Essays on the Economics of Insurance and Healthcare Markets

by

Patricia Alexander Robinson

Department of Economics
Duke University

Date: ______________________

Approved:

____
Frank Sloan, Supervisor

____
James Roberts

____
Peter Arcidiacono

____
Donna Gilleskie

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2017

Abstract

Essays on the Economics of Insurance and Healthcare Markets

by

Patricia Alexander Robinson

Department of Economics
Duke University

Date: ______________________
Approved:

____
Frank Sloan, Supervisor

____
James Roberts

____
Peter Arcidiacono

____
Donna Gilleskie

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics in the Graduate School of Duke University
2017

Abstract

This dissertation contains three chapters that investigate the role of asymmetric information in determining market outcomes in insurance and healthcare markets. The first two chapters focus on the U.S. automobile insurance market and how asymmetric information about consumer characteristics affect consumer and insurer behavior. The third chapter focuses on the U.S. healthcare market and studies how the financial incentives that arise due to price variation in the multi-payer system affect physician behavior when the physician has more information and the discretion to choose a treatment for the patient.

The first chapter, co-authored with Frank Sloan and Lindsey Eldred, quantifies the role of private information in automobile insurance policy choice using data on individuals' subjective beliefs, risk preference, reckless driving, insurer, and insurance policy characteristics merged with insurer-specific quality ratings distributed by independent organizations. We find a zero correlation between ex post accident risk and insurance coverage, reflecting advantageous selection in policy choice offset by moral hazard. Advantageous selection is partly attributable to insurer sorting on consumer attributes known and used by insurers. Our analysis of insurer sorting reveals that lower-risk drivers on attributes observed by insurers obtain coverage from insurers with higher-quality ratings.

The second chapter extends the work of the first chapter to quantifying the welfare impact of private information in an insurance market and evaluating potential policy

interventions. Many studies show that asymmetric information exists in insurance markets, yet there is little consensus on the effectiveness of interventions in these markets. This chapter provides empirical evidence that improving information about risk is not always welfare-improving. I show how the theoretical effect of risk-rating can depend on whether the market is adversely- or advantageously-selected. I then estimate a structural model of insurance choice and reckless behavior to show that in the advantageously-selected U.S. automobile insurance market, risk-rating induces high-risk drivers to drop coverage, creating a negative externality and social welfare loss. Community rating improves welfare despite increasing asymmetric information.

The third chapter considers the impact of cross-payer price variation on physician behavior in the U.S., specifically for the case of births. A key innovation of the paper is to use data from multiple private payers—the Massachusetts All-Payer Claims Database. With these data, I ask how the change in reimbursement from one payer affects the probability that a physician performs a C-section on patients insured by that payer and by other payers. I use a difference-in-differences strategy that takes advantage of variation in contract change dates across the three largest private payers in Massachusetts. The results show that physicians are less likely to perform a C-section on a patient when the relative price the doctor receives from her insurer for a C-section decreases. This effect is concentrated among patients classified as medium-risk based on factors observed before the delivery, specifically women who reportedly experience long labor. These findings suggests that prices can be used as an incentive to change physician behavior, at least in the case of births. Whether this is welfare-improving depends on health impacts and patient preferences. Nevertheless, any policy that will affect prices paid to physicians—such as unilateral price changes by government payers and mergers of insurance companies or provider groups that affect bargaining power—should consider the downstream effects on utilization.

Contents

List of Tables

List of Figures

List of Abbreviations and Symbols

Symbols

Δ	change in value
Λ	logistic function
Φ	CDF of the standard normal distribution
Π	product
Σ	sum
N	number of observations
$E(\cdot)$	expected value
$P(\cdot)$	probability
$\exp(x)$	exponential (e^x)
$\log(x)$	natural logarithm ($\ln(x)$)
$1\{exp\}$	indicator function (=1 if *exp* is true; 0 otherwise)
$\mid$	conditional on
$\in$	in

Abbreviations

AHRQ	Agency for Healthcare Research and Quality
AMA	American Medical Association
API	Application Program Interface
BAC	Blood alcohol content

BRFSS	Behavioral Risk Factor Surveillance System
CDC	Centers for Disease Control and Prevention
CFA	Consumer Federation of America
CHIA	Center for Health Information and Analysis, an agency of the Commonwealth of Massachusetts
CMS	Centers for Medicare & Medicaid Services
CPT	Current Procedural Terminology®
DOJ	Department of Justice
DWI	Driving while intoxicated or under the influence
FTC	Federal Trade Commission
HHI	Herfindahl-Hirschman index
HMO	Health Maintenance Organization
HRS	Health and Retirement Study
ICD-9	International Statistical Classification of Diseases and Related Health Problems, Ninth Revision
IOM	Institute of Medicine
MA APCD	Massachusetts All-Payer Claims Database
MAIS	Maximum Abbreviated Injury Scale
MC	Marginal cost
MEPS	Medical Expenditure Panel Survey
MSB	Marginal social benefit
NAIC	National Association of Insurance Commissioners
NCDOI	North Carolina Department of Insurance
NHTSA	National Highway Traffic Safety Administration
NIAAA	National Institute on Alcohol Abuse and Alcoholism
NPI	National Provider Identifier
OLS	Ordinary least squares

POS	Point-of-service
QALYs	Quality-adjusted life-years
RBRVS	Resource-based relative value scale
SAD	Survey of Alcohol and Driving
SIG-E-CAPS	Mnemonic for characteristics of major depression: Sleep, Interest, Guilt, Energy, Concentration, Appetite, Psychomotor, Suicidal
UM	Uninsured motorists insurance
USDHHS	U.S. Department of Health & Human Services
VBAC	Vaginal birth after Cesarean section
WTP	Willingness-to-pay

Acknowledgements

I am grateful to Frank Sloan for his dedication to my education and for always challenging me to do my best work. I am incredibly thankful for the encouragement and support from Jimmy Roberts, Peter Arcidiacono, and Donna Gilleskie. I also thank Juan Carlos Suárez Serrato, Daniel Xu, my classmates, and seminar participants for their many comments and discussions.

I gratefully acknowledge support from the Duke Social Sciences Research Institute, which funded my use of the Massachusetts All-Payer Claims Database and supported me during the 2016-17 academic year as a fellow of the Program for Advanced Research in the Social Sciences. Chapters 1 and 2 were funded in part by a grant from the National Institute on Alcohol Abuse and Alcoholism (NIAAA, #R01AA017913-01A1). The sponsor had no role in the design or conduct of these studies. I thank my Chapter 1 co-authors Frank Sloan and Lindsey Eldred, as well as Hanming Fang and Ahmed Khwaja, for their collaboration in writing the proposal that led to the grant and for developing the questionnaire used in those studies. Chapter 1 is forthcoming in the *Journal of Risk and Insurance*, published and copyrighted by John Wiley and Sons.

Thank you to the Alexanders and the Kellams for providing a home away from home; my parents for always believing in me; and Christopher Coulter for your actuarial wisdom and so much more.

1

Advantageous Selection, Moral Hazard, and Insurer Sorting on Risk in the U.S. Automobile Insurance Market

Empirical studies of insurance purchase decisions have demonstrated that consumers possess private information when purchasing, and their private information comes in multidimensional forms (e.g., Fang, Keane and Silverman, 2008). This information pertains to risk types (Rothschild and Stiglitz, 1976), but also preferences, such as risk aversion, thrill seeking, tastes for goods leading to risk taking, for example, alcohol, and in subjective beliefs about adverse consequences of risky behaviors. Individuals also differ in cognitive ability, which affects knowledge of the law and adverse consequences of risky behaviors, and, for driving, in their altruism toward other drivers, and in driving skills. While research has documented that multidimensional sources of private information exist, much remains to be learned about what these sources are and how they vary by insurance type, and how they affect functioning of insurance markets.

Even though private information is a source of adverse selection in insurance markets and the notion that purchasers of insurance possess some private information is widely accepted, not all studies have found adverse selection (e.g., Cohen and Siegelman, 2010). A frequently used method for determining whether there is adverse selection is the positive correlation test, where the correlation is between insurance quantity purchased and losses incurred during the policy year. Several studies of automobile insurance have reported a zero correlation between insurance coverage and accident risk (Chiappori and Salanié, 2000; Dionne, Gouriéroux and Vanasse, 2001; Saito, 2006).[1] A deficiency of this test is that a positive correlation can reflect both adverse selection and moral hazard. Disentangling the two is not straightforward (de Meza and Webb, 2001; Bajari, Hong and Khwaja, 2014). Also, a zero correlation may reflect multidimensional private information (Finkelstein and McGarry, 2006).

In the standard model of insurance choice, insurers are passive agents. Insurance markets persist in the presence of private information and with government requirements that certain types of information known to insurers not be explicitly used in premium setting or underwriting. Insurers must have learned how to cope with such private information. Yet research on this topic is still in its infancy (Baker and Swedloff, 2013).

This study uses unique data from a survey of persons who both drove and consumed alcohol conducted for our research in four U.S. states—North Carolina, Pennsylvania, Washington, and Wisconsin. Our survey obtained information on various sources of private information, for example, respondents' subjective probabilities about engaging in future behaviors and of experiencing adverse driving outcomes; risk preference; other preferences including altruism, income, wealth, motor vehicle ownership, and respondents' driving habits; and demographic characteristics. The survey asked for the respondent's automobile insurer and amounts of third-party and

[1] The above studies are based on North American data. In a German study, correlations were either zero or small and positive (Spindler, Winter, and Hagmayer, 2014). Positive correlations were found in studies using French (Chiappori et al., 2006) and Israeli (Cohen, 2005) data.

first-party insurance the person had, and the premium paid. Having the automobile insurer's name, we merge quality ratings obtained from several independent sources with the survey data. To our knowledge, no prior study has had as much information on insurance purchasers and choices as we do. A major strength of our survey is that it obtained objective measures of risk at the follow-up interview that are comparable to measures of subjective beliefs obtained at baseline about a year earlier, which allows for within-sample comparisons.

We find that there is advantageous selection in choice of automobile liability insurance coverage, and there is moral hazard. The advantageous selection reflects market responses on the supply side. In particular, risks are segmented based on factors predictive of the probability of an accident, both observed and used by insurers. The result is that lower-risk drivers are insured by higher-quality insurers and conversely. Thus, even with various forms of government intervention, stratifying on risk eliminates much of the heterogeneity in accident risk among drivers and segregates higher-risk drivers in their own risk pools. Such sorting reduces cross-subsidies from low- to high-risk drivers, and makes automobile insurance relatively attractive to low-risk purchasers, which is reflected in the advantageous selection we observe.

Our analysis proceeds in two stages. First, we assess consumers' private information about risk type and its role in insurance choices and accident risk. To study selection on observed dimensions of private information, we use standard techniques (e.g., Finkelstein and McGarry, 2006). Following others (Chiappori and Salanié, 2000; Finkelstein and Poterba, 2004), we measure the correlation between insurance coverage choice and ex post accident risk. We then measure moral hazard directly.[2] Second, we investigate sorting of drivers among insurers based on predictors of drivers' accident risk as a response to potential selection and moral hazard in

[2] Abbring, Chiappori and Pinquet (2003) is a notable exception—the authors use panel data to test for negative state dependence as a measure of moral hazard. They do not find evidence of moral hazard in the French automobile insurance market.

this market. Lower-risk drivers obtain coverage from insurers rated as having higher quality by independent agencies. Some insurers eschew the risky drivers on criteria observable to insurers using approaches allowable under current law and regulation.

This study makes several important contributions. First, most studies, and to our knowledge all studies of selection in automobile insurance markets, have lacked measures of private information that permit testing for whether the correlation results arise because of no asymmetric information between insurers and consumers or multidimensional private information. Ours is the first study of such markets to have direct measures of private information and is richer in private information than previous studies of selection in any insurance market.

Finkelstein and McGarry (2006) have one measure of the subjective probability of nursing home admission in their analysis of selection in the market for long-term care insurance. The study most similar to ours in having more than one private information measure is Fang, Keane and Silverman (2008), a study of the supplementary Medicare insurance (Medigap) market. They measure financial risk preference, cognitive ability, health behaviors, and subjective risk, but we have these and additional measures of private information. The authors find evidence of advantageous selection in Medigap, but risk preference is not a determinant of insurance choice as economic theory would usually predict. In their study, cognitive ability, measured as we do, drives their negative correlation between insurance coverage and claims risk.

Second, our data allow us to measure risk preference more comprehensively than previously. Risk preference may not be stable across domains, but may be context specific (e.g., Einav et al., 2012). We leverage the multiple ways in which our survey elicited preferences, personality traits, and beliefs to obtain a more general risk preference measure than in previous studies. Third, using our survey, we implement a novel method for quantifying moral hazard, which allows us to separate insurance policy choice from moral hazard. Fourth, given that we can identify respondents' insurers, we examine sorting by insurers of varying quality based on the objective risk of an accident.

The "Background on Automobile Insurance Markets" section describes key institutional features of the U.S. automobile insurance market. The "Data" section describes basic features of our survey. The "Private Information, Insurance Contract Choice, and Accidents" section focuses on the role of private information in the automobile insurance market—particularly relationships between individuals' subjective probabilities of being involved in an automobile accident in the next year and other dimensions of private information, insurance contract choice, and ex post accident risk. In the "Separating Plan Selection from Moral Hazard" section, we use information from hypothetical scenarios posed in our survey to assess the extent of moral hazard. In the "Sorting of Policyholders Based on Their Accident Risk" section, we investigate the relationship between objective and unobserved accident risk and various measures of insurer quality. The "Discussion and Conclusion" section discusses implications of our findings and study conclusions.

1.1 Background on Automobile Insurance Markets

Government intervention in the automobile insurance market reflects financial externalities of such insurance (e.g., absent third-party insurance, accident victims may not recover their losses), perceived consumer ignorance of insurance policy attributes, and the widespread notion that driving is a right, particularly since driving is often essential for employment and performing various household duties. Financial externalities have led to mandatory insurance coverage. Consumer ignorance has led to review of policy forms to assure that the terms of the contract are understandable to consumers. The notion of driving as a right has provided a rationale for premium regulation, community rating, take all comers, guaranteed renewability policies, and formation of public high-risk pools. There is some segmentation of the market by risk class as reflected by the presence of private surplus line insurers in some states.

In the United States, mandatory liability insurance coverage is nearly universal if one includes financial responsibility laws. Liability insurance is often the most practical option for drivers in states with financial responsibility laws. The requirement that drivers purchase minimum amounts of liability insurance prevents unraveling of contracts offering such minimum coverage, which has sometimes occurred in other insurance markets, for example, health insurance (Cutler and Reber, 1998). Even though most drivers are, in effect, required to carry liability insurance, some insurers may not insure high-risk drivers. In all of the four study states except North Carolina, insurers can refuse to write coverage. Of our study states, only North Carolina has a reinsurance facility. All insurers selling automobile insurance in the state must take part, and as a result, an insurer may choose to insure a high-risk driver under its regular plan or transfer it to the reinsurance facility.[3] Pennsylvania, Washington, and Wisconsin utilize an assigned-risk plan. These plans require all automobile insurers to participate. In addition, there is optional automobile insurance coverage, for example, first-party (collision) insurance. Insurance applications require basic demographic information including location, age, gender, race, and marital status. However, in some states, these are protected categories. In all study states, premiums may reflect accident risk in the location in which insurance is sold. The statutes are silent regarding insurer requests for information on driving characteristics, for example, annual mileage, vehicle specifications, or use. The driving record of an insured driver is an important information source on driver risk.

Insurers rarely if ever request some types of private information policyholders are

[3] Although empirical support is lacking, the underlying rationale is that such reinsurance gives drivers free choice of insurance carrier. It also means that individuals do not know when they are insured under the reinsurance facility as the individual's insurance contract is issued by the (primary) insurer. But with sorting on objective risk of an accident, which we argue exists below, insurance agents effectively decide on the risk levels of their clientele and hence on the shares of customers who are insured by a reinsurance facility.

likely to possess because it is impossible to verify statements made from independent sources, for example, on risk preference, use of intoxicating substances, and quality of driving. Insurers might ask about specific chronic conditions known to affect accident risk. Such information can be verified from medical records or health insurance claims.

Automobile insurers commonly use experience rating to increase or decrease premiums according to recent driving history (Lemaire, 2012). Some states regulate experience rating with a statewide point system. Of our study states, North Carolina mandates insurance increases with a state-created system. The other study states regulate surcharges, but do not require specific schedules.

1.2 Data

Battelle Memorial Institute conducted a two-wave survey of drinkers and drivers on our behalf in eight cities in four states during 2010-2012, called the Survey of Alcohol and Driving (SAD). Questionnaire design was guided by questions from prior surveys. This study relies on data from both waves. Wave 1's first part, administered by telephone, included questions on: demographic characteristics/income/wealth, alcohol consumption, accident/traffic violation history, and altruism. Wave 1's second part, administered by computer about a month later, elicited information for which visual displays are helpful or questions involving detailed scenarios (e.g., for eliciting risk preferences, willingness to pay), and details about the respondent's automobile insurance policy. Wave 2, administered by computer, was conducted about a year after Wave 1.

