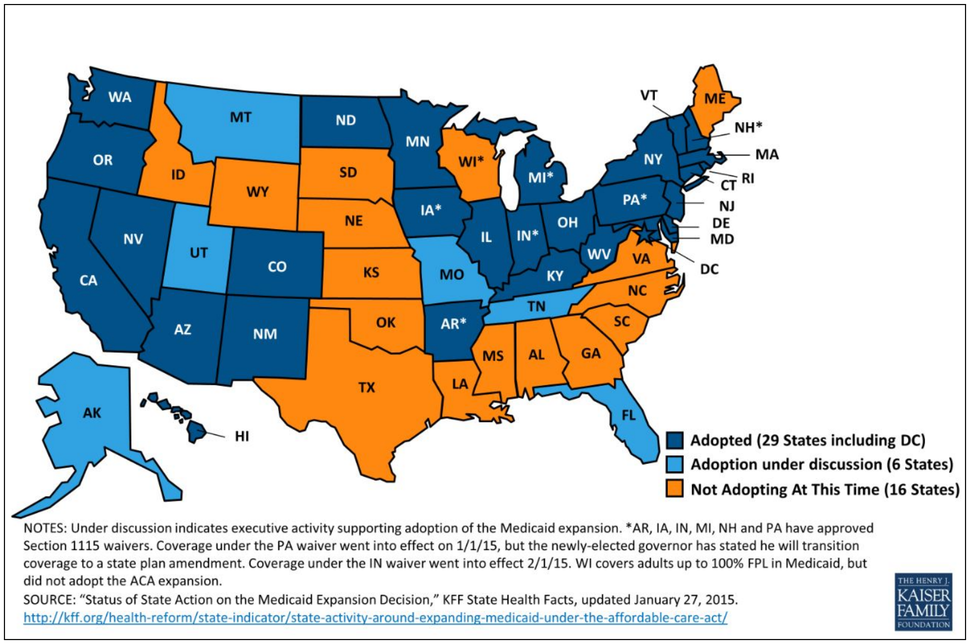

Figure 3.3: January 2015 Medicaid Expansion Status

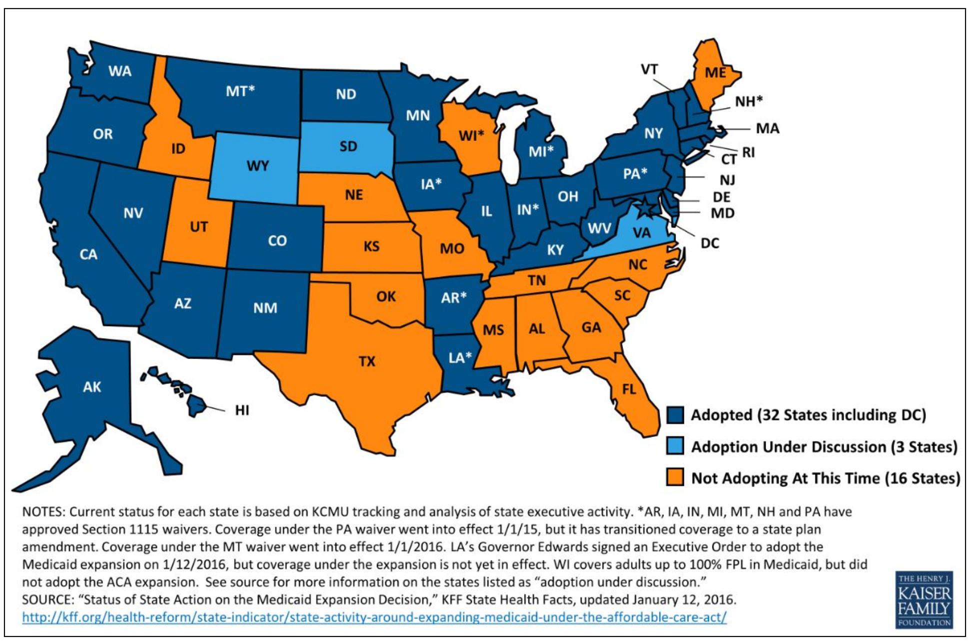

Figure 3.4: January 2016 Medicaid Expansion Status

3.3 Data:

The data for this paper are drawn from a number of sources. The policy details come from the AIS Health Insurance Exchange Database 2014-2016 archive. This is a comprehensive list of all the policies offered in all of the states in 2014-2016. The covariates include the metal level of the policy, and the policy type (HMO, PPO, POS, EPO). The market (typically county-level) covariates are pulled from the county-level data from the Dartmouth Health Atlas and include average Medicare spending per enrollee (to proxy for health care price variation), population, the proportion of people in poverty, and the level of "urban-ness" of the county.[13]

The policy details from the AIS database include metal level, plan type, and price by county, state and the issuing firm. Plan quality is partitioned out into the following categories: Catastrophic, Bronze, Silver, Gold, and Platinum.[14] Plan price means range from \$225.76 to \$371.92 for Bronze and Platinum plans, respectively. Our sample represents EP, HMO, POS and PPO plan types, with the majority being HMO and PPO. Unlike most current studies on the exchanges, we have available information from all of the states as well as for multiple years.

3.4 Methods:

Our initial empirical model is:

$$Prem_{pmst} = \beta Expand_{mst} + \alpha P_{pt} + \gamma M_{mt} + F + S + Year + \epsilon_{pmst}$$

[13]Due to lags in data release, we only use the most recent year (2013) of the DHA data. We apply these single year county-level statistics to all of the years in our data to avoid any interpretation of changes that occured 3 years ago influencing prices currently."Urban-ness" is on a scale of 1-9 with 1 being the most urban.