Eligibility for the SAD required respondents to have driven and to have consumed alcohol during the last month, to reside within a study city, and to be age 18. The recruitment process oversampled persons who consumed large amounts of alcohol in order to study drinking and driving decision making and behaviors in detail. The eight cities were: Raleigh and Hickory, North Carolina; Philadelphia and Wilkes- Barre, Pennsylvania; Seattle and Yakima, Washington; and Milwaukee and La Crosse, Wisconsin. These represent a broad geographic spread of large and small cities. Mean age of study participants is 43.8. Mean educational attainment is 15.6 years, substantially above the U.S. average. Mean household income is $81,470, also above average. Over half are female (54.6 percent), 13.1 percent are nonwhite, and 46.6 percent are married.

1.3 Private Information, Insurance Contract Choice, and Accidents

1.3.1 Overview

We first assess the role of private information on (1) individuals' choices of automobile insurance coverage, (2) their accident risks, and (3) the correlation between the two. Like Finkelstein and McGarry (2006), we divide information sources into three categories: (1) information insurers use in risk classification and thus in premium setting, (2) information insurers often elicit from consumers but specifically used in risk classification, and (3) private information unavailable to insurers. We examine private information about risk preference in financial, driving, and health domains; cognitive ability; impulsivity; altruism toward nonfamily members; and subjective beliefs about the probability the person will engage in reckless driving in the next year—speeding and drinking/driving. The goal is to determine which sorts of private information explain subjective beliefs about the probability of having an accident in the next year, observed insurance purchases, and ex post accident rates.

The system of equations is recursive. Objective information about individuals known and used or known and not used by insurers and information privately held by individuals affect objective and subjective probabilities about having an accident next year. Based on the objective risk of an accident, individuals are offered insurance policies at a premium for the next (policy) year. Individuals select an insurer and policy from that insurer based on their preferences, private information, insurance policy characteristics, and premiums. Once insurance coverage is selected, the individual forms subjective probabilities about her driving behaviors during the policy year. The individual then engages in actual behaviors that affect her accident record during the policy year. At year end, the accident record for the policy year are recorded *ex post*.

1.3.2 Three Types of Information

Type 1: Objective Factors Used by Insurers in Premium Setting. Accident determinants used by insurers for risk classification fall into these categories: demographic characteristics, driving attributes, driving history, and city. In our analysis of objective factors, demographic characteristics are: male <25, female <25, and currently married. For driving attributes, we include binary variables for whether the person reported driving 15,000+ miles/year and age, type, and use of primary vehicle—sedan, sport car, SUV, minivan, truck, or other—whether the person reports driving to work, and driving a motor vehicle regularly for a job. Driving history includes the number of: citations for speeding, arrests for DWI, and prior accidents. The SAD measured accidents at Waves 1 and 2. At Wave 1, the SAD obtained a self- report of the number of speeding citations, DWI arrests, and accidents in the prior 3 years. One year later, the SAD obtained a self-report of accidents since Wave 1. Between Waves 1 and 2, 8.9 percent of respondents reported having had an accident; 21.1 percent reported having had an accident in the past 3 years at Wave 1.

First, to determine the relationship between actual accident risk and information used by the insurer for risk classification, we estimate:

$$Prob(Accident_k = 1 | X_{t-1}) = \Phi(X'_{t-1}\alpha) \tag{1.1}$$

where *Accident* is 1 if the person was involved in an accident in the past 3 years, and X_{t-1} is a vector of characteristics used by insurers for risk classification. Alternatively, *Accident* is 1 if the person was involved in an accident in the year between Waves 1 and 2, and the number of accidents in 3 years prior to Wave 1 is included as covariate in X_{t-1}. Thus, k may be $t-1$ or t.

Table 1.1 reports results for the two dependent variables. Both regressions include covariates for car age and characteristics and city, not shown. In the first specification, young drivers, male and female, those who drive 15,000+ miles/year, and have speeding citations and DWI arrests in the prior 3 years have higher probabilities of having had an accident in the 3 years prior to Wave 1. Persons currently married are less likely to have had an accident. In the second specification, the number of accidents in the past 3 years before Wave 1 has a positive effect on the probability of having had an accident between Waves 1 and 2. Having an accident in the past 3 years before Wave 1 leads to a 0.071 increase in the probability of having an accident in the years between Waves 1 and 2—a large effect size relative to the observational mean, 0.089. The coeffcient for drives for work is also significant. Effect sizes are lower when the dependent variable is for 1 rather than 3 years; fewer coefficients are statistically significant at conventional levels in the second specification.

Type 2: Attributes Observed by Insurers but Not Used in Premium Setting. Attributes often observed by insurers but not used in risk classification are income and/or wealth, educational attainment (in years), and race/ethnicity.

Type 3: Multiple Dimensions of Private Information. The SAD elicited several dimensions of private information. The SAD used a sequence of hypothetical gambles over percent changes in lifetime income, derived from the Health and Retirement Study (HRS), to measure financial risk tolerance (Barsky et al., 1997). Based on the responses, we group respondents into three mutually exclusive categories—least,

Table 1.1: Objective probability of an accident based on characteristics used by the insurer

	Mean (Std. Dev.)		Marginal Effects (Std. Errors): Any accident in last 3 years (wave 1)		Any accident in last year (wave 2)	
	(1)		(2)		(3)	
Dependent variables						
Any accident in last 3 years (wave 1)	0.212	(0.409)			0.071***	(0.018)
Any accident in last year (wave 2)	0.089	(0.284)				
Demographic characteristics						
Male <25	0.020	(0.139)	0.155**	(0.074)	0.006	(0.059)
Female <25	0.041	(0.198)	0.102*	(0.056)	0.061*	(0.036)
Married	0.466	(0.499)	−0.061**	(0.025)	−0.016	(0.018)
Driving characteristics and history						
Miles >15k/yr	0.154	(0.362)	0.104***	(0.030)	0.034	(0.021)
Drives to work	0.759	(0.428)	0.017	(0.029)	0.008	(0.021)
Drives for work	0.321	(0.467)	0.032	(0.026)	0.046***	(0.017)
Speeding violations	0.523	(1.278)	0.027***	(0.009)	0.003	(0.006)
DWI arrests	0.039	(0.299)	0.080**	(0.040)	−0.057	(0.060)
N	1,172		1,171		1,155	
Pseudo R^2			0.06		0.07	

Notes: Marginal effects; standard errors in parentheses in columns (2) and (3). All regressions include car age and fixed effects for car type and city (not shown). All explanatory variables other than accident in last year measured at Wave 1. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

moderately, and most financially risk tolerant.

The series of (standard gamble) questions about risk taking in the medical domain in SAD is: "We want you to keep imagining that you have gotten into an auto accident that leaves you paralyzed. Suppose that doctors could cure you of the paralysis by performing an operation. Without the operation, you would be paralyzed for the rest of your life. But if the operation went well, you would be completely cured of your paralysis. If the operation did not go well, you would die immediately without any pain. Would you choose to have the operation if the chance of dying was (a randomly selected value)?" The initial probability making the person indifferent between having and not having the operation was randomized. The computer program elicited probabilities according to a prespecified formula using a randomly selected value as the starting value until the change in subjective probabilities fell below a threshold.

Impulsivity is a general term describing a tendency to act on a whim, disregarding a rational long-term strategy for maximizing personal welfare (Madden and Johnson, 2010). An impulsive individual may not consider future consequences of present actions and hence be more accident-prone. In psychology, impulsivity is an aspect of personality. In the context of subjective beliefs about having an accident, an impulsive individual may recognize that she is likely to act on a whim in the future, thus exposing her to a higher probability of personal harm (Sloan et al., 2013). A parallel is the sophisticate in the literature on self-control (Gruber and Köszegi, 2001). To measure impulsivity, the SAD incorporated questions developed by (Loewenstein et al., 2001).

Insurers know the individual's prior accident and arrest records, but the individual is likely to know more about her quality of driving and precaution levels than reflected in her driving history. To measure self-rated driving skills, the SAD asked respondents to rate their driving ability relative to others—worse, about the same, better, or much better. To measure precaution levels, Wave 1 elicited the individual's subjective probabilities of speeding 15+ miles/hour over the speeding limit and of driving at least once after having too much to drink in the next year.

We use factor analysis to create an aggregate measure of risk preference based on: preferences for risk in financial and medical domains, impulsivity, whether the individual smokes currently, and whether the individual has ever used illicit drugs or licit drugs without a prescription. The factor analysis yields two factors with eigenvalues over 1. We use the first factor, which loads heavily on impulsivity, smoking, and drug use. Higher values imply higher risk preference.

Cognitive ability may affect the individual's ability to weigh potential benefits and costs of specific choices (Fang, Keane and Silverman, 2008) and/or accuracy of a person's subjective beliefs. The SAD measured recall, memory, and numeracy based on questions from the HRS. We aggregate these measures to create an index that ranges from 0 to 16, increasing in cognitive ability.

In feeling hopelessness, depression may increase subjective beliefs about probabilities of adverse events occurring (Hepburn, Barnhofer and Williams, 2009). If depression increases subjective probabilities of an accident arising from careless driving, it may increase her precaution. But it may also lessen the individual's belief in ability to control one's fate, which may have the opposite effect. The SAD measured depression with questions from the SIG-E-CAPS, a depression screening tool (Wise and Rundell, 1994; Guck et al., 2003; Lieberman, 2003). The SAD included nine symptoms; we measure depression as a count of symptoms the respondent experienced during the year before Wave 1.

Persons who are more altruistic, especially about harming strangers, may be safer drivers and hence less accident-prone. The SAD made statements relative to altruism toward nonfamily members, none of which reference alcohol consumption or driving while intoxicated. Ceteris paribus, we expect more altruistic persons, for example, who internalize the externalities involving harm to others, to be more careful drivers. The SAD included nine statements dealing with altruism without referring to family members, each with three response options—disagree, neutral, agree. We create an altruism index by assigning each response a value of 1-3, where 3 is the most altruistic, and sum across questions.

Persons placing a lower value on avoiding injury and/or disability should be less cautious. The SAD included questions designed to value the nonpecuniary loss from an automobile accident resulting in permanent paralysis. The question design sought to avoid common pitfalls in contingent valuation research and was based on questions one of us has used previously (Sloan et al., 1998; Perreira and Sloan, 2002; Khwaja, Sloan and Wang, 2009).

To elicit maximum willingness to pay, respondents were asked to compare two areas. Area A, which has the same monthly cost of living as the place where the person currently lives. Persons living there are assumed to have a 0.01 annual probability of a person getting into an automobile accident that results in the person becoming paralyzed. Area B has a $X higher cost of living and a 0.008 probability of being involved in an automobile accident resulting in the person being paralyzed. To avoid starting-point bias, SAD assigned random starting values of $X. Based on several rounds of questions, we derive a final value for avoiding a 0.002 per year probability of becoming paralyzed. Finally, the SAD elicited subjective probabilities of speeding 15+ miles/hour over the speed limit, drinking and driving, and having an accident next year. We include subjective probabilities elicited at Wave 1 as additional sources of private information.

Table 1.2 analyzes the same dependent variables as Table 1.1, but adds explanatory variables for attributes of insured individuals known but not used by insurers (Z_{t-1}) and for private information (PI_{t-1}), all at Wave 1, to assess determinants of the probability of an accident from the insurer's perspective if there were no information asymmetries and insurers used all information currently available to them. We also estimate a more parsimonious model containing only two explanatory variables for PI_{t-1}.

$$\begin{aligned} Accident_t =& \kappa_0 + \kappa_1 X_{t-1} + \kappa_2 Z_{t-1} \\ &+ \kappa_3 SubjProbAccident_{t-1} + \kappa_4 RiskPref_{t-1} + \nu_t \end{aligned} \tag{1.2}$$

The dependent variable is a binary variable for whether the person had an accident in the year after Wave 1. The parameters κ_3 and κ_4 relate private information about risk preference and the subjective probability of an accident to actual accident occurrence next year. For both types of private information to be relevant to insurers, $\kappa_3 \neq 0$ and $\kappa_4 \neq 0$. We also expect $\kappa_1 \neq 0$; that is, objective risk as of period $t-1$

predicts actual accident occurrence in t.

The coefficients on characteristics used by insurers in Table 1.2 are very similar to their Table 1.1 counterparts. The coefficients are robust to substantial changes in specification (compare columns (2) and (3) and (4) and (5)). Adding information collected but not used by insurers and individual's private information in Table 1.2 adds 0.03–0.04 to the R^2s in Table 1.1, which implies that consumers possess some information that would improve accuracy in predicting future accident probabilities if insurers had the information and actually used it.

When the dependent variable is a binary for having an accident in the last 3 years, higher net worth drivers have a lower probability of an accident, but, *ceteris paribus*, more highly educated persons and nonwhites have a higher probability on average. Among variables for private information, coefficients on self-assessed driving ability—worse self-assessed driving ability leading to more accidents and for risk preference—more risk tolerant persons being more likely to have an accident are highly significant. A one standard deviation increase in risk preference, that is, becoming substantially more risk tolerant, would lead to a 0.044 increase in the accident probability in a 3-year period, slightly over 20 percent of the corresponding observational mean. The coefficients for the subjective probability of an accident are positive and nearly statistically significant when the dependent variable is any accident in the last 3 years ($p = 0.091$, column (2), $p = 0.069$, column (3)). When the dependent variable is a binary variable for having had an accident between Waves 1 and 2, the coefficient on the subjective probability of an accident next year, recorded at Wave 1, is estimated very imprecisely. But given the stochastic property of accidents, an accident history for a year represents much less information about the underlying quality of driving than a 3-year period does.

Table 1.2: Objective probability of an accident based on characteristics used by individual drivers

	Mean (std. dev.)		Any accident in last 3 years (wave 1)				Any accident in last year (wave 2)			
	(1)		(2)		(3)		(4)		(5)	
Dependent variables										
Any accident last 3 yrs. (wave 1)	0.212	(0.409)								
Any accident last yr. (wave 2)	0.090	(0.286)								
Demographic characteristics										
Male <25	0.020	(0.141)	0.191***	(0.074)	0.176**	(0.074)	0.051	(0.057)	0.028	(0.058)
Female <25	0.041	(0.199)	0.109*	(0.056)	0.124**	(0.055)	0.070*	(0.037)	0.075**	(0.037)
Married	0.461	(0.499)	−0.034	(0.027)	−0.036	(0.027)	0.002	(0.019)	−0.002	(0.019)
Driving characteristics and history										
Miles >15k/yr	0.154	(0.361)	0.112***	(0.031)	0.111***	(0.030)	0.044**	(0.022)	0.043**	(0.021)
Drives to work	0.761	(0.427)	0.006	(0.029)	0.004	(0.029)	0.013	(0.022)	0.011	(0.022)
Drives for work	0.325	(0.468)	0.032	(0.026)	0.030	(0.025)	0.047***	(0.018)	0.051***	(0.018)
Speeding violations	0.521	(1.283)	0.026***	(0.009)	0.027***	(0.009)	0.001	(0.006)	0.005	(0.006)
DWI arrests	0.039	(0.303)	0.068*	(0.040)	0.066*	(0.040)	−0.071	(0.070)	−0.069	(0.069)
Demographic characteristics and wealth, not used										
Net worth ($100k)	3.406	(5.989)	−0.004*	(0.002)	−0.004*	(0.002)	−0.000	(0.002)	0.000	(0.002)
Education (years)	15.616	(1.940)	0.022***	(0.007)	0.024***	(0.007)	0.002	(0.005)	0.003	(0.005)
Non-white	0.130	(0.336)	0.073**	(0.035)	0.063*	(0.034)	0.039	(0.025)	0.035	(0.024)
Private information										
Driving: worse than others	0.023	(0.150)	0.208***	(0.076)			0.083*	(0.045)		
Driving: about the same as others	0.246	(0.431)	0.048	(0.033)			−0.035	(0.025)		
Driving: better than others	0.493	(0.500)	−0.001	(0.030)			−0.003	(0.021)		
Risk preference factor	0.003	(0.997)	0.044***	(0.013)	0.047***	(0.012)	0.021**	(0.009)	0.023***	(0.009)
Cognitive ability (0-16)	14.474	(1.743)	−0.002	(0.007)			−0.002	(0.005)		
Depressed (binary)	0.182	(0.386)	−0.019	(0.032)			0.035	(0.022)		
Altruism (non-familial)	22.994	(2.623)	−0.002	(0.005)			0.005	(0.003)		
WTP to avoid paralysis ($/mo)	36.802	(33.929)	−0.000	(0.000)			0.001***	(0.000)		
Subj. prob. of speeding	0.449	(0.399)	−0.033	(0.032)			0.033	(0.023)		
Subj. prob. of drink and drive	0.165	(0.292)	0.011	(0.042)			−0.021	(0.030)		
Subj. prob. of accident	0.134	(0.145)	0.135*	(0.079)	0.142*	(0.078)	0.032	(0.057)	0.037	(0.056)
N	1,134		1,133		1,152		1,120		1,136	
Pseudo R^2			0.10		0.09		0.10		0.07	

Notes: Marginal effects; standard errors in parentheses in columns (2) through (5). All regressions include car age and fixed effects for car type and city (not shown). * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

1.3.3 Subjective Probability of an Accident in the Next Year

Specification. The subjective probability of an accident is based on (1) information used by insurers, (2) other information available to insurers but not used, and (3) private information. To examine the relationship between the subjective probability of an accident next year and specific sources of private information, conditional on objective information used and not used by insurers, we estimate:

$$SubjProbAccident_{t-1} = \beta_0 + \beta_1 X_{t-1} + \beta_2 Z_{t-1} + \beta_3 PI_{t-1} + \varepsilon_{t-1} \tag{1.3}$$

All variables come from Wave 1. Measures of X_{t-1} are the predicted probability of having an accident in the past 3 years elicited at Wave 1 based on Table 1.1 coefficients and a binary variable for actually having had an accident during the same 3 years. Measures of Z_{t-1} and PI_{t-1} are the same as in Table 1.2.

Results. Overall, there are several statistically significant and plausible relationships between the subjective probability of having an accident next year and other types of objective and private information. The subjective probability of an accident reflects the objective probability of having had an accident during the 3 years prior to Wave 1. Net worth (for Z_{t-1}) negatively affects the subjective probability. Self-assessed driving ability negatively affects the subjective probability of an accident, and the two subjective probabilities of reckless driving both have positive effects (for PI_{t-1}). Although positive, the coefficient on risk preference is not statistically significant, which suggests that although risk-tolerant persons are more accident-prone, such persons on average do not think that they have a higher accident risk. A one-unit change in the factor (approximately the standard deviation of the risk preference factor), results in less than a 0.01 change in the subjective probability of having an accident next year.

In sum, the subjective probability of an accident reflects information known and

Table 1.3: Formation of subjective beliefs about the probability of an accident in the next year

	Mean (std. dev.) (1)		Subjective prob. of accident (2)	
Dependent variables				
Subjective prob. of accident	0.135	(0.146)		
Accident probability, used				
Objective prob. of accident (wave 1)	0.214	(0.105)	0.149***	(0.043)
Any accident in last 3 years (wave 1)	0.214	(0.410)	0.019*	(0.011)
Demographic characteristics				
Net worth (\$100k)	3.415	(5.986)	−0.003***	(0.001)
Education (years)	15.619	(1.939)	−0.001	(0.002)
Non-white	0.129	(0.335)	−0.012	(0.013)
Private information				
Driving: worse than others	0.022	(0.147)	0.104***	(0.030)
Driving: about the same as others	0.248	(0.432)	0.038***	(0.012)
Driving: better than others	0.493	(0.500)	0.031***	(0.011)
Driving: much better than others	0.238	(0.426)		
Risk preference factor	0.004	(0.997)	0.008	(0.005)
Cognitive ability (0-16)	14.476	(1.743)	−0.004	(0.003)
Depressed (binary)	0.181	(0.386)	−0.019	(0.012)
Altruism (non-familial)	22.994	(2.621)	0.002	(0.002)
WTP to avoid paralysis (monthly)	36.758	(33.923)	−0.000	(0.000)
Subj. prob. of speeding	0.450	(0.399)	0.024**	(0.011)
Subj. prob. of drinking and driving	0.165	(0.292)	0.026*	(0.015)
Constant			0.090	(0.065)
N	1,135		1,135	
R^2			0.07	

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$

used by insurers in underwriting and premium setting, some other information known to insurers, and private information. The next step is to determine the extent to which private information affects insurance choice.

1.3.4 Choice of Insurance Coverage

At Wave 1, the SAD asked whether the person had liability insurance, and if so, what the liability limits were. Liability limits measure quantity of liability insurance purchased. A feature of tort law in the United States is that when the defendant's liability obligation exceeds the person's wealth, the defendant is considered to be judgment proof (Shavell, 2005), which should decrease lower-income households' demand on insurance. But as seen above, wealthier persons have lower objective and subjective accident risk, for example, exercise self-protection given their greater financial exposure (Ehrlich and Becker, 1972). Also for this reason, we expect the individual's choice of liability limits to depend on wealth. If richer individuals are less risk averse, demand for insurance would decrease, *ceteris paribus*.

To measure private information, we include covariates for risk preference and the subjective probability of having an accident next year elicited at Wave 1. Although, as seen in Table 1.3, risk preference at most has a minimal effect on subjective beliefs of an accident, risk preference is key to the individual's insurance purchase decision. We estimate:

$$LiabLimit_{t-1} = \gamma_0 + \gamma_1 X_{t-1} + \gamma_2 NetWorth_{t-1} + \gamma_3 PI_{t-1} + \eta_{t-1} \tag{1.4}$$

We base the empirical specification of Equation (4) on variables for which the theoretical case for inclusion is greatest. The dependent variable is based on the higher reported liability limit. The liability limit is set to 0 for the 39 persons who lacked liability insurance. For private information, we use the subjective probability of an accident and our risk preference measure, which is decreasing in risk aversion. For X_{t-1}, we use the predicted objective probability of an accident next year (from Equation (1)) and the binary variable for whether the individual actually had an accident during the 3 years before Wave 1.

The coefficients for X_{t-1} are jointly insignificant ($p = 0.399$), reflecting higher premiums for higher objective accident risk, and possibly sorting on objective risk (Table 1.4). The coefficient on net worth is positive and statistically significant, indicating that wealthier persons demand deeper liability coverage. Each additional \$100,000 increase in net worth leads to an \$8,100 increase in the upper liability limit, much lower than the "rule of thumb" to carry liability coverage at least equal to total

Table 1.4: Relationship between insurance coverage and subjective and objective probability of an accident

	Mean (Std. Dev.)				OLS		Probit	
	Liability limit	Liability limit | Has insurance	Has collision | Has liability	Deductible ⩽ \$500 | Has collision	Liability limit	Liability limit | Has insurance	Has collision | Has liability	Deductible ⩽ \$500 | Has collision
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variables								
Upper liability limit (\$100k)	2.943	3.067						
	(2.168)	(2.125)						
Has collision insurance			0.917					
			(0.276)					
Deductible ⩽ \$500 | Has collision				0.837				
				(0.369)				
Accident probability, used								
Obj. prob. of accident (wave 1)	0.209	0.209	0.214	0.210	−0.857	−0.972	−0.104	−0.180
	(0.099)	(0.099)	(0.104)	(0.100)	(0.694)	(0.699)	(0.075)	(0.121)
Any accident last 3 yrs. (wave 1)	0.213	0.213	0.214	0.215	0.289*	0.267	0.033	0.003
	(0.410)	(0.409)	(0.410)	(0.411)	(0.168)	(0.170)	(0.021)	(0.030)
Wealth, not used								
Net worth (\$100k)	3.738	3.833	3.509	3.821	0.081***	0.078***	0.021***	−0.001
	(6.344)	(6.396)	(5.991)	(6.350)	(0.011)	(0.011)	(0.005)	(0.002)
Private information								
Risk preference factor	−0.054	−0.093	−0.033	−0.097	−0.442***	−0.365***	−0.032***	0.009
	(0.958)	(0.925)	(0.972)	(0.915)	(0.071)	(0.074)	(0.007)	(0.014)
Subjective prob. of accident	0.130	0.129	0.135	0.134	0.106	0.146	0.090	−0.073
	(0.142)	(0.140)	(0.144)	(0.143)	(0.473)	(0.485)	(0.056)	(0.084)
N	966	927	1,130	934	966	927	1,130	934
R^2 or Pseudo R^2					0.114	0.098	0.103	0.011

Notes: Liability limit regressions exclude individuals with indeterminate liability limits. Marginal effects and associated standard errors shown in columns (7) and (8). The objective probability of an accident and the binary variable for any accident are not jointly significant in any of the regressions. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

assets.[4] The coefficient of interest is γ_3, which measures the effect of an individual's private information on insurance coverage. The classic model of unidimensional asymmetric information about risk type implies that an individual with a higher subjective probability of an accident will choose more complete insurance. With multidimensional private information, more risk-averse individuals may choose more complete insurance while having a lower accident risk.

Conditional on objective risk, we find a statistically significant negative relationship between risk preference and liability limits (Table 1.4, columns (5) and (6)) and risk preference and having collision insurance (column (7))—more risk-averse persons demand more insurance. But risk preference has no effect on whether the person has collision insurance with a deductible of \$500 or less (column (8). The subjective probability of an accident has no statistically significant influence on liability limits or on collision coverage.

In sum, private information, particularly risk preference, and net worth affect individuals' insurance choices. The importance of this finding depends on whether this private information also predicts ex post accident risk, but as we saw in Table 1.2, more risk-tolerant persons at Wave 1 have a higher probability of reporting having had an accident ex post, for example, at Wave 2.

1.3.5 Private Information and Ex Post Accident Risk: Positive Correlation Test

Together, Tables 1.2 and 1.4 suggest advantageous selection. Although we consider many possibilities, other dimensions of private information may not be captured by our measures of PI_{t-1}. For this reason, and to compare our results with previous studies, we perform a positive correlation test. The positive correlation test does not require observing private information. Yet, results of this test may be misleading if

[4] http://guides.wsj.com/personal-finance/insurance/how-much-car-insurance-do-you-need/, accessed 3/19/15.

there are offsetting sources of information that yield a zero correlation (Finkelstein and McGarry, 2006).

We use probit to estimate parameters of a model of accident occurrence—a function of insurance coverage—controlling for risk classification, similar to Finkelstein and Poterba (2004):

$$Prob(Accident_t = 1 | X_{t-1}, Insurance_{t-1}) = \Phi(\theta_1 X_{t-1} + \theta_2 Insurance_{t-1}) \tag{1.5}$$

where Insurance refers to insurance policy characteristics. With asymmetric information only about risk type, $\theta_2 > 0$, and with asymmetric information only about risk preference (unobserved by insurers), $\theta_2 < 0$. With multidimensional, offsetting private information, we may find $\theta_2 = 0$.

The results show a zero correlation between insurance coverage and accident risk (Table 1.5). The zero correlation could mask a combination of advantageous selection, a phenomenon leading to a negative correlation, and moral hazard, leading to a positive correlation (de Meza and Webb, 2001). Hence, we develop a method for separating selection from moral hazard.

1.3.6 Separating Plan Selection From Moral Hazard

Overview. We use questions from Waves 1 and 2 that elicit how individuals' subjective probabilities of reckless driving change in response to changes in financial penalties imposed by the insurer for a traffic violation. Our survey randomized financial penalties for each respondent. We use these results in combination with relationships between subjective probabilities of reckless driving and actual realizations of reckless driving, and between reckless driving and *ex post* accident risk, to quantify moral hazard. Our approach involves three steps.

Table 1.5: Relationship between accident occurrence and insurance coverage

	Mean (Std. Dev.)	Any accident in last year (wave 2)	
	(1)	(2)	(3)
Dependent variable			
Any accident in last year (wave 2)	0.089		
	(0.285)		
Insurance policy characteristics			
Upper liability limit ($100k)	2.939	0.004	0.004
	(2.168)	(0.004)	(0.004)
Has collision insurance	0.891	−0.004	0.003
	(0.312)	(0.031)	(0.031)
Accident probability, used			
Obj. prob. of accident (wave 1)	0.209		0.225***
	(0.099)		(0.085)
Any accident last 3 years (wave 1)	0.212		0.075***
	(0.409)		(0.020)
N	962	962	962
Pseudo R^2		0.00	0.05

Notes: Excludes individuals with indeterminate liability limits. The objective probability of an accident and binary variable for any accident in the last 3 years are jointly significant in col. 3, $p < 0.001$.
* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

1.3.7 Step 1. Relationship Between Expected Premium Increase and Subjective Probability of Reckless Driving

We first relate an individual's subjective probability of reckless driving, measured alternatively by the probability of speeding given his expected penalty from speeding and driving after having had too much to drink. The SAD asked respondents to estimate the (subjective) probability of speeding next year and the premium increase conditional on being convicted for speeding. The SAD then posed two scenarios in which the expected premium was increased by a randomized amount and elicited new (subjective) probabilities of speeding. Since respondents were asked to state subjective probabilities of engaging in reckless driving next year in response to randomly selected penalties presented successively at each wave, respondents may have selected a lower probability of reckless driving in response to an increase in the expected penalty posed by the survey just to give a logically plausible response. The same sequence of questions but with different randomly selected values was posed in both survey waves. To eliminate this potential source of bias, for each individual we randomly select one expected premium increase-probability of reckless driving pair from each of the two waves, based on the assumption that respondents could not remember their precise Wave 1 responses at Wave 2. We then estimate the following equation with individual fixed effects and standard errors clustered at the individual level:

$$SubjProbSpeeding_t = \phi_0 + \phi_1 E_t[\Delta Premium|Speeding] + \psi_t,\ t \in \{-1, 0\} \quad (1.6)$$

where $E_t[\Delta Premium|Speeding]$ is the expected change in premium conditional on speeding $\geqslant$15 mph. The coefficient of interest is ϕ_1, the mean individual change in the subjective probability of speeding due to a given percent increase in premiums. The null hypothesis is $\phi_1 = 0$—that is, no moral hazard.

In addition, the SAD elicited the individual's subjective belief about the expected premium increase following a DWI conviction and the probability of drinking and driving next year given the person's subjective belief about a DWI conviction. Respondents were also asked the probability of drinking and driving if the expected premium increase doubled. We randomly select one pair of responses from each wave and follow the same approach as for the speeding scenarios to estimate Equation (6) with *DrinkingAndDriving* replacing *Speeding*.

The relationship between the change in expected premiums and in the subjective probabilities can be interpreted as causal for two reasons. The first is the randomization of the hypothetical changes by SAD in expected premium increases following a conviction for speeding or drunk driving (randomized because the probability of being pulled over by police conditional on speeding or drunk driving was randomized). Second, we include individual fixed effects, which account for time-invariant heterogeneity in formation of subjective beliefs.

1.3.8 Step 2. Relationship Between Subjective Probability of Reckless Driving and Actual Reckless Driving Behavior

Step 2 measures the relationship between the subjective probability of reckless driving and realizations of reckless driving. We regress the number of times an individual reports having driven after drinking too much during the past year at Wave 2 on the subjective probability of driving after drinking too much during the next year at least once from Wave 1. We lack a measure of actual speeding behavior from Wave 2. We assume that the subjective probability of speeding is accurate; the coefficient in this step is set to 1, assuming that individuals are rational in their subjective beliefs about future speeding.

1.3.9 Step 3. Relationship Between Reckless Driving and Ex Post Accident Risk

Third, we measure the relationship between actual reckless driving behavior and the ex post probability of an accident. We use probit to estimate the probability of any accident realization as a function of whether an individual drove after drinking too much in the same year.[5] We also use probit to estimate the probability of any accident realization during the year before Wave 2, as a function of the subjective probability of speeding next year reported in Wave 1, again assuming that the subjective probability of speeding next year at Wave 1 equals the actual probability of speeding during the year before Wave 2.

[5] We pool observations from both waves and include a wave fixed effect to account for the fact that Wave 1 counted accident realizations from the past 3 years and Wave 2 counted accidents from the past year. We also cluster standard errors at the individual level.

1.3.10 Results

Descriptive statistics are shown in Appendix A (Table A.1). In Step 1, doubling the expected premium increase from a conviction for speeding leads to a 0.332 lower subjective probability of speeding 15+ mph during the next year (Table A.2, row 3, column 2), and doubling the expected premium increase from drinking and driving leads to a 0.089 lower subjective probability of drinking and driving (Table A.2, row 4, column 5). In Step 2, the change in the objective probability of reckless driving from a change in the subjective probability is assumed to be 1 for speeding and is estimated to be 0.693 for drinking and driving (Table A.3, row 1). In Step 3, the change in the objective probability of an accident due to a change in the objective probability of reckless driving is 0.046 for speeding and 0.048 for drinking and driving (Table A.3, rows 2 and 3).

The product of the estimates from these steps for speeding is 0.015, or 16.9 percent relative to the sample mean objective probability of having had an accident, and for drinking and driving it is 0.003, or 3.4 percent relative to the sample mean (Table 1.6). Moral hazard, particularly for speeding, is nontrivial.

Alternatively, in sensitivity analysis, we take 100 random draws of stated changes in speeding and drinking-driving behaviors in response to changes in premium penalties imposed for reckless driving responses to the hypothetical change in the premium for engaging in/being pulled over for/convicted of reckless driving. The estimates are robust to changes in the method used for computing the effect of engaging in reckless driving on insurance premiums (Table A.2). In addition, we perform robustness checks of Step 3 by adding/dropping binary variables for city and alternatively including individual fixed effects. The parameter estimates do not change much and do not alter our conclusions on moral hazard. Underlying the use of individual fixed effects is the assumption that other behaviors (e.g., texting while driving, drag rac-

Table 1.6: Moral hazard

Step	Speeding	Drinking & Driving
1: Δ Subj. prob. of reckless driving/Δ Penalty	−0.332***	−0.089***
2: Δ Obj. prob. of reckless driving/Δ Subj. prob. of reckless driving	1.000§	0.693***
3: Δ Obj. prob. of accident/Δ Obj. prob. of reckless driving	0.046**	0.048***
Δ Obj. prob. of accident/Δ Penalty	−0.015	−0.003
Effect of moral hazard on accident rate between Waves 1 and 2 (mean = 0.089)	−16.9%	−3.4%

Notes: Marginal effects shown. Sources of marginal effects (Appendix A): Step 1—Table A.2, row 6, column 9 and row 7, column 11; Step 2—Table A.3, row 1, column 1; Step 3—Table A.3, rows 2 and 3, column 3. Δ Penalty=100 percent increase in premium. Excludes individuals without liability insurance. §Assumed value. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

ing) correlated with speeding or drunk driving are time invariant. If these other behaviors potentially affecting accident risk are time varying and covary[6] with the behaviors observed by SAD, the Step 3 result is not interpretable as a causal effect.