[14]Due to the low availability and differing regulations on catastrophic plans we drop them from our analysis.

Where, for the Equation, $Prem_{pmt}$ is the premium for plan p, in market m, in year t; $Expand_{mt}$ is binary representing one if the market is in a state that expanded Medicaid before year t and zero else; P_p and M_m are vectors of the plan and market covariates (respectively); F and S are firm and state effects (respectively); ϵ_{pm} is the mean zero, exogenous error term.[15] A potential problem is the existence of unobservable market characteristics correlating with a state's decision whether or not to expand Medicaid. We pursue two different avenues to address this endogeneity concern: a geographic disconinuity avenue and a difference in differences analysis.

Our geographic discontinuity strategy consists of limiting our sample to counties that border across state lines where the state pairs differ in their decision to expand Medicaid. This approach has direct precedent with Sen and Deleire (2016) as well as some indirect precedent with Gowrisankaran and Krainer (2011) examining ATM surcharges using differing laws in Minnesota and Iowa and with Dube et al. (2010) using it to examine minimum wage effects on job growth. Our strategy differs slightly to the one used by Sen and Deleire (2016), however our results are generally equivalent. A key difference is that we use data on multiple years and more states, as well as we do not include states which define rating areas as MSA + 1 (such as Texas) since the geographic identification is clouded for the rating area that covers the rural areas of the entire state. We also define our unit of analysis as individual plans rather than using, as much of the literature does, the second-lowest price silver plan in each area as our unit.

Bordering county characteristics tend to vary little across state lines, and given this limited difference we expect there to be fewer selection issues that vastly differ between these neighboring counties. An issue with this strategy is that not all states define their

[15] P_p comprises the following plan characteristics: plan type (HMO, PPO, etc.) and metal level.[16] M_m consists of the following market characteristics: Log(Population); average Medicare spending per enrollee, to account for geographic variation in both health and cost of healthcare; urban level; and percent of county in poverty.

rating areas based on geographic areas. States have defined their rating areas based on either counties (or small collections of geographically connected counties), 3-digit zip codes, area codes, or MSAs + 1 where each metropolitan area is its own rating area and then the rest of the state is combined into one.[17,18] We run the same model above with the exclusion of state effects. The reason for this exclusion is that none of the states who switch (allowing for state effects) create a new border by doing so.[19]

Our difference in differences strategy relies on variation derived from several states that initially opted out of the exchange decided to expand later. Specifically, our identification is derived from Michigan, New Hampshire, and Pennsylvania expanding during 2014 and from Montana and Indiana expanding during 2015.[20] Our analysis is slightly different from a traditional diff-in-diff (defining expansion as treatment) in that we have three groups: states who are never treated, states who we observe before and after treatment, and states who we only observe after treatment. Our model is therefore defined as:

$$Prem_{pmt} = \beta AlwaysExpand_m + \delta Switch_m + Year + \theta Switch_m * Expand_{mt} + \alpha P_p + \gamma M_m + F + \epsilon_{pm}$$

Where $AlwaysExpand_m$ represents the states that expanded Medicaid prior to the 2014 markets and $Switch_m$ defines the states that had not expanded at the start of 2014, but had by the start of 2016. θ is our coefficient of interest, defining the effect of expanding Medicaid on premium levels. All errors for all models presented are clustered at the state level.

[17] A list of each state's rating area definitions can be found at https://www.cms.gov/cciio/programs-and-initiatives/health-insurance-market-reforms/state-gra.html

[18] The "natural" way to define ratings areas was MSAs + 1. States had to apply to obtain a different definition.

[19] The only new borders created are from states with MSA + 1 ratings areas such as the North Dakota - Montana border. Other states, such as Pennsylvania, were surrounded by expansion states already.

[20] Open enrollment for the exchanges is from October through December. We define states that expand on January 1st as having expanded in the year prior to account for expectations.

3.5 Results:

	Dependent variable:				
	Monthly Premium				
	All	Bronze	Silver	Gold	Platinum
	(1)	(2)	(3)	(4)	(5)
Expand	−24.30**	−12.52*	−24.07***	−29.42*	−37.72
	(9.92)	(7.59)	(8.68)	(15.74)	(24.76)
HMO	−2.56	−1.82	2.41	−5.03	−28.41**
	(7.88)	(6.16)	(7.45)	(7.39)	(12.17)
POS	13.65*	12.51*	16.12**	16.62**	9.09
	(6.98)	(7.27)	(6.33)	(8.24)	(10.28)
PPO	37.30***	32.42***	41.15***	40.67***	17.53
	(9.93)	(10.23)	(10.15)	(9.52)	(14.57)
Log(Population)	−0.07	−0.49	−0.65	0.46	2.90*
	(1.33)	(1.03)	(1.11)	(1.09)	(1.59)
MSPE ($000s)	2.25	2.10	1.94	1.76	5.62**
	(1.83)	(1.42)	(1.77)	(1.95)	(2.67)
Urban	2.17***	2.13***	2.00**	2.58***	1.83*
	(0.73)	(0.66)	(0.79)	(0.86)	(1.06)
Percent Poverty	−5.05***	−4.63***	−4.85***	−6.40***	−47.59
	(0.98)	(0.72)	(1.14)	(1.45)	(67.47)
Year 2015	3.38	−5.29	9.47	6.01	5.85
	(14.66)	(14.18)	(12.64)	(14.43)	(37.90)
Year 2016	30.25**	10.83	24.62*	50.89***	99.88***
	(13.32)	(16.22)	(12.73)	(13.63)	(36.16)
Constant	138.01***	152.18***	195.84***	230.96***	319.79***
	(24.72)	(24.76)	(27.03)	(27.73)	(31.56)
Metal Levels	Yes	No	No	No	No
Insurer Effects?	Yes	Yes	Yes	Yes	Yes
State Effects?	Yes	Yes	Yes	Yes	Yes
Observations	57,372	17,409	22,903	13,634	3,426
R^2	0.55	0.44	0.41	0.46	0.69
Adjusted R^2	0.55	0.43	0.40	0.45	0.68
Residual Std. Error	59.28	68.09	50.55	58.55	47.16
F Statistic	237.84***	47.70***	53.78***	38.72***	60.12***

Note: *p<0.1; **p<0.05; ***p<0.01
This table represents OLS regressions for individual plans by metal level in every exchange except Hawaii and Alaska for 2014 - 2016. Premiums are defined for a 27 year old male on the individual market. MSPE is the average Medicare spending per enrollee in the rating area ($000s). Cluster robust (at the State level) standard errors are in parenthesis.

Table 3.1: OLS Regression Results by Metal Level

Table 3.1 provides the results of our initial analysis using OLS. Our initial analysis indicates a positive relationship between price and the decision to opt out of the Medicaid expansion, indicating higher private insurance costs in the states that opted out of the Medicaid expansion. When aggregating all plan levels across the exchanges, we find that on average states who expand Medicaid see a monthly premium decrease of \$24.3, a decrease of 8.89%. This estimated effect is similar to Sen and Deleire's (2016) effect of 7%. This aggregate state analysis is evidence to suggest that the effects of opting out of the Medicaid expansion may be driving in riskier entrants towards the private market exchanges. The positive relationship between price and opting out provides evidence to support our "crowding in" hypothesis.

To differentiate the effect opting out has on low and high coverage plans, the latter columns of Table 3.1 perform the same analysis sub-setting by metal level. The larger effect on the the gold plans relative to bronze plans (9.28% for gold vs. 5.55% for bronze) supports the notion of an adverse selection issue. However, the large effect on silver plans (8.96%) as well as the still significant effect on bronze plans lends credence to the negative income health correlation hypothesis.

Our preferred specification is our difference in differences model, the results of which are presented in Table 3.2. The results from our initial regression were biased towards zero, leading to increased economic (and statistical) significance of our key results. This attenuation bias was esspecially strong for platinum plans, when we control for the fact that states who switch are different from states who do not we see expansion has a large, negative effect on monthly platinum plan premiums (\$113.33 or 30.47%). This result strongly suggests that expansion of Medicaid leads to a reduction in to the adverse selection problems that are plaguing the exchanges.

Using the same diff-in-diff strategy, we also examine differences in expansion effects for different age groups as well as individual vs. small group (SHOP) plans. These results are

	Dependent variable:				
	Monthly Premium				
	All	Bronze	Silver	Gold	Platinum
	(1)	(2)	(3)	(4)	(5)
Expand Effect	−32.40***	−19.43***	−30.89***	−39.20**	−113.33***
	(9.38)	(7.27)	(8.47)	(15.35)	(43.63)
Switch	16.41	8.52	19.01	15.48	48.02
	(14.07)	(10.04)	(12.87)	(20.93)	(55.43)
Always Expand	−11.93	−13.17*	−10.20	−18.49*	17.69
	(8.45)	(7.02)	(9.18)	(10.27)	(17.94)
Year 2015	6.04	−2.53	13.85	6.42	16.41
	(20.20)	(19.79)	(15.59)	(16.67)	(49.43)
Year 2016	30.29**	10.79	23.80*	51.73***	100.11***
	(13.41)	(16.36)	(12.81)	(13.50)	(35.45)
HMO	−8.80	−5.84	−4.45	−12.88	−30.92***
	(7.42)	(5.28)	(6.88)	(10.95)	(10.69)
POS	13.28	14.13*	15.65*	14.13	7.06
	(8.42)	(8.06)	(8.38)	(12.56)	(9.62)
PPO	24.69**	21.76**	27.63***	27.56**	1.21
	(9.68)	(9.03)	(9.46)	(12.49)	(13.84)
Log(Population)	−0.35	−0.79	−1.13	0.40	1.83
	(1.72)	(1.27)	(1.32)	(1.02)	(3.47)
MSPE ($000s)	−2.29	−2.84	−2.61	−2.37	2.88
	(2.17)	(1.86)	(2.25)	(2.18)	(3.37)
Urban	1.73**	1.62**	1.58**	2.25**	1.52
	(0.74)	(0.66)	(0.78)	(0.88)	(1.30)
Percent Poverty	−3.48***	−3.28***	−3.30***	−4.59***	−32.13
	(0.74)	(0.47)	(0.98)	(1.04)	(82.91)
Constant	227.85***	241.99***	283.24***	315.60***	418.06***
	(38.03)	(37.95)	(39.14)	(38.39)	(60.07)
Metal Levels	Yes	No	No	No	No
Insurer Effects?	Yes	Yes	Yes	Yes	Yes
Observations	57,372	17,409	22,903	13,634	3,426
R^2	0.52	0.42	0.36	0.42	0.64
Adjusted R^2	0.52	0.41	0.35	0.41	0.63
Residual Std. Error	61.05	69.39	52.55	60.72	50.57
F Statistic	243.57***	50.04***	50.22***	37.40***	61.05***

Note: *p<0.1; **p<0.05; ***p<0.01
This table represents Diff - in - Diff regressions for individual plans by metal level in every exchange except Hawaii and Alaska for 2014 - 2016. Premiums are defined for a 27 year old male on the individual market. MSPE is the average Medicare spending per enrollee in the rating area ($000s). Cluster robust (at the State level) standard errors are in parenthesis.