1.3.11 Summary

Individuals possess private information about their driving abilities, risk preference, and beliefs about their precaution levels next year, which are systematically and plausibly related to their subjective probabilities of having an accident next year. Private information about risk preference affects demand for insurance and accident risk. Our results imply advantageous selection, particularly in level of liability coverage. But the positive moral hazard effect offsets advantageous selection, consistent with our result from the correlation test.

[6] We also include a wave fixed effect in specifications with individual fixed effects. This covariate should account for time-varying behaviors unrelated to the attributes of reckless driving we measure.

1.4 Sorting of Policyholders Based on Their Accident Risk

1.4.1 Rationale for Sorting

Our finding of advantageous selection is surprising, especially given state governments' role in regulating terms of the sale of automobile insurance. State laws and regulations prohibit insurers from using all information they possess. State laws limit underwriting. Laws mandating insurance coverage work to combat adverse selection, but such laws disproportionately attract high-risk persons willing to bet that if sued, they would be declared judgment proof.

Thus far, we have focused on the demand side. Insurers are not passive agents; even within the law they can affect outcomes through underwriting, marketing practices, and policy design. One way automobile insurers can combat potential adverse selection and its effects is through sorting, as explained below.

Rejections of applications for automobile insurance are rare. Only 2.2 percent of SAD respondents had been rejected by an insurer in the last 3 years at Wave 1; only 1.7 percent were rejected in the year before Wave 2. Very few respondents were covered by surplus-line insurers or high-risk pools. This suggests that there must be another mechanism for insuring high-risk drivers.

A mechanism for dealing with high-risk drivers is for insurers to sort drivers based on factors associated with future accident risk; drivers are accepted for some level of coverage rather than being rejected outright, but the high-risk drivers "pay" for their additional risk in part by being insured by lower-quality insurers, where quality is measured by such attributes as poorer customer service—for example, greater hassles in collecting from insurers, fewer repair firms in the insurer's network, and less access to the loaner vehicles. By specializing, lower-quality insurers become expert in gauging not-easily observed consumer characteristics. Methods for sorting include adjusting the content of advertisements (e.g., emphasis on specializing in insuring drivers with bad driving records or bad credit), office location, and simply welcoming prospective customers in the risk class in which the insurer specializes.

Our empirical analysis of sorting in the automobile insurance market involves three steps. In Step 1, we show that higher-risk drivers measured on objective characteristics used by insurers for risk classification tend to be covered by lower-quality insurers. Estimating parameters of a hedonic premium equation in Step 2 reveals that premiums fall with increases in consumer ratings of insurers (excluding consumer satisfaction with premiums). Absent some sort of non-price rationing, one would expect price and quality to be positively related. In Step 3, we investigate whether high-risk drivers possess more private information about their accident risk. If so, this could explain why premiums paid by the high-risk group are disproportionally high in Step 2.

1.4.2 Choice of Insurer

We obtain data on insurers from three additional sources: J.D. Power, Insure.com, and Consumer Reports. Each source reports consumer ratings on multiple dimensions of quality for large automobile insurance companies. Insure.com also reports coverage options and discounts offered by each insurer included in its study. Consumer ratings reported in these studies cover 83 percent of our sample for J.D. Power ratings and 77 percent of our sample for Consumer Reports and Insure.com ratings. Consumer Reports publishes automobile insurance ratings based on a survey of 102,207 subscribers. The "reader score" measures overall satisfaction with claims handling based on the subsample of 29,116 survey respondents with a claim, 2009 to mid-2012. Scores range from 0 to 100, where 100 indicates all respondents are completely satisfied; 80, very satisfied; 60, fairly well satisfied; and 40, somewhat dissatisfied on average. "Premium satisfaction" is on a scale 1- 5, increasing in satisfaction. Insure.com published ratings of 20 large automobile insurance companies based on a survey of 5,600 insurance customers conducted in 2012. The "overall score," ranging from 0 to 100, is an aggregate measure of satisfaction with: claims processing, customer service, premium paid given the coverage, the percent of respondents who would renew their coverage ("plan to renew"), and the percent of respondents who would recommend or already recommended the insurer. Insure.com also reports coverage options, including whether or not the company offers accident forgiveness ("accident forgiveness offered"), and a distribution of reasons respondents bought from the company (e.g., "saw commercial"). J.D. Power ratings include: (1) "overall claims satisfaction"; (2) "overall purchase experience"; (3) "claims service interaction," based on claimants' ratings of the insurer representative or agent handling the claim; and (4) "local agent interaction," reflecting purchasers' experiences interacting with the insurer's local agent or staff. All ratings vary from 2 to 5, with 5 among the best, 4 better than most, 3 average, and 2, the rest.

Given that private information affects individuals' choice of level of insurance coverage, a question remains whether private information affects individuals' choice of insurer based on other dimensions of insurance contracts such as coverage options, discounts, and insurer quality scores. We estimate:

$$Attribute_{t-1} = \pi_0 + \pi_1 X_{t-1} + \pi_2 PI_{t-1} + \omega_{t-1} \tag{1.7}$$

where Attribute alternatively measures consumers' ratings, policy options, or discounts reported in Consumer Reports, Insure.com, or J.D. Power for the SAD respondent's insurer. If $\pi_1 \neq 0$, then the observable accident probability, measured by the predicted objective accident probability in the last 3 years before Wave 1 and a binary for having actually having had an accident during this period, is related to the attribute represented by the dependent variable, which suggests sorting by accident risk along that dimension of insurer quality. If $\pi_2 \neq 0$, then individuals' residual private information also affects choice of insurer.

Table 1.7: Private information and quality of insurer

	Reader Score		Overall Score		Overall Claims Satisfaction		Overall Purchase Experience		Plan To Renew	
Source:	*Consumer Reports*		*Insure.com*		*J.D. Power*		*J.D. Power*		*Insure.com*	
Panel A: Overall Quality	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Obj. prob. of accident		−1.814***		−1.975		−0.319		−0.661**		−5.857***
		(0.674)		(1.229)		(0.207)		(0.315)		(1.827)
Any accident last 3 yrs.		0.021		0.096		−0.012		−0.048		0.288
		(0.174)		(0.318)		(0.054)		(0.081)		(0.472)
Risk preference factor	−0.117	−0.100	0.136	0.151	−0.030	−0.026	−0.037	−0.028	0.141	0.184
	(0.076)	(0.076)	(0.135)	(0.137)	(0.023)	(0.023)	(0.035)	(0.035)	(0.202)	(0.203)
Subj. prob. of accident	0.505	0.676	0.209	0.358	0.100	0.135	0.114	0.191	−0.591	−0.150
	(0.466)	(0.469)	(0.848)	(0.857)	(0.142)	(0.144)	(0.218)	(0.220)	(1.267)	(1.274)
N	792	792	786	786	801	801	819	819	786	786
R^2	0.02	0.03	0.01	0.01	0.01	0.01	0.02	0.03	0.01	0.03
Joint test significance level[a]		0.006		0.117		0.100		0.021		0.002

	Premium Satisfaction		Claims Service Interaction		Local Agent Interaction		Offers Accident Forgiveness		Reason for Buying: Saw Commercial	
Source:	*Consumer Reports*		*J.D. Power*		*J.D. Power*		*Insure.com*		*Insure.com*	
Panel B: Specific Attributes	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Obj. prob. of accident		−0.480**		−0.598		−0.561		0.341**		2.189*
		(0.191)		(0.407)		(0.382)		(0.159)		(1.254)
Any accident last 3 yrs.		−0.003		−0.035		−0.046		0.025		0.445
		(0.049)		(0.105)		(0.099)		(0.041)		(0.324)
Risk preference factor	−0.041*	−0.036*	−0.062	−0.055	−0.040	−0.032	0.025	0.021	0.267*	0.221
	(0.021)	(0.022)	(0.045)	(0.046)	(0.042)	(0.042)	(0.018)	(0.018)	(0.139)	(0.139)
Subj. prob. of accident	0.225*	0.272**	−0.119	−0.051	−0.228	−0.164	0.130	0.090	−0.416	−0.745
	(0.132)	(0.133)	(0.280)	(0.283)	(0.264)	(0.267)	(0.109)	(0.110)	(0.867)	(0.875)
N	792	792	801	801	794	794	820	820	786	786
R^2	0.02	0.03	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.03
Joint test significance level[a]		0.009		0.110		0.103		0.019		0.031

Table 1.7: Private information and quality of insurer

Source:	Overall Quality *Aggregate Measure*	
Panel C: Aggregate Quality	(1)	(2)
Obj. prob. of accident		−1.814***
		(0.674)
Any accident last 3 yrs.		0.021
		(0.174)
Risk preference factor	−0.117	−0.100
	(0.076)	(0.076)
Subj. prob. of accident	0.505	0.676
	(0.466)	(0.469)
N	792	792
R^2	0.02	0.03
Joint test significance level[a]		0.006

Notes: Standard errors in parentheses. All specifications also include income, educational attainment, and race/ethnicity not shown.
[a]Joint test is for combined effect of objective probability of accident and any accident, both in the 3 years before Wave 1.
* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

An individual with a higher predicted objective probability of an accident is more likely to have a low-quality insurer, as measured by summary indicators of quality (Table 1.7, Panel A, columns 2, 8, and 10). The coefficients on the binary for actual accidents vary in sign and are uniformly smaller than their associated standard errors, but the coefficients on the two components of X_{t-1} are jointly significant for the reader score, overall purchase experience, and plan to renew—in all these cases implying the persons with higher accident risk tend to obtain coverage from lower-quality insurers.

Policyholders with a higher objective accident probability and a prior accident record are less likely to be satisfied with their premiums (Panel B, column 2 of Table 1.7), which is plausible since such persons are likely to pay higher premiums on average. But the effect is partly offset by higher satisfaction with premiums among policyholders with private information that they are more likely to be involved in an accident. The negative coefficient on the risk preference factor implies that premium satisfaction decreases as risk tolerance increases.

Persons with high objective accident probabilities are much more likely to be influenced by insurers' advertisements (Panel B, column 10 of Table 1.7). Not only do persons with a higher objective risk face higher premiums, which would encourage search, but they are also likely to be less welcome by insurers when they seek to purchase coverage. Friends might be reluctant to encourage an accident-prone friend to buy coverage from their insurer. As expected, accident-prone persons are more likely to be satisfied with accident forgiveness (Panel B, column 8) because they are more likely to have had their accidents forgiven. In Panel C, the dependent variable is a consumer satisfaction index reflecting all of the satisfaction measures available for the respondent's insurer excluding premium satisfaction. Like Panel A, persons with a higher objective accident probability and with a prior accident have lower assessments of overall insurer quality. Neither private information covariate has a statistically significant impact.

Taken together, these findings suggest that more accident-prone persons, measured by factors that insurers consider, are less satisfied with their insurers, even on dimensions other than the premium. Proportions of high- and low-risk drivers are systematically related to the quality of the insurer.

1.4.3 The Effect of Insurer Quality on Premiums Paid

Ceteris paribus, maximum willingness to pay for a good or service should be increasing in the good or service's quality. Failure to find this relationship empirically may reflect some type of nonprice rationing, in this context sorting by insurers based on the person's objective probability of an accident. We estimate:

$$\begin{aligned} Premium_{t-1} = &\tau_0 + \tau_1 X_{t-1} \\ &+ \tau_2 InsCoverage_{t-1} + \tau_3 InsQuality_{t-1} + \xi_{t-1} \end{aligned} \tag{1.8}$$

In Equation (8), the annual premium paid by the SAD respondent for the household's automobile insurance policy depends on the objective accident risk, insurance contract parameters—the upper liability limit (in \$100,000s), a binary variable for whether the respondent had collision coverage, and insurer quality, the same measure as in Table 1.7, Panel C; insurer quality and other factors. All specifications use a Heckman correction to account for sample selection, that is, for whether or not the respondent had liability insurance.

In Table 1.8, column 1, based on the full sample, as expected, premiums rise with increases in the predicted objective probability of an accident and the binary variable for an accident in the 3 years before Wave 1. Premiums are higher for persons with collision coverage and higher liability limits, although the latter is not statistically significant. Higher-quality insurers have lower premiums, evidence for sorting. Results for other covariates are generally plausible but not shown.

Table 1.8: Premiums and quality of insurer

	Premium				
	Full sample	Low Risk/No Accident	Medium Risk/No Accident	High Risk	Low or Medium Risk/Accident
	(1)	(2)	(3)	(4)	(5)
Accident probability, used					
Objective prob. of accident (wave 1)	682.895**				
	(345.904)				
Any accident in last 3 years (wave 1)	47.262				
	(80.775)				
Insurance contract parameters					
Upper liability limit ($100k)	20.994	21.567	−7.371	51.866**	20.098
	(15.549)	(27.953)	(36.672)	(23.495)	(34.024)
Has collision insurance	431.843***	348.713	611.180**	428.229**	441.939
	(138.340)	(269.442)	(287.864)	(208.295)	(378.079)
Insurer quality	−102.236**	−161.252**	2.549	−104.468	−209.367**
	(45.361)	(81.903)	(105.805)	(68.731)	(95.038)
N	760	238	198	237	87

Notes: The dependent variable is the respondent's reported annual premium. All specifications use a Heckman correction to account for selection into insurance, where the selection equation includes chargeable accidents, DWI arrests, cars in household, miles driven per week, drives to work, and risk preference, not shown. The premium equation also includes the number of adult (age$\geqslant$25) and young (age$<$25) drivers on the policy, the number of cars in the household, and the collision insurance deductible. Standard errors in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Next, we stratify the sample based on terciles of the predicted objective probability of an accident. We remove observations with an accident in the 3 years before Wave 1 if they are in the low-risk or medium-risk tercile and place them in a fourth group—low or medium risk and accident. The fourth group contains relatively few individuals. The high-risk group contains individuals with high predicted accident risk and an accident if an accident was reported.

Two results are particularly noteworthy. First, persons in the high-risk group pay high premiums for higher liability limits. The parameter estimates for upper liability limit for the other three groups are not statistically significant. Second, the coefficients on insurer quality are negative for most groups, but only statistically significant in the low-risk/no-accident and low- or medium-risk/accident analyses. The latter result implies that lower-risk drivers do not pay higher premiums for having had an accident in the past 3 years.

1.4.4 Private Information by Risk Type

Thus far, we have seen that persons with higher objective probabilities of an accident tend to purchase automobile insurance from firms assessed by consumers to offer poorer quality of service. Premiums obtained by firms with lower assessed quality charge higher, not lower premiums. Both findings suggest sorting. Hendren (2013), using data from other types of insurance, finds that high-risk persons possess more private information than low-risk types do. If so, high-risk drivers based on factors observed and used by insurers are less desirable customers. Thus, insurers must be compensated for this in some way.

As a final step, we explore a possible reason for insurer sorting based on accident risk. We hypothesize that high-risk persons, based on their objective probabilities of an accident and actual accident histories, possess more private information about their actual riskiness than others. To gauge the importance of private information for persons with different objective accident risk, we reestimate Equation (3) with the sample split into the same risk groups as in the previous section. The dependent variable is a binary for whether or not the respondent had an accident between Waves 1 and 2.

Examining mean values by risk group (Table 1.9), the mean fractions of persons with an accident in the follow-up year rises monotonically from the low-risk/no-accident group to the low- or medium-risk/accident group. The fractions of persons with an accident during follow-up for the high-risk and low- or medium-risk/accident group are statistically different from the low-risk/no-accident group. Risk tolerance rises monotonically from the low-risk/no-accident to the low- or medium-risk/ accident group. The subjective probability of an accident is statistically higher for the high-risk and low- or medium-risk accident group than for the low-risk/no-accident group. Persons in the low- or medium-risk/accident group are far more risk tolerant than those in the low- and medium-risk groups who did not have an accident in the 3 years prior to Wave 1.

In regression analysis with the dependent variable, a binary variable for whether or not the person had an accident in the year before Wave 2, the only statistically significant result for the high-risk group is for the subjective probability of an accident, where the marginal effect is 0.258. Among the other measures of private information, the marginal effect of risk preference is 0.040 and statistically significant in the regression for the low-risk/no-accident group. Since the marginal effect applies to a factor for which units have no natural meaning, consider that the difference between the means on risk preference between low-risk/no-accident and the high- risk group is 0.272. Multiplying this mean value by the low-risk/no-accident's marginal effect of risk preference yields 0.011, which is substantially below the marginal effect on accident risk from the subjective probability of an accident for the high-risk group, supporting the view that private information is relatively important among high-

Table 1.9: Private information by objective risk groups

	Means (Std. Dev.)				Any Accident in Last Year (Wave 2)			
	Low Risk/No Accident	Medium Risk/No Accident	High Risk	Low-Med Risk/ Accident	Low Risk/No Accident	Medium Risk/No Accident	High Risk	Low-Med Risk/ Accident
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any accident in last year (wave 2)	0.049	0.071	0.119**	0.152				
	(0.217)	(0.257)	(0.324)	(0.360)				
Objective prob. of accident (wave 1)	0.121	0.193***	0.327***	0.170***				
	(0.030)	(0.020)	(0.096)	(0.039)				
Any accident in last 3 years (wave 1)	0.000	0.000	0.318***	1.000***				
	(0.000)	(0.000)	(0.466)	(0.000)				
Net worth ($100k)	4.021	3.011**	3.456	2.858	0.002	−0.002	−0.001	0.003
	(5.672)	(5.469)	(6.916)	(4.436)	(0.002)	(0.003)	(0.003)	(0.007)
Risk preference factor	−0.169	−0.005**	0.103	0.131	0.040***	−0.016	0.023	0.046
	(0.834)	(0.982)	(1.076)	(1.118)	(0.013)	(0.017)	(0.014)	(0.029)
Subjective prob. of accident	0.117	0.120	0.160***	0.143	−0.077	−0.489**	0.258***	0.013
	(0.130)	(0.137)	(0.156)	(0.153)	(0.090)	(0.237)	(0.096)	(0.232)
N	344	310	387	125	337	305	381	119
Pseudo R^2					0.10	0.06	0.04	0.02

Notes: Marginal effects and associated standard errors from probit in columns (5)–(8). Significance tests in columns (2)–(4) are pairwise t-tests. The reference group is the low-risk/no-accident group. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

objective-risk consumers,[7] which in turn should be reflected in premiums. However, the coefficient on the subjective probability of an accident for the medium-risk/no-accident group is negative, implying that persons in this group were overly pessimistic about their future accident experiences.