Table 3.2: Diff - in - Diff Regression Results by Metal Level

	Dependent variable:			
	Monthly Premium			
	Age 27 Ind.	Age 40 Ind.	Age 27 SHOP	Age 40 SHOP
	(1)	(2)	(3)	(4)
Expand Effect	-32.40^{***}	-46.73^{***}	-24.93^{**}	-13.77
	(9.38)	(14.44)	(12.71)	(28.32)
Switch	16.41	28.73	12.05	-1.40
	(14.07)	(18.85)	(22.29)	(37.64)
Always Expand	-11.93	-8.87	-4.53	-6.00
	(8.45)	(11.29)	(16.11)	(22.29)
Year 2015	6.04	-14.06	51.65^{***}	-25.29
	(20.20)	(27.46)	(14.69)	(17.61)
Year 2016	30.29^{**}	17.98^{***}	26.67^{***}	-62.73^{***}
	(13.41)	(2.16)	(7.90)	(12.03)
Metal Level Gold	98.42^{***}	123.98^{***}	108.81^{***}	133.12^{***}
	(3.22)	(2.67)	(2.57)	(3.38)
Metal Level Platinum	145.08^{***}	192.14^{***}	160.03^{***}	200.44^{***}
	(9.29)	(7.39)	(6.65)	(8.59)
Metal Level Silver	45.76^{***}	56.30^{***}	54.99^{***}	65.64^{***}
	(1.97)	(2.38)	(1.92)	(2.51)
HMO	-8.80	-10.19	-10.31	-13.29
	(7.42)	(8.43)	(14.35)	(18.57)
POS	13.28	17.32^{*}	8.93	11.40
	(8.42)	(9.98)	(14.72)	(19.54)
PPO	24.69^{**}	24.66^{***}	27.02^{**}	34.01^{*}
	(9.68)	(9.27)	(13.70)	(17.72)
Log(Population)	-0.35	-0.19	-3.25^{***}	-4.47^{***}
	(1.72)	(2.74)	(1.20)	(1.31)
MSPE ($000s)	-2.29	-3.62	4.65^{***}	5.79^{***}
	(2.17)	(2.79)	(1.32)	(1.82)
Urban	1.73^{**}	2.91^{**}	1.14	1.91^{*}
	(0.74)	(1.20)	(0.89)	(1.16)
Percent Poverty	-3.48^{***}	-2.80^{**}	-3.42^{***}	-7.54^{**}
	(0.74)	(1.32)	(1.25)	(3.17)
Constant	227.85^{***}	275.33^{***}	165.61^{***}	297.06^{***}
	(38.03)	(34.63)	(32.77)	(32.93)
Insurer Effects?	Yes	Yes	Yes	Yes
Observations	57,372	42,922	31,692	24,462
R^2	0.52	0.70	0.74	0.73
Adjusted R^2	0.52	0.70	0.74	0.73
Residual Std. Error	61.05	46.35	36.87	45.86
F Statistic	243.57^{***}	545.24^{***}	526.12^{***}	536.26^{***}
Note:	*p<0.1; **p<0.05; ***p<0.01 This table represents Diff - in - Diff regressions for individual and small group (SHOP) plans in every exchange except Hawaii and Alaska for 2014 - 2016. premiums are defined for both 27 and 40 year old males. Cluster robust (at the State level) standard errors are in parenthesis. MSPE is the average Medicare spending per enrollee in the rating area ($000s).			

Table 3.3: Diff - in - Diff Regression Results Individual vs. Shop

presented in Table 3.3. If adverse selection is the primary issue, we would expect to see a difference in effect between individual and SHOP plans with individual plans seeing a greater reduction in premiums due to expansion. We would also expect to see this difference be larger for the the older age group, as there is likely to be a greater prevalence of assymetric knowledge of health status.[21] The point estimates for the effects of differences between individual and SHOP for age 27 have the predicted relationship ($|\beta_{ind}| > |\beta_{SHOP}|$), however the coefficients are not statistically different from each other. The difference in coefficients for age 40 is again as predicted and, while larger, this difference is also not statistically significant.

As a robustness check on our diff-in-diff identification strategy, we separately control for endogeneity in the decision to expand or not expand Medicaid by analyzing border counties in neighboring opt-in/opt-out states. The results from this analysis are presented in Table 3.4.[22] This analysis loses quite a bit of power, as sample restriction is ought to do, but the results remain generally consistent. Interestingly, we lose significance (though barely) on the silver plans. The strongly negative effects on both bronze and gold plans signals that Medicaid expansion helps to alleviate adverse selection issues in the private market as well as increases the overall health of the private pool by sending the least healthy to Medicaid.

[21]Largely due to many chronic health issues usually being discovered during your 30s. An example would be the common recommendation that men have their prostate checked when they turn 30.

[22]We do not include platinum plans in our analysis due to their rarity. Many matched counties do not both contain platinum level plans to compare.