Thus, overall from the evidence obtained from Step 3, we cannot conclude that high- risk persons possess more private information about their accident risk. Perhaps the results on premiums reflect insurers' beliefs about private information, which have not been sufficiently validated empirically. Or one year's experience is insufficient to measure individual driving quality.

1.5 Discussion and Conclusion

Our empirical analysis yields these major findings. (1) Information insurers routinely collect on their customers predicts the objective probability of an accident during the policy year. (2) The subjective probability of an accident in the following year reflects the risk of an accident based on factors observed and used by insurers, but it also contains additional information not used by insurers in premium setting and underwriting. (3) Risk preference is an important source of private information. However, given that risk preference is multidimensional and may be domain specific, the overall effect of risk preferences on demand for insurance is an empirical question. Risk preference affects demand for insurance through the traditional channel as a means of consumption smoothing, but because risk preference also affects driving and other behaviors, risk-tolerant persons may demand more, not less, insurance. In our empirical analysis, risk-tolerant persons have lower demand for liability insurance even after taking account of other dimensions of risk preference, providing evidence of advantageous selection. (4) The combination of advantageous selection and moral

[7] For Tercile 2, the coefficient on the subjective probability of an accident is an implausible -0.210, which may partially reflect multicollinearity with the objective probability of an accident. Note that the coefficient on the objective probability of an accident is implausibly high.

hazard is consistent with the zero correlation between insurance choices and ex post accident risk we obtain. (5) There is evidence of sorting of policyholders across insurers based on factors insurers use in underwriting/premium setting. By sorting based on objective probabilities of an accident, people find themselves in pools with persons of similar accident risk. Since such sorting reduces premiums to low-risk drivers, they are understandably motivated to purchase deeper coverage.

According to the old adage, "If it ain't broke, don't fix it." In the U.S. automobile insurance market, institutional responses by private firms have done the fixing.

2

Information in Insurance Markets: When Less is More

Asymmetric information is a widely recognized potential failure in insurance markets. Although many studies have shown varying degrees of selection resulting from asymmetric information in insurance markets, there is no clear consensus about how to best deal with this market failure. Typical government interventions, including mandatory insurance, subsidized insurance, and premium regulations, can have theoretically ambiguous effects on welfare (Einav and Finkelstein, 2011). Some empirical studies support the view that risk-rating—whereby insurers are allowed to adjust premiums based on certain individual risk factors—improves welfare in an adversely selected market (e.g., Bundorf, Levin and Mahoney, 2012; Einav et al., 2013). However, recent work suggests that other markets are advantageously selected: individuals who purchase insurance are lower-risk than non-purchasers. For example, Finkelstein and McGarry (2006) and Fang, Keane and Silverman (2008) find evidence of advantageous selection in long-term care insurance and Medigap. In this article, I first provide new evidence of advantageous selection in the U.S. automobile insurance market. I then use this setting to ask how the effects of government interventions—specifically premium regulations—differ in an advantageously-selected insurance market.

This article argues that for an advantageously-selected insurance market, using more information to set premiums is not necessarily better. In the U.S. automobile insurance market, allowing insurers to adjust premiums based on an additional dimension of individual risk—propensity to drink and drive—causes high risk-drivers, who tend to have low levels of risk aversion, to drop coverage. If these individuals continue to drive but do not have sufficient wealth to compensate victims for damages in an accident, then they impose a negative externality on other drivers. This externality offsets the welfare gains from reducing asymmetric information. Community rating, on the other hand, reduces the share of high-risk, uninsured drivers without causing low-risk drivers, who tend to be highly risk-averse, to drop coverage, leading to an overall increase in welfare despite increasing asymmetric information—an illustration of the theory of second best (Lipsey and Lancaster, 1956).

Theoretically, the amount of regulation that achieves the socially optimal level of insurance coverage depends on the type and degree of selection in the market and the size of the externality from insurance. In the canonical model of selection in insurance markets, advantageous selection leads to overinsurance in equilibrium (e.g., de Meza and Webb, 2001). In this case, risk-rating can improve welfare by reducing the welfare loss due to overinsurance. However, this ignores the benefit to other drivers from being insured. I show that accounting for this externality can lead to an equilibrium characterized by underinsurance rather than overinsurance. In this case, risk-rating can cause further underinsurance. On the other hand, restricting insurers from using information about individual risk in premium-setting, as through community rating, may increase welfare by allowing high-risk drivers, who previously drove without insurance, to purchase insurance at a lower price.

To test these predictions and quantify the welfare effects of premium regulation, I develop a structural model of insurance choice, reckless driving behavior, and ex post accidents and driving violations. The model is sequential: individuals first choose an insurance policy from a set of discrete alternatives. Then they choose whether or not to drink and drive, conditional on the chosen policy and their subjective probabilities of an accident and of being convicted for driving while intoxicated (DWI). To model the future value term, I make the key assumption that individuals act in accordance with these subjective probabilities. This is a divergence from the assumption in most choice models that people act in accordance with objective probabilities, which requires that people not only act on, but also know the objective probabilities. Using subjective probabilities, on the other hand, is consistent with individuals' perceptions of probabilities and only requires the assumption that they act on these perceptions (Zafar, 2011).

I estimate the model using panel data from a survey designed specifically for this study to elicit information about individuals' reckless driving behavior, risk preference, and other variables typically unobserved by researchers. The data also include information about insurance coverage, accidents, and other violations. Administrative data have the advantage of a larger sample size and objective accident reports. However, the survey data are extremely valuable for two reasons: (1) they include measures of risk preference and subjective probabilities, which otherwise must be deduced under strong modeling assumptions; and (2) they contain self-reports of reckless driving behavior, such as drinking and driving, even when the individual was not caught. Observing both risk preference and reckless driving behavior allows me to test an idea that is pervasive in theory but has little empirical evidence: risk-averse consumers are more likely to buy insurance and more likely to take precautionary behavior, making them lower risk and inducing advantageous selection.

The estimated model parameters establish that the automobile insurance market is advantageously selected: risk-averse drivers purchase more complete insurance coverage and are less likely to drink and drive. This finding contrasts with earlier studies of automobile insurance that found no evidence of selection (e.g., Chiappori and Salanié, 2000) or some evidence of adverse selection (e.g., Cohen, 2005; Chiappori et al., 2006). However, these studies relied on variants of the positive correlation test, which Finkelstein and McGarry (2006) showed can be misleading in the presence of multi-dimensional private information. My results suggest that this is the case in the automobile insurance market.

Using the estimated parameters, I simulate the effects of three insurance pricing regimes: risk-rating, community rating, and experience rating. Under risk-rating, the insurer adjusts premiums for individual-level risk factors, which causes high-risk individuals to pay more for the same insurance coverage. Under community rating, the insurer must offer all individuals the same premium for the same coverage. Under experience rating, the insurer updates an individual's premium for the next contract period based on new information from the most recent policy period, such as accidents or DWI arrests. Experience rating is thus a dynamic form of risk-rating, as well as a punitive measure to deter reckless driving.

I find that these three insurance pricing policies each have very different effects on market outcomes and consumer welfare. Risk-rating on drinking and driving behavior increases the share of uninsured drivers by 7%, which offsets any welfare gain from reducing asymmetric information after accounting for the externality of additional uninsured motorists. This contrasts with community rating, which induces previously uninsured, high-risk drivers to buy insurance, whereas many low-risk drivers retain coverage, despite their large premium increases, because they are risk averse. Experience rating has no deterrent effect on drinking and driving, which suggests that moral hazard is not manifested through this reckless behavior. Rather, the model parameters show that individual preferences—for drinking and avoiding non-pecuniary loss—are the key determinants of drinking and driving behavior.

Drinking and driving is a particularly relevant reckless driving behavior to study because it is a source of asymmetric information in automobile insurance that can have devastating consequences: Levitt and Porter (2001) estimate that drinker-drivers are 7-13 times more likely to cause a fatal crash, depending on their blood-alcohol content (BAC). Yet, insurers only know whether an individual has been drinking and driving if he is convicted for DWI, and such convictions are rare relative to the prevalence of drinking and driving—Levitt and Porter (2001) estimate that nearly 20% of drivers between 8pm and 5am are under the influence of alcohol, but the probability of being arrested for DWI conditional on driving with a BAC above 0.1% is 0.005 (Beitel, Sharp and Glauz, 2000).

Section 2.1 provides background on automobile insurance coverages and premiums. Section 2.2 uses a graphical framework to illustrate the potential consequences of premium regulation. Section 2.3 lays out a theoretical model of insurance choice, reckless driving behavior, and *ex post* risk, which I estimate using the panel data and empirical strategy described in Sections 2.4 and 2.5. Section 2.6 presents the estimation results, establishing that risk-averse drivers buy more complete insurance and are less likely to drink and drive. Section 2.7 uses counterfactual simulations to show that restricting the information used to set premiums can improve welfare in an advantageously selected market.

2.1 Auto Insurance Coverage and Premium-Setting

To model demand for automobile insurance, we measure its value to consumers: how much it reduces exposure to financial loss from an accident. There are four potential sources of financial loss: bodily injury to others, property damage to others, own property damage, and injuries to oneself. The actual cost that the individual faces for each of these losses may differ from the total cost for two reasons: (1) insurance covers a portion or all of the cost, or (2) the amount an individual can be forced to pay is limited by his net worth (Mahoney, 2015).

Liability insurance protects drivers from the largest financial risk: injury and property damage to others. The insurance policy specifies upper limits on the amount the insurer will pay for such losses. The insured driver is responsible—and can be sued—for any damages above that amount. These "liability limits" are typically quoted in three numbers, e.g., 30/60/25, which represent the maximum dollar amounts, in thousands, the insurer will pay for injuries per person, injuries per accident, and property damage per accident. States regulate the minimum liability limits on an insurance policy. The state minimums range from $20,000 to $100,000 for bodily injury and $5,000 to $25,000 for property damage (Insurance Information Institute, 2011).

Liability insurance is compulsory, yet an estimated 14% of drivers are uninsured nationwide (Insurance Information Institute, 2011), exposing them to substantial financial risk and legal penalties.[1] The legal penalties for driving without insurance vary by state but include fines, license suspension, probation, and imprisonment.[2] An uninsured driver also imposes a negative financial externality on other drivers if he is unable to pay for damages in an accident he causes. Hence, insurers offer uninsured motorists coverage, which protects the insured driver in an accident caused by an uninsured driver. Some states require that all policies include this coverage.

Another reason for a driver to purchase automobile insurance is to insure against damage to himself and his vehicle when he causes an accident. Loss to the individual's vehicle is covered by collision insurance, which is optional. Drivers who do buy

[1] The estimated percentages of uninsured drivers in the study states, NC, PA, WA, and WI, are 12%, 7%, 16%, and 15%, respectively, which are based on the ratio of uninsured motorists claims to bodily injury claims in 2007 (Insurance Information Institute, 2011).

[2] As of January 2014, the maximum penalties for a first offense were: NC: $100 in fines, 30 day suspension of registration, with proof of financial responsibility required for restoration, and 1-45 days probation; PA: $350 in fines and 3-month suspension of license and registration, with proof of financial responsibility required for restoration; WA: $287 in fines or community restitution; WI: $500 fine (Consumer Federation of America, 2014)

collision insurance choose a deductible (e.g., \$500), which is the amount the driver must pay before the insurer pays the rest. A driver without collision insurance is exposed to loss up to the value of his vehicle. The medical costs of injuries a driver causes to himself are typically covered by health insurance, so drivers who have health insurance have little reason to purchase the additional coverage insurers offer to cover such injuries.

Unlike health insurance, automobile insurers are allowed to use a rich set of risk factors to set premiums, which reduces the asymmetry of information about risk in this market. Of course, insurers can only use information that they can reliably observe. Observable risk factors of individuals include age, gender, past accidents, speeding violations, DWI convictions, miles driven, and driving uses, though the latter two are not easily verifiable. Household-level risk factors insurers use are the number of cars and the area where they are driven—which affect the probability of an accident—and the car models—which affect expected cost if the policy includes collision coverage.

Some states use insurance market premiums as an instrument for public policy interventions. For example, North Carolina has a "Safe Driver Incentive Plan" that dictates how much insurers are required to increase premiums in response to driving violations and criminal offenses (North Carolina Department of Insurance, 2010). In theory this system should allay moral hazard: the threat of punishment through higher premiums is an incentive for drivers to be cautious.

2.2 Theoretical model of advantageous selection

This section presents a graphical framework to illustrate the welfare effects of pricing regulation in the automobile insurance market.

In the standard model of advantageous selection in a competitive insurance market, there is over-insurance in equilibrium (see, e.g., de Meza and Webb, 2001).

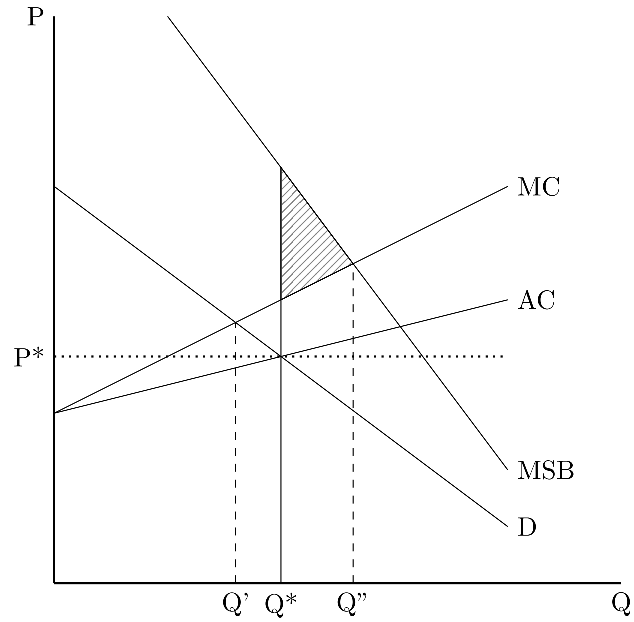

FIGURE 2.1: Welfare loss under advantageous selection with a positive externality

Low-risk drivers have a strong preference for insurance, which reduces the mean price of insurance and allows high-risk drivers to purchase insurance at a price below marginal cost. This can be represented graphically by an upward-sloping marginal cost curve, where price is on the vertical axis and the number of insured persons is on the horizontal axis (Figure 2.1).[3] The average cost curve then lies below marginal cost. Marginal cost can exceed the marginal benefit of insurance because insurers markup premiums above the actuarial value. As a result, the equilibrium level of insurance, determined by the intersection of the demand and average cost curves (a zero-profit condition on insurers), occurs at a level where marginal cost exceeds marginal benefit, which is deemed inefficient, over-provision of insurance (Q*-Q' in Figure 2.1).

However, the standard model ignores the social benefit of insuring high-risk

[3] See also Einav and Finkelstein (2011), who also illustrate the results for an adversely selected market, and Hackmann, Kolstad and Kowalski (2015), who use a similar theoretical framework to study how an insurance mandate affects equilibrium insurance coverage and welfare in the adversely-selected individual health insurance market.

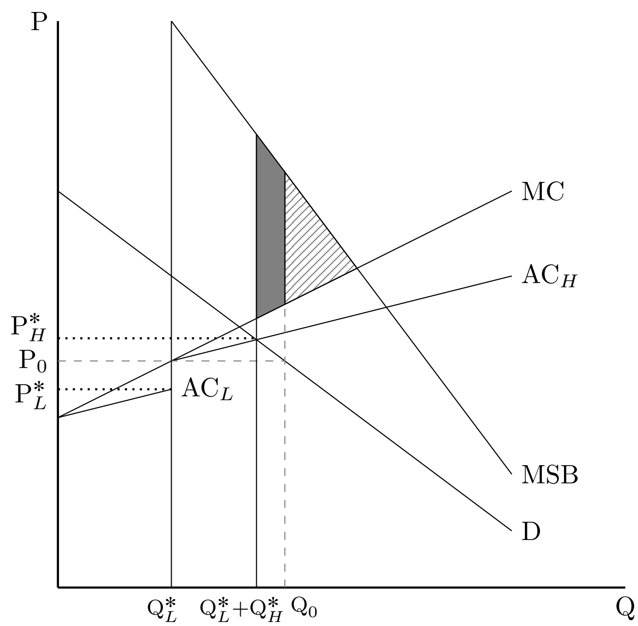

FIGURE 2.2: Impact of risk-rating on welfare under advantageous selection

drivers—when the social benefit of insurance exceeds the private benefit, an advantageously selected market can exhibit under-insurance. The equilibrium quantity of insured drivers, Q*, is still determined by the intersection of private demand and average cost, but at this quantity marginal cost can lie below marginal social benefit. The resulting welfare loss is given by the shaded, triangle-shaped area between the MC and MSB curves, to the right of the equilibrium quantity Q* and to the left of the optimal quantity Q". Of course, it is possible that the inefficiencies arising from advantageous selection and from the positive externality perfectly offset each other so that the market attains the optimal level of insurance in equilibrium. This would happen if the MSB curve intersected MC at Q*. Similarly, the externality may not be large enough to generate under-insurance, in which case the MSB curve would intersect MC at a level below Q*. Thus, whether the market exhibits over- or under-insurance is ultimately an empirical question.

Risk-rating in an advantageously-selected market may reduce the welfare loss due to asymmetric information, but it increases the welfare loss due to the externality of insurance (Figure 2.2). To see this, suppose insurers learn new information that allows them to perfectly distinguish between drivers with a marginal cost below the equilibrium price and those with a marginal cost above the equilibrium price P_0. The insurer can now offer insurance to low-risk drivers at a lower price, P_L^*, and to high-risk drivers at a higher price, P_H^*. At these prices, the quantity of low-risk, insured drivers, Q_L^*, does not change; these drivers bought insurance at a price P_0 so they still buy at the lower price P_L^*. However, high-risk drivers face a price increase, which causes the quantity of insured, high-risk drivers to decrease from $Q_0 - Q_L^*$ to Q_H^*. The total quantity of insured drivers decreases from Q_0 to $Q_L^* + Q_H^*$, causing the loss in social welfare due to underinsurance to increase to include the dark shaded area.[4,5]

Risk-rating also affects the distribution of welfare among consumers—the surplus to low-risk drivers increases and the surplus to high-risk drivers decreases. The change in consumer surplus for low-risk drivers is given by the area below P_0 and above P_L^* between 0 and Q_L^*. The change in consumer surplus for high-risk drivers is given by the area below P_H^* and D, above P_0, and between Q_L^* and Q_0.

To quantify these changes in welfare, I first estimate a model of insurance pricing, demand, and accident risk. I use the estimated model parameters to calculate the compensating variation to consumers under counterfactual pricing scenarios.[6]

[4] If risk-rating is imperfect, so insurers cannot cleanly separate consumers at a certain level of risk, Q, then the welfare analysis is less straightforward—we must separate the MC and demand curves for each risk category, consider each its own market, and sum across markets to obtain the overall welfare effect. Einav and Finkelstein (2011) give a numerical example to show how the welfare effect of risk-rating is ambiguous in this case. This occurs if risk rating is based on a variable that is not highly correlated with marginal cost. In the empirical analysis, I focus on a variable shown in previous research to be highly correlated with severe accidents: propensity to drink and drive (e.g., Levitt and Porter, 2001).

[5] Using a downward-sloping MC curve, it is straightforward to show that in adversely-selected market, perfect risk-rating improves welfare both with and without a positive externality by increasing the quantity of insured drivers. However, unlike the case of advantageous selection, it is never possible to achieve the socially optimal outcome by risk-rating alone when there is a positive externality.

[6] Compensating variation approximates the change in consumer surplus under the Marshallian

Finally, I estimate the externality from insurance and use this to approximate the overall change in social welfare under risk-rating and other pricing regulations. The next section describes the empirical model.

2.3 Model of insurance and reckless behavior

In this section, I develop a structural model of insurer pricing and consumer insurance choice, reckless driving behavior, and ex post accidents and driving violations based on a finite horizon, discrete choice framework. The model provides a framework for evaluating the welfare effects of pricing policies.

2.3.1 Insurer pricing

I assume a competitive insurance market. In 2011 there were 77 sellers of private passenger automobile insurance in the United States, with 37, 46, 33, and 51 sellers in the study states of NC, PA, WA, and WI, respectively (National Association of Insurance Commissioners, 2012*a*).[7] The respective HHIs in those states were 935, 1,000, 787, and 958—all well below the upper threshold of 1,500 used by the DOJ and FTC to define concentrated markets. The highest single share of a seller in the study states is 21.7% in WI; all other sellers in all four states have shares below 20% (National Association of Insurance Commissioners, 2012*b*).[8] This contrasts with the health insurance industry, which in the same year had a state average HHI of 4,231 in the individual health insurance market, with similar HHIs in the small and large group markets (Kaiser Family Foundation, 2014)—a very high degree of concentration. A Maryland Insurance Administration report concludes that the

demand curve, and is a reasonable approximation when the price change is small. The measurement is exact when the demand curve is Hicksian—or when there are no income effects (Mas-Colell, Whinston and Green, 1995).

[7] Sellers are defined as insurance groups, where a group often owns many subsidiary companies.

[8] The four-firm concentration ratios were 52% (NC), 56% (PA), 45% (WA), and 53% (WI).

state's automobile insurance market is competitive—despite moderate concentration (HHI=1,226)—because there exists a high number of sellers (155 companies) and a declining share of the residual market, which would normally serve consumers who cannot find insurance at a reasonable price (Maryland Insurance Administration, 2015). The view that automobile insurance markets in the U.S. are competitive is further supported by data from the NAIC and industry experts (e.g., Insurance Information Institute, 2012).

The auto insurance industry also differs from health insurance in that there is better information among insurers and consumers, and consumers make more active choices, as auto insurance is typically not employer-sponsored. This is supported by the fact that auto insurers spend significantly more on advertising—the top six auto insurers spent over three billion dollars on advertising in 2015, compared to only 185 million by the highest-spending health insurer (Statista, 2015)—which suggests that auto insurers must work to retain existing customers and can induce others to switch. Indeed, popular websites suggest that consumers shop around to find lower prices for automobile insurance[9], as do state and national insurance associations (e.g., National Association of Insurance Commissioners, 2011; North Carolina Department of Insurance, 2010).

I assume insurers price plans according to a standard actuarial pricing model, where price is determined by the product of a base price and a set of rating factors S that reflect the average increase in expected cost due to plan characteristics—including liability limits and deductible—and insured driver characteristics—such as number of drivers, demographics, and driving violations. The pricing equation is given by

$$p_{ij} = Base \cdot \prod_{s=1}^{S} \Big[1\big\{Factor_{ijs}\big\} \cdot Weight_s \Big] \tag{2.1}$$

[9] See, e.g., edmunds.com, dmv.org, and nerdwallet.com articles about how to buy car insurance.

where p_{ij} is the premium quoted to individual i for policy j; $Base$ is the base price for the rating period; s indexes rating factors; $1\{Factor_{ijs}\}$ is one if factor s applies to individual i and policy j; and $Weight_s$ is the average increase in expected cost due to factor s.

2.3.2 Individual choices

Taking prices as given, individuals make a sequence of two choices: first, whether to purchase automobile insurance and how much, and second, whether to drive cautiously or carelessly, measured by driving after drinking too much during the policy year. Individuals are forward-looking: they have expectations about their future driving behavior and recognize both the potential consequences of this behavior and how insurance can protect them against financial loss. If a driver causes an accident, his *ex ante* choice of insurance determines his level of financial exposure because the insurance policy specifies the maximum amount the insurer will pay for an accident.

The model has a finite horizon, consistent with automobile insurance policies being written on six- or twelve-month terms. There are three stages of the individual choice model:

1. The individual chooses an insurance policy based on his risk preference, expected driving behavior, and expected accident costs.
2. Conditional on his insurance policy and expected accident costs, the individual chooses whether or not to drink and drive.
3. The individual enters the road, and driving events are realized.

The model can be solved via backward recursion. The following subsections describe in detail each stage of the model, beginning with stage 3 and working backward to stage 1. The preference parameters in each stage are denoted by α, subscripted by the stage number. Subscripts i, j, and d index individuals, insurance policies, and driving behavior.

Stage 3. Consequences of drinking and driving

Once an individual enters the road, he makes no more drinking decisions and receives no further utility from drinking. The individual receives (dis)utility in stage 3 only through the cost of realized accidents and arrests. Thus, the monetized flow utility in stage 3 is given by

$$\begin{aligned} U_{ijd3} &= u_{ijd3} + \varepsilon_{ijd3} \\ &= \sum_{k=1}^{K} 1\{\text{incur consequence } k\} \cdot Cost_{ijk} + \varepsilon_{ijd3} \end{aligned} \tag{2.2}$$

where k indexes the potential pecuniary consequences of accidents and arrests, $Cost_{ijk}$ is the cost that individual i with insurance policy j faces if consequence k if realized, and ε_{ijd3} is a mean-zero error term. The indicator function $1\{\cdot\}$ is 1 if the argument is true and 0 otherwise. The individual's costs, $Cost_{ijk}$, reflect how insurance protects him from financial loss. Expected utility in stage 3 is given by the expected value of equation 2.2:

$$E\big[U_{ijd3}\big] = \sum_{k=1}^{K} P_{idk} \cdot Cost_{ijk} \tag{2.3}$$

where P_{idk} is the probability that individual i incurs consequence k conditional on drinking behavior d chosen in stage 2. The expected cost of stage 3 depends on drinking and driving behavior through these probabilities.

Stage 2. Choice of drinking and driving

In the second stage, the individual chooses whether to drink and drive, taking into account the effect his decision has on the expected cost of an accident or DWI conviction given his insurance coverage. Flow utility is

$$\begin{aligned} U_{ijd2} &= u_{ijd2} + \varepsilon_{ijd2} \\ &= \alpha_{d2} X_{i2} + \lambda_d \eta_{i2} + \varepsilon_{ijd2} \end{aligned} \tag{2.4}$$

where X_{i2} is a vector of individual characteristics and η_{i2} is a time-invariant, individualspecific "type" known to the individual but not to the econometrician. The term ε_{ijd2} represents a stochastic shock that is observed by the individual in stage 2 and affects his preference for drinking and driving, such as an impromptu invitation to a party. I assume ε_{ijd2} is distributed type I extreme value to obtain the familiar, logit formula for the conditional choice probability of stage 2.

The conditional value function in stage 2 is

$$\begin{aligned} v_{ijd2} &= u_{ijd2} + \beta_2 E\big[u_{ijd3} + \varepsilon_{ijd3}\big] \\ &= \alpha_{d2} X_{i2} + \lambda_d \eta_{i2} + \beta_2 \sum_{k=1}^{K} P_{id}\big(\text{incur consequence } k\big) \cdot Cost_{ijk} \end{aligned} \tag{2.5}$$

The future value term in v_{ijd2} is the discounted utility value of equation (2.3)—the expected pecuniary costs of driving conditional on insurance policy j and drinking and driving behavior d. The coefficient β_2 is a composite term that combines the short-term discount rate, which measures how the individual weighs future consequences at the time he makes his drinking and driving decision, and the marginal utility of income.

The variables in X_{i2} reflect how attractive drinking and driving is to the individual: the importance of drinking to his social life; preference for driving—whether he drives to work; and demographic variables such as age and gender. X_{i2} also includes subjective valuation of non-pecuniary consequences of driving: changes in quality of life due to paralysis from an accident or due to criminal punishment for DWI. The individual chooses the level of drinking and driving that yields the highest value $V_{ijd2} = v_{ijd2} + \varepsilon_{ijd2}$.

Stage 1. Choice of insurance policy

In the first stage, the individual chooses from a menu of J automobile insurance policies plus the outside option of no insurance, $j = 0$, to maximize expected utility. The flow utility to individual i from policy j is

$$\begin{aligned} U_{ij1} &= u_{ij1} + \varepsilon_{ij1} \\ &= \alpha_{j1} X_{i1} + \delta(y_i - p_{ij}) + \varepsilon_{ij1} \end{aligned} \tag{2.6}$$

where X_{i1} is a vector of individual characteristics, y_i is income, p_{ij} is the premium individual i pays for policy j, and ε_{ij1} is an individual-policy specific preference shock that I assume follows a type I extreme value distribution. The conditional value function for policy j is

$$\begin{aligned} v_{ij1} &= u_{ij1} + \beta E\big[V_d|ij\big], \\ \text{where} \qquad E\big[V_d|ij\big] &= \ln\left[\sum_d \exp(v_{ijd2})\right] + \gamma \end{aligned} \tag{2.7}$$

is the expected value of the best drinking and driving choice at Stage 2 conditional on choosing policy j in Stage 1, and γ is the Euler-Mascheroni constant. This follows from the assumption that ε_{ijd2} is distributed type 1 extreme value (Rust, 1987). The individual chooses the insurance policy $j \in \{0, 1, \ldots, J\}$ that maximizes his expected utility—that is, that yields the highest $v_{ij1} + \varepsilon_{ij1}$.

Individuals receive varying degrees of flow utility from the purchase of insurance based on their preferences. The most obvious characteristic to affect flow utility from having insurance is risk aversion. This and other characteristics that affect insurance preferences (such as number of drivers and cars in the household) enter X_{i1}. Private information about risk is present at the time of the insurance choice in stage 1 because individuals are forward-looking about their future driving behavior (stage 2) and expected accident costs (stage 3). Premiums reflect observable information about risk to the extent that the insurer observes—and is allowed to use—information from driving records, the insurance application, and past claims experience.

2.4 Data

The primary data source is a unique panel survey focused on drinking and driving behavior in the U.S.—the Survey of Alcohol and Driving (SAD)—which elicited many dimensions of private information, such as risk preference and subjective probabilities of accidents and future driving behavior, in addition to objective risk and insurance coverage. I use three supplementary data sources to obtain a distribution of actual accident costs in the U.S., market values of the cars individuals report driving, and health insurance coverage to use in calculations of expected cost.

2.4.1 Survey of Alcohol and Driving

The SAD is a two-period panel survey of 1,634 individuals in four states (NC, PA, WA, WI) who had driven and had consumed alcohol in the month prior to the start of the survey. The administration of the survey is described in detail in Appendix B.

While administrative databases—e.g., state driving records—are large, they lack information on individual beliefs, preferences, behavior, and attributes of insurance policies. The rich survey data allows me to include direct measures of variables like risk preference and non-pecuniary loss in the model rather than having to impose additional structure to estimate them. The survey also spans multiple states and insurers. This contrasts with insurance studies using administrative data, which is commonly collected from only a single state, insurer, or segment of the population, such as Medicare enrollees.

One concern with the survey data is that the sample is not representative. In Appendix Table B.1, I compare the final estimation sample from the SAD to national data from the Behavioral Risk Factor Surveillance System. There are significant demographic differences: on average, the SAD sample is more educated, more likely to be employed, has higher income, has a smaller household size in terms of both adults and children, and is disproportionately white. However, there are no significant differences in terms of the likelihood of drinking and driving. A potential bias for this study is that higher income households are less likely to be judgment proof, and therefore more likely to purchase insurance, which may cause the uninsured rate in this study to be disproportionately low and bias the price elasticity of demand toward zero. This is likely to attenuate the welfare effects from pricing policies.

Insurance choice set The ten policies in the insurance choice set, including no insurance, reflect state laws and observed policies in the data. All policies must include liability insurance by law. The minimum bodily injury and property damage liability limits are based on state laws in effect at the time the SAD was conducted. The per-accident bodily injury minimum liability limits for PA, WA, NC and WI are \$30k, \$50k, \$60k, and \$100k. Options for higher per-accident bodily injury liability limits are based on frequently observed limits in the data: \$100k (except WI), \$300k, \$500k, and \$1M.[10] For each liability limit, the individual can choose collision or no collision coverage, with the exception that a \$1M limit without collision is not an option because this combination is rarely observed.

The most common insurance choice in the sample is a \$300k liability limit with collision coverage (Appendix Table B.2). A majority buy collision coverage (89%). Only 15% choose the minimum level of liability coverage allowed in their state, and only 4% are uninsured, which is much lower than the 13.5% of drivers uninsured

[10] I do not explicitly model the choice of property damage liability limits, but they are needed to calculate the expected loss under each insurance policy. For the minimum policies, I use the state laws—\$5k, \$10k, \$25k, and \$15k for PA, WA, NC, and WI, respectively. Property damage liability limits were not collected in the SAD, so for policies in the choice set that are above the state minimums, I assume the property damage liability limit is 25% of the per-accident bodily injury liability limit.

nationally in 2011 (Brobeck, Best and Feltner, 2013).

Premium-setting variables The proportion of uninsured drivers with accidents and speeding violations is similar to insured drivers, but a striking 48% of uninsured drivers had a DWI arrest in the past three years, while only 1% of insured drivers did (Table 2.1). This suggests that arrest for DWI is an important determinant of whether an individual has automobile insurance, which could be due to high experience-rated premiums, actual or expected denial due to insurer underwriting, or legal consequences of a DWI arrest (e.g., license suspension). However, few individuals in the sample report being denied coverage, so I assume all drivers have the option of buying insurance.[11]

Drinking and driving behavior Individuals have substantial private information about their drinking and driving behavior since the probability of arrest for this behavior is so low. In the twelve months following the baseline survey, 17% of the sample did drink and drive, measured as driving after having 5 or more drinks for men or 4 or more drinks for women (Table 2.1). This is much higher than the proportion of the sample that was arrested for DWI during that time, which was only about 1% (not shown). The high percentage of drinker-drivers provides statistical power in the empirical analysis, which is based on reported drinking and driving behavior, not DWI arrests.