	Dependent variable:			
	Monthly Premium			
	All	Bronze	Silver	Gold
	(1)	(2)	(3)	(4)
Expand	−26.05**	−23.01*	−28.39	−38.66**
	(12.98)	(12.88)	(17.74)	(18.79)
HMO	−13.94	9.98	8.38	4.38
	(15.10)	(9.05)	(11.84)	(15.71)
POS	20.11	35.80**	41.70***	41.53**
	(13.97)	(17.47)	(15.99)	(18.38)
PPO	7.85	27.36**	30.49**	38.15**
	(15.03)	(10.61)	(13.40)	(16.04)
Log(Population)	−6.32	0.25	1.80	2.79
	(3.91)	(3.68)	(4.72)	(4.57)
MSPE ($000s)	−7.13	−5.48	−8.91**	−6.57
	(5.03)	(3.71)	(3.98)	(4.82)
Urban	−2.10	0.95	0.28	2.59
	(2.46)	(1.57)	(1.87)	(2.13)
Percent Poverty	−2.50	−108.01	−109.92	−196.50**
	(139.69)	(69.11)	(85.58)	(95.54)
Year 2015	74.41*	−3.53	−10.03	−16.36
	(43.09)	(39.45)	(47.31)	(50.95)
Year 2016	38.56***	28.03**	38.09***	57.05***
	(8.78)	(11.46)	(11.91)	(19.55)
Constant	310.45***	254.04***	329.46***	360.01***
	(48.76)	(30.25)	(34.89)	(46.82)
Metal Levels	Yes	No	No	No
Insurer Effects?	No	Yes	Yes	Yes
Observations	5,262	1,727	2,071	1,284
R^2	0.50	0.65	0.63	0.66
Adjusted R^2	0.50	0.62	0.61	0.63
Residual Std. Error	47.85	28.13	32.62	36.40
F Statistic	404.85***	25.01***	27.04***	18.65***

Note: *p<0.1; **p<0.05; ***p<0.01
This table represents OLS regressions for individual plans by metal level in exchanges located in border counties where states differ in their expansion decision. Cluster robust (at the State level) standard errors are in parenthesis. MSPE is the average Medicare spending per enrollee in the rating area (in $000s). States who use MSA + 1 for ratings areas are not included in the matching procedure.

Table 3.4: Georgraphic Discontinuity Regression

3.6 Conclusions:

While the majority of the ACA held up to judicial scrutiny, the Supreme Court ruled that mandatory expansion of Medicaid was unconstitutional. This gave states the option of "opting out" of the expansion. The debate, in both congressional and judicial halls, over the expansion largely weighed the direct benefits to the individuals newly covered vs. the accounting cost of doing so Harrington (2010b). In this article, we examine potential effects that a state's expansion of Medicaid may have on the private market for health insurance. We find that expanding Medicaid reduces average monthly premiums by \$32.4; a decrease of 11.86%.

We examine two different avenues for this reduction in premiums to occur: adverse selection reduction and the income/health correlation, finding evidence that both are occuring simultaneously. The large increase in prices in gold and platinum plans from 2015 to 2016 provides limited evidence of an adverse selection "death spiral" occurring in the higher coverage plans. Our evidence suggests that this could be partially alleviated by all of the states expanding Medicaid. The ACA is still hotly debated, with many conservatives seeking to repeal the law. Our analysis is informative on the cause and extent of second-order effects of Medicaid expansion, but is too limited to make any broad policy recommendations; especially since we do not consider the effect of price increases on the increased size of subsidies. Instead, we seek to expand on the standard welfare analysis currently being debated with regard to the Medicaid expansion. An exhaustive examination of the true welfare effects of expanding public insurance to the poor is an excellent avenue for further research for which the ACA provides a unique laboratory for.

References

Abraham, Jean M and Pinar Karaca-Mandic, "Regulating the Medical Loss Ratio: Implications for the Individual Market," *Am J Manag Care*, 2011, *17* (3), 211–218.

Anscombe, Frank J and John W Tukey, "The Examination and Analysis of Residuals," *Technometrics*, 1963, *5* (2), 141–160.

Arkes, Hal R and Catherine Blumer, "The Psychology of Sunk Cost," *Organizational Behavior and Human Decision Processes*, 1985, *35* (1), 124–140.

Basu, Anirban and NB Coe, "2SLS vs 2SRI: Appropriate Methods for Rare Outcomes and/or Rare Exposures," *Working Paper*, 2015.

Becker, Gary S, "A Theory of the Allocation of Time," *The Economic Journal*, 1965, pp. 493–517.

Bellman, Richard, "Dynamic Programming," *Princeton University Press*, 1957.

_ **and Stuart Dreyfus**, "Functional Approximations and Dynamic Programming," *Mathematical Tables and Other Aids to Computation*, 1959, *13* (68), 247–251.

Belloni, Alexandre, Daniel Chen, Victor Chernozhukov, and Christian Hansen, "Sparse Models and Methods for Optimal Instruments with an Application to Eminent Domain," *Econometrica*, 2012, *80* (6), 2369–2429.

_ **, Victor Chernozhukov, and Christian Hansen**, "High-Dimensional Methods and Inference on Structural and Treatment Effects," *The Journal of Economic Perspectives*, 2014, *28* (2), 29–50.

_ **,** _ **, and** _ , "Inference on Treatment Effects after Selection among High-Dimensional Controls," *The Review of Economic Studies*, 2014, *81* (2), 608–650.

_ **,** _ **,** _ **, and Damian Kozbur**, "Inference in High-Dimensional Panel Models With an Application to Gun Control," *Journal of Business & Economic Statistics*, 2016, *34* (4), 590–605.

Bertsekas, Dimitri P and John N Tsitsiklis, "Neuro-dynamic Programming (Optimization and Neural Computation series, 3)," *Athena Scientific*, 1996, *7*, 15–23.

Blackstone, Erwin A, "Restrictive Practices in the Marketing of Electrofax Copying Machines and Supplies: the SCM Corporation Case," *The Journal of Industrial Economics*, 1975, *23* (3), 189–202.

Bloom, David E and David Canning, "The Health and Wealth of Nations," *Science*, 2000, *287* (5456), 1207–1209.

Bowman, Ward S, "Tying Arrangements and the Leverage Problem," *The Yale Law Journal*, 1957, *67* (1), 19–36.

Breiman, Leo, "Random Forests," *Machine Learning*, 2001, *45* (1), 5–32.

Brown, Jeffrey R and Austan Goolsbee, "Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry," *Journal of Political Economy*, 2002, *110* (3), 481–507.

Bucks, Brian K, Jesse Bricker et al., "Household Mobility Over the Great Recession: Evidence from the US 2007-09 Survey of Consumer Finances Panel," *Finance and Economics Discussion Series*, 2013, *53.*

Campbell, Ritchie A, "The Demand for Life Insurance: An Application of the Economics of Uncertainty," *The Journal of Finance*, 1980, *35* (5), 1155–1172.

Carson, James M, "Backdating of Life Insurance Contracts: An Examination," *Journal of Insurance Regulation*, 1994, *13* (2), 185–200.

_ **and Krzysztof Ostaszewski**, "The Actuarial Value of Life Insurance Backdating," *Journal of Actuarial Practice*, 2004, *11* (1), 63–77.

_ **and Mark D Forster**, "Suitability and Life Insurance Policy Replacement: An Analytical Tool," *Journal of Insurance Regulation*, 2000, *18* (4), 427–447.

_ **, Christopher Clark, and Krzysztof Ostaszewski**, "The Incentives and Welfare Economics of Life Insurance Backdating.," *Journal of Insurance Regulation*, 2012, *31*, 91 – 104.

Case, Anne, Darren Lubotsky, and Christina Paxson, "Economic status and health in childhood: The origins of the gradient," *The American Economic Review*, 2002, *92* (5), 1308–1334.

Cawley, John and Tomas Philipson, "An Empirical Examination of Information Barriers to Trade in Insurance," *American Economic Review*, 1999, *89* (4), 827–846.

Chernozhukov, Victor, Christian Hansen, and Martin Spindler, "Post-Selection and Post-Regularization Inference in Linear Models with many Controls and Instruments," *The American Economic Review*, 2015, *105* (5), 486–490.

Chiappori, Pierre-André, Bruno Jullien, Bernard Salanié, and Francois Salanie, "Asymmetric Information in Insurance: General Testable Implications," *The RAND Journal of Economics*, 2006, *37* (4), 783–798.

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

Cutler, David M, Amy Finkelstein, and Kathleen McGarry, "Preference Heterogeneity and Insurance Markets: Explaining a Puzzle of Insurance," *The American Economic Review*, 2008, *98* (2), 157–162.

_ **and Jonathan Gruber**, "The Effect of Medicaid Expansions on Public Insurance, Private Insurance, and Redistribution," *The American Economic Review*, 1996, pp. 378–383.

_ **and Sarah J Reber**, "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection," *The Quarterly Journal of Economics*, 1998, *113* (2), 433–466.

Cybenko, George, "Approximation by Superpositions of a Sigmoidal Function," *Mathematics of Control, Signals and Systems*, 1989, *2* (4), 303–314.

Dafny, Leemore S, "Are Health Insurance Markets Competitive?," *The American Economic Review*, 2010, pp. 1399–1431.

Dai, Hengchen, Katherine L Milkman, and Jason Riis, "The Fresh Start Effect: Temporal Landmarks Motivate Aspirational Behavior," *Management Science*, 2014, *60* (10), 2563–2582.

Deaton, Angus, "Policy Implications of the Gradient of Health and Wealth," *Health Affairs*, 2002, *21* (2), 13–30.

Deaton, Angus S and Christina H Paxson, "Aging and Inequality in Income and Health," *The American Economic Review*, 1998, *88* (2), 248–253.

_ **and Christina Paxson**, "Mortality, Education, Income, and Inequality among American Cohorts," in "Themes in the Economics of Aging," University of Chicago Press, 2001, pp. 129–170.

Depew, Briggs and James Bailey, "Did the Affordable Care Act's Dependent Coverage Mandate Increase Premiums?," *Journal of Health Economics*, 2015.

Doherty, Neil A, "Portfolio Efficient Insurance Buying Strategies," *Journal of Risk and Insurance*, 1984, pp. 205–224.

_ **and Hal J Singer**, "The Benefits of a Secondary Market for Life Insurance Policies," *Real Property, Probate and Trust Journal*, 2003, pp. 449–478.

Doherty, Neil A. and Hal J. Singer, "Regulating the Secondary Market for Life Insurance Policies.," *Journal of Insurance Regulation*, 2003, *21* (4).

Dube, Arindrajit, T William Lester, and Michael Reich, "Minimum Wage Effects across State Borders: Estimates using Contiguous Counties," *The review of economics and statistics*, 2010, *92* (4), 945–964.

Einav, Liran and Amy Finkelstein, "Selection in Insurance Markets: Theory and Empirics in Pictures," *The Journal of Economic Perspectives*, 2011, *25* (1), 115.

Eling, Martin and Michael Kochanski, "Research on Lapse in Life Insurance: What has been Done and What Needs to be Done?," *The Journal of Risk Finance*, 2013, *14* (4), 392–413.

Fang, Hanming and Edward Kung, "How Does Life Settlement Affect the Primary Life Insurance Market?," *NBER Working Paper Series*, 2010, p. 15761.

Fei, Wenan, Claude Fluet, and Harris Schlesinger, "Uncertain Bequest Needs and Long-Term Insurance Contracts," *Journal of Risk and Insurance*, 2015, *82* (1), 125–148.

Fischer, Stanley, "A Life Cycle Model of Life Insurance Purchases," *International Economic Review*, 1973, pp. 132–152.