Individuals also have heterogeneous preferences that likely influence drinking and driving behavior, but are unobserved by the insurer. The most common response to

[11] This could be because very high-risk drivers believe they would be denied coverage, and therefore do not attempt to purchase it. However, states have various policies in place to ensure that high-risk drivers can obtain insurance, such as take-all-comers, surplus line insurers, and high-risk pools. Again, few drivers report obtaining coverage through these means, but in some states, such as North Carolina, drivers will not be aware if they are in a high-risk pool because they can purchase insurance from any carrier, and the carrier decides internally whether to place the driver in the high-risk pool.

Table 2.1: Summary statistics

	Insured		Uninsured	
Variable	Mean	Std. Dev.	Mean	Std. Dev.
Insurance cost				
Annual premium	1,242	935	—	—
Observable individual risk				
Chargeable accidents	0.03	0.17	0.07	0.26
Speeding $\geqslant$ 15mph	0.15	0.42	0.12	0.45
DWI arrests	0.01	0.15	0.48***	1.05
Miles driven per week	159	165	64.66***	82.07
Drives to work	0.76	0.42	0.32***	0.46
Observable household risk				
Adult drivers ($\geqslant$ 25)	1.59	0.63	1.10***	0.48
Young drivers ($<$ 25)	0.45	0.84	0.19**	0.45
Cars in household	2.22	1.15	1.29***	0.67
Risk preference				
Risk preference factor	-0.04	0.96	0.97***	1.30
Reckless behavior (past 12 mos.)				
Did drink and drive	0.17	0.38	0.26	0.45
Subjective beliefs (in next year)				
Prob. of automobile accident	0.13	0.14	0.14	0.18
Prob. of accident with injury	0.08	0.11	0.11*	0.16
Odds of accident: 4 vs. 0 drinks	1.50	0.36	1.39**	0.34
Prob. of drinking and driving	0.19	0.31	0.15	0.23
Prob. of conviction $\mid$ DWI	0.09	0.14	0.19***	0.24
Premium % incr. $\mid$ accident	1.39	0.39	1.47	0.28
Premium % incr. $\mid$ DWI convic.	1.80	1.20	1.83	0.00
Demographic characteristics				
Family net worth ($1,000s)	382	642	139.11**	427
Own car value ($1,000s)	9.61	7.50	8.73	9.23
Has health insurance	0.91	0.28	0.60***	0.50
Observations	924		42	

Notes: Annual premium excludes observations that are missing (n=19) or that lie in the first and last 200-quantiles (n=7). The risk preference factor is positively correlated with risk-taking behavior. Asterisks indicate statistical significance (* 0.10 ** 0.05 *** 0.01) in pairwise t-tests of the difference in means between the insured and uninsured subsamples.

the question, "How important is it for your social life to be able to enjoy a few drinks with your friends?" is "Slightly important" (44%), whereas 7% of the sample reports that drinking is "Very important" to their social life (not shown). In response to the question, "How great an effect would it have on your ability to conduct your personal life if you were to be arrested for drunk driving?", 17% report that "It would almost ruin my life".

Risk preference I use factor analysis to construct a measure of risk preference.[12] The factor, which loads heavily on impulsivity, smoking, and drug use, and is also positively related to risk-taking in the health domain. The range of the factor is -1.68 to 3.95, and the mean in the insured subsample is -0.04. The uninsured subsample has a statistically significantly higher mean of 1.30. This may partially reflect a correlation between the factor and past risk-taking behavior in the driving domain. However, it likely also reflects the importance of risk preference for insurance choice, which is widely recognized as a source of private information in insurance markets (e.g., Finkelstein and McGarry, 2006).

Subjective beliefs Comparing mean subjective probabilities to the objective counterparts in the sample shows that individual perceptions of risk are different from objective risk in many domains, which motivates the use of subjective probabilities in estimation (Manski, 2004). The mean subjective probability of getting into any accident in the next year is 0.13, which is higher than the mean accident rate observed in the sample (about 0.09 in the follow-up year). The mean expected increase in the odds of an accident when drinking and driving is 1.5 times, which is lower than the actual increase in odds in the sample of 2.5. The mean expected premium

[12] The variables included in the factor analysis are: risk preference in the financial domain; risk preference in the health domain; impulsivity; whether the individual smokes; and whether the individual has ever used illicit drugs or licit drugs without a prescription. The factor analysis yields two factors with eigenvalues over 1. I use the first factor.

increases following an accident and DWI are 39% and 80%, both lower than the actual increases in the sample: 63% and 83% (not shown).

On average, uninsured individuals report a higher subjective probability of being arrested for DWI in the next year than insured individuals but a lower subjective probability of getting into an accident while drinking and driving relative to driving sober. Uninsured individuals also have lower net worth than insured individuals and are less likely to have health insurance.

2.4.2 Additional data

There are three important data elements that are not elicited in the SAD: the dollar cost of an accident, the value of individuals' vehicles, and out-of-pocket exposure to healthcare costs.[13] I use different sources for each of these three measures.

First, I obtain the distribution of accident costs incurred in the U.S. in 2010 by accident severity level from the National Highway Traffic Safety Administration (NHTSA) and use it to calculate the expected cost of an accident under each insurance plan. The NHTSA reports accident costs at eight severity levels: property damage only, MAIS 0–MAIS 5[14], and fatal. Costs include property damage, medical cost, other economic costs, and lost life or quality-of-life. The value of lost life is life expectancy at time of death minus age, multiplied by the value of a year of statistical life of $8.86 million. For injuries, the value of lost quality-of-life is based on the severity and duration of the injury using quality-adjusted life years (Blincoe et al., 2014).

The decision about how much automobile insurance to purchase is not trivial, given that mean accident costs range from $4,544 to $25.9 million, depending on the

[13] The SAD does elicit accident costs, but very few individuals were actually involved in an accident, making this data element too sparse to be useful in estimation.

[14] MAIS stands for "Maximum Abbreviated Injury Scale", which is a measure of the severity of injuries, where zero is least severe and five is most severe.

severity of the accident (Appendix Table C.1). The highest-cost accidents are the least likely to occur; however, many individuals choose bodily injury liability limits well above the expected cost of an accident with bodily injury of $160k (Table B.2). This suggests that other factors, such as perception of and preference for risk, affect insurance choice.

Second, I estimate the value of each individual's primary car by looking up the Edmunds.com True Market Value® for a typically equipped vehicle by make, model, and year reported in the SAD.[15] I adjust prices back to 2011 by accounting for four years of depreciation: 25% in the car's first year, 15% in years two through five, and 7.5% in all later years.[16]

Finally, I obtain average health insurance deductibles from the Medical Expenditure Panel Survey. The mean deductible for an employee at a private-sector establishment in 2011 was $1,123 for single coverage and $2,220 for family coverage (Agency for Healthcare Research and Quality, 2011). I use these values, along with car prices and accident costs, to calculate expected costs for stage 3 of the model.

2.5 Estimation strategy

The estimation strategy involves several steps: (1) a hedonic model of premium-setting; (2) precise calculation of expected costs; (3) a new approach to estimating unobserved heterogeneity; and (4) a maximum likelihood routine. The first three steps recover values needed to specify the model for the optimization routine, whereas the fourth recovers the preference parameters of the sequential model within the optimization routine.

[15] The Edmunds.com True Market Value® (Edmunds.com TMV®) price is Edmunds.com's determination of the current average base dealer retail price, unadjusted for color or any options. Prices in this study were accessed 5/15/15. TMV® prices are copyrighted by Edmunds.com, Inc., which reserves all rights. Edmunds®, Edmunds.com True Market Value® and TMV® are registered trademarks of Edmunds.com, Inc.

[16] www.edmunds.com/car-buying/how-fast-does-my-new-car-lose-value-infographic.html

The estimation strategy takes advantage of the observed individual subjective probabilities described in section 2.4 to explicitly specify parts of the model that are typically unobserved, which has two distinct advantages. First, it explicitly accounts for individuals' subjective probabilities, which were shown in the previous section to differ from objective probabilities. Second, it allows for precise estimation of parameters that would otherwise rely on low probability outcomes (chargeable accidents and DWI convictions).

2.5.1 Step 1. Hedonic premium regression

In the first step, I linearize the pricing model given in equation 2.1 and estimate the weights for each rating factor as the coefficients of a hedonic regression of premium on plan and driver characteristics. I use these results to estimate premiums for unobserved plans—the SAD elicited the insurer, premium, and characteristics for the policy that each household chooses, but did not collect premiums for policies that the household could have chosen but did not. I simulate premiums for the ten insurance policies in Appendix Table B.2 using the estimated weights (implicit prices) of the plan components. The hedonic regression is

$$Premium_i = \beta_0 + \beta_1 Liability_i + \beta_2 Collision_i + \beta_3 Ded_i + \beta_4 X_i + \alpha_k + \varepsilon_i \quad (2.8)$$

where $Premium_i$, $Liability_i$, $Collision_i$, and Ded_i are the monthly premium, upper liability limit, collision indicator, and collision insurance deductible of the policy that household i owns; X_i is a vector of household driver characteristics; α_k is an insurer-specific fixed effect; and $\beta = [\beta_0\ \beta_1\ \beta_2\ \beta_3\ \beta_4]$ is the vector of parameters of interest. The sum of β_0 and α_k is an estimate of the base price for insurer k. The vector X_i contains the following variables that reflect the risk of the insured drivers on household i's policy: the respondent's chargeable accidents, DWI arrests, annual mileage, and whether he drives to work, as well as the total number of adult drivers (age $\geqslant$ 25), young drivers (age $\leqslant$ 24), and cars in the household.[17]

Equation (2.8) can only be estimated for insured respondents. To correct for sample selection, I use a Heckman correction. The selection equation is

$$Insured_i = \gamma_0 + \gamma_1 X_i + \gamma_2 Z_i + \varepsilon_i \tag{2.9}$$

where $Insured_i$ is a binary variable for whether or not the respondent reports having any auto insurance coverage. X_i contains explanatory variables that affect premiums and that effect the probability that the respondent has any insurance coverage: the number of accidents and arrests of the respondent in the past three years and the number of cars and drivers in the household. I use risk preference as the instrumental variable Z_i because it affects willingness-to-pay for insurance, but is excluded from the premium regression because it is private information. The outcome equation is given by equation (2.8), plus an additional term for the inverse Mills ratio.

2.5.2 Step 2. Compute individual expectations of stage 3 costs

I compute expected costs for each individual under each insurance policy and drinking and driving behavior combination—i.e., the value of equation 2.3 for all i, j, and d. The equation depends on probabilities and costs. I use individuals' subjective beliefs about accidents and DWI arrests from the SAD for the probabilities. (The relevant probabilities are summarized in Table 2.1).

The calculation of the cost components of equation 2.3 is more complex because the cost faced by the individual depends on many things. I combine variables from the SAD with the additional data on accident costs, health insurance deductibles,

[17] The variables included in equation 2.8 capture all the policyholder characteristics that an insurer would typically use when setting premiums, which was part of the intentional design of the SAD. There are some minor insurance plan characteristics that are not observed, such as personal injury protection and roadside assistance. However, these are second-order features with relatively low actuarial value. It is also possible that there are nonlinearities in how premium changes with the liability limit. I performed robustness checks including various polynomial terms for the liability limit, and the estimated non-linear coefficients were not statistically significant.

Table 2.2: Potential pecuniary losses from driving

Cost	Formula	Notes
Bodily injury to others	$\min(\max(0, OthersInjuries - BodilyLimit), NetWorth)$	$BodilyLimit = 0$ if no liability insurance
Bodily injury to self	$\min(HealthDeductible, OwnInjury, NetWorth)$	$HealthDeductible = \infty$ if no health insurance
Damage to others' property	$\min(\max(0, OthersDamages - PropertyLimit), NetWorth)$	$PropertyLimit = 0$ if no liability insurance
Damage to own property	$\min(OwnDamage, CollisionDeductible, OwnCarValue)$	$CollisionDeductible = \infty$ if no collision insurance
Premium increase due to accident	$Premium \cdot (\%PremiumIncrease \mid Accident)$	$Premium = 0$ if no auto insurance
Premium increase due to DWI	$Premium \cdot (\%PremiumIncrease \mid DWI)$	$Premium = 0$ if no auto insurance
Fine for DWI	$Fine \mid DWI$	

Notes: The table defines the pecuniary costs to a driver of driving, conditional on insurance coverage. The first column lists the potential costs. The second column defines a formula for how much of the actual cost the driver is held responsible for. The third column notes how the formulas change for special cases with limited insurance coverage. When individuals choose their reckless driving behavior, the following are random variables: *OthersInjuries*, *OwnInjury*, *OthersDamages*, and *OwnDamage*. These are realized at the time of an accident and can take a wide range of values depending on the severity of the accident. Variables that are known at the time of the reckless driving decision are: *BodilyLimit*, *PropertyLimit*, *CollisionDeductible*, *Premium*, *HealthDeductible*, *OwnCarValue*, and *NetWorth*, which are based on the individual's insurance coverage and household demographics. Individuals also have expectations about $\%PremiumIncrease \mid Accident$, $\%PremiumIncrease \mid DWI$, and $Fine \mid DWI$.

and average car prices described in Section 2.4. The formulas for calculating each of the cost components are given in Table 2.2. I use the NHTSA distribution of accident costs in the U.S. for the four components of total accident cost: *OthersInjuries*, *OwnInjury*, *OthersDamages*, and *OwnDamage*. The values of *HealthDeductible* and *OwnCarValue* are approximated using data from MEPS and Edmunds.com (see Section 2.4). The automobile insurance characteristics *BodilyLimit*, *PropertyLimit*, and *CollisionDeductible* are defined by the insurance policy j. *Premium* is based on the estimated premium for policy j from Step 1. The remaining variables, *NetWorth*, $\%PremiumIncrease \,|\, Accident$, $\%PremiumIncrease \,|\, DWI$, and $Fine \,|\, DWI$ are from the SAD, where the latter three were elicited as subjective beliefs.

A key feature to observe about the formulas in Table 2.2 is that the accident cost to the individual is capped by the value of his car (for own property damage) and his net worth. Thus, the assumption that people know the national distribution of accident costs is more innocuous than it may seem at first. It turns out that net worth is binding for most individuals except at the lowest accident severity levels. In forming an expectation about costs and choosing an insurance policy, an individual is likely to think about the maximum amount he could lose, which is readily known to him as the value of his car and his net worth.

2.5.3 Step 3. Estimate time-invariant unobserved heterogeneity

I recover time-invariant unobserved preference for drinking and driving by making the observation that subjective probabilities and actual choice probabilities of drinking and driving from the model should be equal. We can derive a standard expression for the choice probability that depends only on data and parameters and then set this equal to the subjective probability in the data and use OLS to recover a residual term that affects the probability of drinking and driving but is not observed by the econometrician.

Given that ε_{ijd2} is distributed type I extreme value, the probability that individual i with insurance policy j chooses drinking and driving level d is:

$$P_{ijd2} = \frac{\exp(v_{ijd2})}{\sum_{d'} \exp(v_{ijd'2})} \tag{2.10}$$

The log of the ratio of the probability of drinking and driving to the probability of not drinking and driving is given by:

$$\begin{aligned}
\ln(P_{ijd2}) - \ln(P_{ij02}) &= v_{ijd2} - v_{ij02} \\
&= \Big(\alpha_{d2}X_{i2} + \lambda_d\eta_{i2} + \beta E\big[u_{ijd3}\big]\Big) - \Big(\alpha_{02}X_{i2} + \lambda_0\eta_{i2} + \beta E\big[u_{ij03}\big]\Big) \\
&= \Big(\alpha_{d2} - \alpha_{02}\Big)X_{i2} + \beta\Big(E\big[u_{ijd3}\big] - E\big[u_{ij03}\big]\Big) + \Big(\lambda_d - \lambda_0\Big)\eta_{i2} \\
\tilde{y}_i &= \tilde{\alpha}X_{i2} + \beta\tilde{X}_{i3} + \tilde{\lambda}\eta_{i2} \qquad (2.11)
\end{aligned}$$

If we assume the choice probability is equal to the subjective probability, then both sides of this equation are observed: the left-hand side is a transformation of the probability of drinking and driving, and the right hand side is the difference between the choice-specific conditional value functions for drinking and driving and not drinking and driving in stage 2 (eq. 2.5). I calculate $\tilde{X}_{i3} = E\big[u_{ijd3}\big] - E\big[u_{ij03}\big]$ with respect to the insurance policy j that the individual chooses (if the individual is uninsured, then $\tilde{X}_{i3} = 0$) and then estimate equation (2.11) by OLS and recover the residuals for each individual, $\widehat{\tilde{\lambda}\eta_{i2}}$.[18]

2.5.4 Step 4. Maximum likelihood estimation

I use maximum likelihood to estimate the preference parameters. Assuming that ε_{ij1} and ε_{ij2} are distributed type I extreme value and independent, the probability that

[18] Because $\ln(x)$ is not defined for $x = 0$, and 38% of the sample reports a subjective probability of drinking and driving of zero, and another 7% reports a subjective probability of one, I set the probabilities for these observations to 0.00001 and 0.99999. Varying the approximations to 0 and 1 and estimating this equation with quantile regression both do not substantially impact the results.

individual i drinks and drives and the probability that individual i chooses insurance policy j are

$$P_{ij1} = \frac{\exp(v_{ij1})}{\sum_k \exp(v_{ik1})} \qquad \text{and} \qquad P_{ijd2} = \frac{\exp(v_{ijd2})}{\sum_{d'} \exp(v_{ijd'2})} \tag{2.12}$$

The log likelihood function is

$$L = \sum_i \left(\sum_{j=0}^{J} 1(d_{i1} = j) \log(P_{ij1}(X_i, y_i, p_{ij}, \alpha_1, \delta)) + \sum_{d=0}^{D} 1(d_{i2} = d) \log(P_{ijd2}(X_i, \alpha_2)) \right)$$

where $1(d_{i1} = j)$ is one if the individual chose insurance policy j and zero otherwise, $1(d_{i2} = d)$ is one if the individual chose to drink and drive at level d and zero otherwise, and $P_{ij1}(X_i, y_i, p_{ij}, \alpha_1, \delta)$ and $P_{ijd2}(X_i, \alpha_2)$ are the choice probabilities given in equation (2.12). Maximizing the log likelihood function over α_1, α_2, and δ yields estimates of these preference parameters. I estimate bootstrapped standard errors to incorporate estimation error from each step of the estimation strategy.