Gabor, Andre, "A Note on Block Tariffs," *The Review of Economic Studies*, 1955, *23* (1), 32–41.

Gatzert, Nadine, Gudrun Hoermann, and Schmeiser' Hato, "The Impact of the Secondary Market on Life Insurers' Surrender Profits," *Journal of Risk and Insurance*, 2009, *76* (4), 887–908.

Goeman, Jelle J, "L1 Penalized Estimation in the Cox Proportional Hazards Model," *Biometrical Journal*, 2010, *52* (1), 70–84.

Golberg, David E, "Genetic Algorithms in Search, Optimization, and Machine Learning," *Addion wesley*, 1989, *1989*, 102.

Gottlieb, Daniel, "Prospect Theory, Life Insurance, and Annuities," *The Wharton School Research Paper*, 2012, (44).

_ , "Prospect Theory, Life Insurance, and Annuities," *Working Paper*, 2012.

_ **and Kent Smetters**, "Lapse-Based Insurance," *Manuscript submitted for publication*, 2014.

Gowrisankaran, Gautam and John Krainer, "Entry and Pricing in a Differentiated Products Industry: Evidence from the ATM Market," *The RAND Journal of Economics*, 2011, *42* (1), 1–22.

Gruber, Jonathan and Kosali Simon, "Crowd-out 10 years later: Have Recent Public Insurance Expansions Crowded out Private Health Insurance?," *Journal of Health Economics*, 2008, *27* (2), 201–217.

Gutter, Michael S and Charles B Hatcher, "Racial Differences in the Demand for Life Insurance," *Journal of Risk and Insurance*, 2008, *75* (3), 677–689.

Ham, John C and Lara Shore-Sheppard, "The Effect of Medicaid Expansions for Low-income Children on Medicaid Participation and Private Insurance Coverage: Evidence from the SIPP," *Journal of Public Economics*, 2005, *89* (1), 57–83.

Handel, Ben, Igal Hendel, and Michael D Whinston, "Equilibria in Health Exchanges: Adverse Selection Versus Reclassification Risk," *Econometrica*, 2015, *83* (4), 1261–1313.

Harrington, Scott E, "The Health Insurance Reform Debate," *Journal of Risk and Insurance*, 2010, *77* (1), 5–38.

_ , "US Healthcare Reform: The Patient Protection and Affordable Care Act," *Journal of Risk and Insurance*, 2010, *77* (3), 703–708.

Heckman, James J and V Joseph Hotz, "Choosing Among Alternative Nonexperimental Methods for Estimating the Impact of Social Programs: The Case of Manpower Training," *Journal of the American statistical Association*, 1989, *84* (408), 862–874.

Hendel, Igal and Alessandro Lizzeri, "The Role of Commitment in Dynamic Contracts: Evidence from Life Insurance," *The Quarterly Journal of Economics*, 2003, *118* (1), 299–328.

Herring, Bradley and Mark V Pauly, "Incentive-Compatible Guaranteed Renewable Health Insurance Premiums," *Journal of Health Economics*, 2006, *25* (3), 395–417.

Hilliard, J. I., A. P. Liebenberg, I. A. Liebenberg, and J. S. Ruhland, "The Market Impact of the Supreme Court Decision Regarding the Patient Protection and Affordable Care Act: Evidence from the Health Insurance Industry," *Working Paper*, 2013.

Ho, Teck-Hua, Ivan PL Png, and Sadat Reza, "Sunk Cost Fallacy in Driving the World's Costliest Cars," *Management Science*, 2017.

Hofmann, Annette and Mark Browne, "One-Sided Commitment in Dynamic Insurance Contracts: Evidence from Private Health Insurance in Germany," *Journal of Risk and Uncertainty*, 2013, *46* (1), 81–112.

Holahan, John, "The Cost and Coverage Implications of the ACA Medicaid Expansion: National and State-by-state Analysis," 2012.

Hornik, Kurt, "Approximation Capabilities of Multilayer Feedforward Networks," *Neural Networks*, 1991, *4* (2), 251–257.

Joskow, Paul L, "Cartels, Competition and Regulation in the Property-Liability Insurance Industry," *The Bell Journal of Economics and Management Science*, 1973, pp. 375–427.

Judd, Kenneth L, *Numerical Methods in Economics*, MIT press, 1998.

Karni, Edi and Itzhak Zilcha, "Risk Aversion in the Theory of Life Insurance: the Fisherian Model," *Journal of Risk and Insurance*, 1986, pp. 606–620.

Kennickell, Arthur B and Martha Starr-Mccluer, "Household Saving and Portfolio Change: Evidence from the 1983–89 SCF Panel," *Review of Income and Wealth*, 1997, *43* (4), 381–399.

Kuo, Weiyu, Chenghsien Tsai, and Wei-Kuang Chen, "An Empirical Study on the Lapse Rate: The Cointegration Approach," *Journal of Risk and Insurance*, 2003, *70* (3), 489–508.

Lewis, Frank D, "Dependents and the Demand for Life Insurance," *The American Economic Review*, 1989, *79* (3), 452–467.

Liebenberg, Andre P, James M Carson, and Randy E Dumm, "A Dynamic Analysis of the Demand for Life Insurance," *Journal of Risk and Insurance*, 2012, *79* (3), 619–644.

_ , _ , and Robert E Hoyt, "The Demand for Life Insurance Policy Loans," *Journal of Risk and Insurance*, 2010, *77* (3), 651–666.

Lin, Yijia and Martin F Grace, "Household Life Cycle Protection: Life Insurance Holdings, Financial Vulnerability, and Portfolio Implications," *Journal of Risk and Insurance*, 2007, *74* (1), 141–173.

Linton, MA, "Panics and Cash Values," *Transactions of the Actuarial Society of America*, 1932, *33*, 265–394.

Mayers, David and Clifford W Smith Jr, "The Interdependence of Individual Portfolio Decisions and the Demand for Insurance," *The Journal of Political Economy*, 1983, pp. 304–311.

McAfee, R Preston, Hugo M Mialon, and Sue H Mialon, "Do Sunk Costs Matter?," *Economic Inquiry*, 2010, *48* (2), 323–336.

McCulloch, Warren S and Walter Pitts, "A Logical Calculus of the Ideas Immanent in Nervous Activity," *The Bulletin of Mathematical Biophysics*, 1943, *5* (4), 115–133.

Meza, David De and David C Webb, "Advantageous Selection in Insurance Markets," *RAND Journal of Economics*, 2001, pp. 249–262.

Mulholland, Barry S and Michael S Finke, "Does Cognitive Ability Impact Life Insurance Policy Lapsation?," *Available at SSRN 2512076*, 2014.

Nahata, Babu, Krzysztof Ostaszewski, and Prasanna Sahoo, "Buffet Pricing," *The Journal of Business*, 1999, *72* (2), 215–228.

Nardi, Mariacristina De, Eric French, and John B Jones, "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 2010, *118* (1), 39–75.

Oi, Walter Y, "A Disneyland Dilemma: Two-Part Tariffs for a Mickey Mouse Monopoly," *The Quarterly Journal of Economics*, 1971, pp. 77–96.

Outreville, J Francois, "Whole-life Insurance Lapse Rates and the Emergency Fund Hypothesis," *Insurance: Mathematics and Economics*, 1990, *9* (4), 249–255.

Pashchenko, Svetlana, "Accounting for Non-Annuitization," *Journal of Public Economics*, 2013, *98*, 53–67.

_ **and Ponpoje Porapakkarm**, "Welfare Costs of Reclassification Risk in the Health Insurance Market," *Journal of Macroeconomics*, 2015, *45*, 21–44.

Pauly, Mark, Adam Leive, and Scott Harrington, "The Price of Responsibility: The Impact of Health Reform on Non-Poor Uninsureds," Working Paper 21565, National Bureau of Economic Research September 2015.

Pesando, James E, "The Interest Sensitivity of the Flow of Funds through Life Insurance Companies: an Econometric Analysis," *Journal of Finance*, 1974, *29* (4), 1105–1121.

Powell, Warren B, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Vol. 703, John Wiley & Sons, 2007.

Purushotham, Marianne, "US Individual Life Persistency Update," *LIMRAS Marketfacts Quarterly*, 2006, *25* (1), 36.

Puterman, Martin L, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley & Sons, 2014.

Rothschild, Michael and Joseph Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *The Quarterly Journal of Economics*, 1976, *90* (4), 629–649.

Rust, John, "Using Randomization to Break the Curse of Dimensionality," *Econometrica: Journal of the Econometric Society*, 1997, pp. 487–516.

Schmalense, Richard, "Monopolistic Two-Part Pricing Arrangements," *The Bell Journal of Economics*, 1981, *12* (2), 445–466.

_ , "Pricing the Razor: A Note on Two-Part Tariffs," *International journal of industrial organization*, 2015, *42*, 19–22.

Schott, Francis H, "Disintermediation Through Policy Loans at Life Insurance Companies," *The Journal of Finance*, 1971, *26* (3), 719–729.

Selden, Thomas M, Jessica S Banthin, and Joel W Cohen, "Medicaid's Problem Children: Eligible but not Enrolled," *Health Affairs*, 1998, *17* (3), 192–200.

Sen, Aditi P. and Thomas Deleire, "The Effect of Medicaid Expansion on Marketplace Premiums," *ASPE Issue Brief*, 2016.

Shefrin, Hersh, *Beyond Greed and Fear: Understanding Behavioral Finance and the Psychology of Investing*, Oxford University Press on Demand, 2002.

Sommers, Benjamin D, Thomas Buchmueller, Sandra L Decker, Colleen Carey, and Richard Kronick, "The Affordable Care Act has led to Significant Gains in Health Insurance and Access to Care for Young Adults," *Health affairs*, 2013, *32* (1), 165–174.

Stole, Lars A, "Price Discrimination and Competition," *Handbook of industrial organization*, 2007, *3*, 2221–2299.

Terza, Joseph V, Anirban Basu, and Paul J Rathouz, "Two-Stage Residual Inclusion Estimation: Addressing Endogeneity in Health Econometric Modeling," *Journal of Health Economics*, 2008, *27* (3), 531–543.

Tibshirani, Robert, "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, 1996, pp. 267–288.

_ **et al.**, "The Lasso Method for Variable Selection in the Cox Model," *Statistics in medicine*, 1997, *16* (4), 385–395.

Tirole, Jean, *The Theory of Industrial Organization*, MIT press, 1988.

Tsitsiklis, John N, "Asynchronous Stochastic Approximation and Q-learning," *Machine Learning*, 1994, *16* (3), 185–202.

_ **and Benjamin Van Roy**, "An Analysis of Temporal-Difference Learning with Function Approximation," *IEEE Transactions on Automatic Control*, 1997, *42* (5), 674–690.

Wager, Stefan and Susan Athey, "Estimation and Inference of Heterogeneous Treatment Effects using Random Forests," *arXiv preprint arXiv:1510.04342*, 2015.

Werbos, Paul, "Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences," 1974.

Yaari, Menahem E, "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer," *The Review of Economic Studies*, 1965, *32* (2), 137–150.