2.6 Results

The estimation results establish that the automobile insurance market is advantageously selected and show that preferences and expected disutility of non-pecuniary losses are as important in an individual's decision to engage in reckless driving behavior as pecuniary losses. These findings affect how individual choices and welfare respond to counterfactual pricing policies.

2.6.1 Model results

Premiums More risk tolerant people are less likely to be insured, as expected (Table 2.3). Premiums are higher for individuals who have had past driving violations—each accident increases annual premium by \$406—which confirms that insurers use observable information about risk. These high premiums may induce high-risk drivers

Table 2.3: Hedonic premium regression

	No selection adjustment (1)		With selection adjustment (2)	
Panel A. Selection equation			Has auto insurance	
Chargeable accidents			−0.04	(0.48)
DWI arrests			−0.81***	(0.21)
Cars in household			0.51***	(0.14)
Miles per week (100s)			0.29**	(0.00)
Drives to work			0.85***	(0.19)
Risk preference factor			−0.28***	(0.08)
Inverse Mills ratio			−479.89	(372.08)
Panel B. Premium equation	Annual premium		Annual premium	
Plan characteristics				
Liability limit ($100k)	30.49**	(14.65)	29.08**	(14.36)
Has collision insurance	440.83***	(127.08)	422.09***	(123.87)
Collision deductible	−0.04	(0.09)	−0.04	(0.09)
Insurer observables				
Chargeable accidents	388.55**	(167.80)	406.03**	(165.76)
Speeding $\geqslant$ 15mph	95.61	(70.63)	90.20	(69.40)
DWI arrests	−0.36	(195.71)	154.91	(223.07)
Adult drivers ($\geqslant$ 25)	126.30**	(51.72)	121.32**	(50.60)
Young drivers (< 25)	208.94***	(37.99)	210.64***	(37.23)
Cars in household	119.88***	(29.81)	105.37***	(31.52)
Miles driven per week	−0.32*	(0.18)	−0.38**	(0.18)
Drives to work	−55.86	(72.02)	−118.02	(85.99)
Observations	899		941	

Notes: The table reports the results of estimating equation (2.8) on the subsample of insured individuals with reported annual premium that lies between the first and last 200-quantiles. Column (2) uses a Heckman correction to account for selection in the insurance purchase decision, identified by private information about risk preference. The dependent variable in the selection equation is an indicator for whether the respondent has any automobile insurance. The dependent variable in the premium equation is annual automobile insurance premium. All specifications include area, car type, and insurer fixed effects in the premium equation. Standard errors are in parentheses. Asterisks indicate statistical significance at the * 0.10 ** 0.05 *** 0.01 levels.

to drop coverage. The negative relationship between having a DWI arrest in the past three years and having insurance biases the coefficient on DWI arrests toward zero (Panel B, col. 1), motivating the Heckman correction. This yields a higher premium penalty of $155 per DWI arrest, though this is not statistically significant due to lack of power (col. 2).[19]

The coefficients on insurance coverage characteristics are as expected and align with industry averages: premium increases with the liability limit at a rate of $29 per $100k of liability coverage; is higher by $422 per year for plans with collision coverage; and is decreasing in the deductible (col. 2). The estimated effect of adding collision coverage is remarkably close to the mean combined expenditures of $437 on collision and comprehensive coverage in 2007 reported in Insurance Information Institute (2011). It is likely that the measure of collision insurance from the SAD data also reflects comprehensive coverage as the two are often bundled.[20]

Expected costs There is substantial variation across insurance policies and drinking and driving behavior in the expected costs of driving—from $325 when the mean driver is covered with the most comprehensive insurance plan and does not drink and drive to $2,638 when the mean driver is uninsured and drinks and drives (Appendix Table C.2). Expected costs vary across drivers depending on their subjective probabilities of accidents with and without bodily injury, of being convicted for DWI, and their net worth, car value, and health insurance coverage. In general, however, higher liability limits, collision insurance, and not drinking and driving all reduce drivers' expected costs.

[19] It is also possible that an individual arrested for DWI was not actually convicted, or was at least able to get the charges reduced. Either of these situations would lead to a smaller impact on the premium than a conviction. Unfortunately, the survey did not explicitly ask about past DWI convictions, only arrests. This is one variable for which the information I have might deviate from what the insurer has, biasing the estimate toward zero.

[20] http://www.progressive.com/understanding-insurance/entries/2009/9/1/can_you_have_compre/

Table 2.4: Insurance policy choice

	Insurance policy								
Collision coverage:	No collision				Collision				
Upper liability limit:	Min.	$100k	$300k	$500k	Min.	$100k	$300k	$500k	$1M
Panel A: Coefficients on individual characteristics									
Risk preference factor	0.08	−0.18	−0.27	−0.76*	−0.37**	−0.47***	−0.66***	−0.79***	−1.14***
	(0.21)	(0.39)	(0.30)	(0.44)	(0.18)	(0.17)	(0.16)	(0.17)	(0.23)
Subj. prob. of accident	−3.53	−1.27	−1.76	5.27	0.94	0.79	−0.05	−1.68	0.78
	(2.80)	(4.59)	(2.79)	(3.38)	(1.47)	(1.60)	(1.52)	(1.53)	(1.86)
Cars in household	0.62	0.57	0.42	1.55***	0.71**	1.01***	0.81**	0.95***	1.11***
	(0.41)	(0.57)	(0.56)	(0.39)	(0.35)	(0.32)	(0.34)	(0.36)	(0.38)
Adult drivers	0.73	0.12	0.54	1.31	1.05	0.81	1.46*	1.36	1.79**
	(0.99)	(1.20)	(1.30)	(0.99)	(0.90)	(0.87)	(0.87)	(0.87)	(0.89)
Young drivers	0.42	−0.38	−0.69	0.30	0.44	0.20	0.27	0.39	−0.55
	(0.63)	(2.19)	(2.65)	(1.26)	(0.52)	(0.53)	(0.52)	(0.52)	(0.65)
Observations	32	15	11	8	107	134	389	188	40
Panel B: Coefficients on plan characteristics									
Income net of premium	16.25*								
	(9.64)								
Discount factor (assumed)	0.97								
Panel C: Implied premium elasticities									
Premium elasticity	−1.03	−1.06	−1.16	−1.26	−1.53	−1.50	−1.10	−1.55	−2.07

Notes: The table reports estimated parameters for the insurance choice term of the likelihood function—equation (2.5.4). "Min." is the legal minimum liability limit in the respondent's state. The omitted category is uninsured (n=42). Risk preference is positively correlated with risk-taking behavior in the health domain (see text for details). Premium is the individual-specific predicted value from the hedonic regression for each plan. Implied elasticities are given by the percentage change in the probability of choosing each plan with respect to a percentage change in income net of premium. Standard errors based on 500 bootstrap samples are in parentheses. * 0.10 ** 0.05 *** 0.01.

Insurance choice Individuals with stronger preferences for risk are less likely to be insured overall and are less likely to choose high levels of automobile insurance coverage (Table 2.4). This can lead to advantageous selection if risk preference also affects risk-taking behavior that influences accident risk. Despite concerns over private information about risk type, there is no evidence that individuals with higher subjective probabilities of an accident buy more complete insurance. Risk preference is the dominant source of private information in this market.

Premiums are quite relevant to individuals' insurance policy choice: the implied elasticities of the probability of choosing each plan (a liability-collision combination) with respect to premium are sizable, ranging from -1.03 to -2.07, with higher elasticities associated with richer plans (Panel C). Households with more cars and adult drivers are more likely to choose plans with collision coverage and liability limits above the minimum.

Drinking and driving More risk tolerant individuals are more likely to drink and drive, confirming that there is advantageous selection: risk-tolerant individuals are less likely to buy insurance, but more likely to engage in reckless driving behavior—and therefore more likely to be involved in an accident (Table 2.5). Individuals who report that drinking is a "very important" part of their social life are also more likely to drink and drive. The probability of drinking and driving decreases as the importance of drinking to social life decreases.

Expected pecuniary cost has a negative and statistically significant effect on the probability of drinking and driving, but the magnitude of the marginal effect is small, implying that there is relatively little moral hazard in drinking and driving behavior. In particular, a $1,000 increase in the expected loss leads to, on average, a 0.3% decrease in the probability of drinking and driving, where the expected losses in the sample range from $0-50k (not shown). In contrast, the marginal effect of non-

Table 2.5: Drinking and driving choice

	Drink and drive	
Preferences		
Risk preference factor	0.73***	(0.28)
Importance of drinking to social life:		
Not at all important (omitted)		
Slightly important	2.76***	(0.68)
Quite important	4.54***	(0.82)
Very important	6.57***	(1.18)
Valuation of consequences		
Expected pecuniary loss ($10k)	−6.18***	(2.18)
WTP to avoid paralysis ($/month)	−0.02**	(0.01)
Effect of DWI arrest on life:		
It would almost ruin my life (omitted)		
It would hurt me badly	0.36	(0.82)
It would hurt me some/none	0.74	(0.82)
Demographic and driving characteristics		
log(*age*)	−2.35***	(0.71)
Female	−2.03***	(0.31)
Drives to work	2.52***	(0.86)
Observations	966	

Notes: The dependent variable is an indicator for whether the respondent drove after drinking four or five drinks at least once during the follow-up year. Standard errors are in parentheses. Asterisks indicate statistical significance at the * 0.10 ** 0.05 *** 0.01 levels.

pecuniary cost, as measured by willingness-to-pay to avoid paralysis, is about three times as large as that of pecuniary cost in the expected direction: the higher one's willingness to pay to avoid paralysis from an automobile accident, the less likely he is to drink and drive. This intangible consequence cannot be insured against by the market, so the best decision for individuals with a high willingness-to-pay to avoid injury is to take precaution by not driving after drinking, regardless of automobile insurance status, a point illustrated by Ehrlich and Becker (1972).

2.6.2 Model fit

The predicted shares of each insurance policy and the number of drinker-drivers from the model match the data quite well. In both the model and the data, 4% of drivers are uninsured and 18% report drinking and driving during the follow-up year. The full distribution of insurance choices is very close. The share of drivers without collision matches at every liability limit, while the share of drivers with collision deviates just slightly: 0.12, 0.13, 0.40, 0.19, and 0.04 in the data (Table B.2) versus 0.11, 0.13, 0.41, 0.20, and 0.04 in the model.

2.6.3 Discussion

It is worth pausing to consider how these results align with some recent issues in theoretical and empirical literature on asymmetric information in insurance markets.

Market structure and asymmetric information One branch of the literature is focused on understanding how market structure interacts with asymmetric information in insurance markets. Chiappori et al. (2006) laid the groundwork for this work. One of their key theoretical results is that a competitive market must exhibit a positive correlation between insurance coverage and accident risk (i.e., adverse selection). This seemingly creates a puzzle for my work because I find advantageous selection in a competitive market. However, Spinnewijn (2013) shows that violations of Chiappori et al. (2006) "realistic expectations" assumption—i.e., the individual's subjective assessment of risk is equal to the objective risk—can break the positive correlation, opening up the possibility for advantageous selection. In particular, advantageous selection can arise when a group of individuals are both overly optimistic about their true risk (so-called "baseline optimistic" in Spinnewijn (2013)) and pessimistic about how much their behaviors affect that risk ("control pessimistic"). Chiappori et al. (2006) argue that the realistic expectations assumption is "indispensable for empirical applications" because subjective beliefs are not observable. In the SAD, however, subjective beliefs are observed. In Appendix D, I use subjective beliefs from the SAD together with other features of the survey to show that in my sample, risk tolerant individuals are baseline optimistic and control pessimistic, which can explain why we observe advantageous selection in a competitive market.

Moral hazard In another branch of the literature, researchers are searching for better methods for empirically measuring moral hazard, with some recent applications to automobile insurance markets. While most early empirical work did not find evidence of moral hazard in automobile insurance (e.g., Chiappori and Salanié, 2000; Abbring, Chiappori and Pinquet, 2003), some recent papers have found evidence of moral hazard. A key takeaway of these papers, though, is that the findings do not apply to all driving behavior—rather, moral hazard is highly context-specific.[21] For example, Dionne, Gouriéroux and Vanasse (2001) find moral hazard in traffic violations but not accidents; Weisburd (2015) only finds evidence of moral hazard in "small collisions"—defined as collisions related to parking rather than driving on roads; and Robinson, Sloan and Eldred (forthcoming) find that moral hazard is much more likely to be manifested through speeding than drinking and driving behavior, which is far less likely to lead to severe accidents. These papers are consistent with my finding of little moral hazard in drinking and driving. My model does have the limitation that it does not allow for moral hazard to be manifested through driving behaviors other than drinking and driving. However, to the extent that moral hazard only affects minor accidents—where a significant portion of the cost of the accident is property damage to one's own car—omitting channels other than drinking and driving should have

[21] In their widely-cited survey of the literature, Cohen and Siegelman (2010) look at a broad set of markets beyond automobile insurance and also conclude that moral hazard is context-specific. Among the papers on automobile insurance, only one finds evidence of moral hazard: Israel (2004) finds statistically significant evidence of moral hazard using the panel-data approach developed by Abbring et al. (2003), but the magnitude of the moral hazard effect is small.

little impact on my main counterfactual findings, which are unrelated to first-party losses.[22]

Overall, the finding that moral hazard primarily affects less risky behaviors and more minor accidents is intuitive because, as shown in my results, individuals care about non-pecuniary consequences like self-injury—they won't engage in infinitely more risky behavior simply because the pecuniary cost is lower since they also bear the non-pecuniary costs of such behavior.

2.7 Counterfactual pricing policies

I simulate three counterfactual pricing policies—risk-rating, community rating, and experience rating—to measure the effects of asymmetric information on insurance choice, drinking and driving behavior, and welfare.

Risk-rating relaxes information constraints and allows for symmetric information about drinking and driving behavior. The insurer observes and sets premiums based on drivers' propensities to drink and drive. Community rating restricts the insurer from using any individual characteristics to set premium—all drivers face the exact same menu of plans and premiums.[23] Policymakers may favor community rating for its egalitarian appeal—driving is often viewed a right, because being able to operate an automobile is necessary for maintaining employment, participating in household tasks, etc. The third counterfactual policy requires strict experience rating: insurers must increase premiums substantially following accidents and convictions.

To calculate the welfare effects of these counterfactual policies, I first adjust premiums according to the counterfactual policy and then calculate changes in expected utility using the estimated model parameters.

[22] In particular, the cost of uninsured motorists' coverage is not affected by a change in expected loss due to own property damage, which is covered entirely by first-party collision insurance.

[23] In practice, some community rating systems use rating bands (min. and max. premiums), and some allow for the use of a few attributes of the insured such as age category.

2.7.1 Constructing counterfactual premiums

I construct counterfactual premiums for the risk-rating simulation by redistributing premium from non-drinker-drivers to drinker-drivers based on the increase in expected accident costs for drinker-drivers relative to the average expected accident cost, conditional on risk factors observed by the insurer (see Appendix E for details). If drinker-drivers are more accident-prone than non-drinker-drivers conditional on information that the insurer possesses, then there is a true private information component to drinking and driving behavior, and some non-drinker-drivers must pay more than the actuarial value of their policies under the status quo.

Constructing counterfactual premiums under pure community rating is straightforward in a partial equilibrium framework: I take the mean premium of each plan and set each individual's premiums to those mean values. This simulates the short-run effect of community rating by not allowing individuals to re-sort among plans or insurers to dynamically adjust prices to reflect the change in composition of drivers in each plan. In reality, regulation of rate-setting restricts how quickly insurers can update rates, so this is a reasonable approximation in the short run.

To simulate the effect of experience rating, I change the expected premium increase from the individual's subjective belief to the strict statutory rates used in North Carolina: 60% for an accident and 340% for driving with a BAC over 0.08% (North Carolina Department of Insurance, 2010). Without changing the distribution of premium in the first stage, this simulation measures the effect of experience rating on the probability of drinking and driving—i.e., moral hazard.

2.7.2 Calculating welfare under a premium change

The compensating variation for each individual under a policy change is:

$$CV_i = \frac{\ln\left(\sum_j \exp v_{ij1}^1\right) - \ln\left(\sum_j \exp v_{ij1}^0\right)}{\alpha} \tag{2.13}$$

where v_{ij1}^0 is individual i's conditional value function for insurance policy j under the status quo, v_{ij1}^1 is individual i's conditional value function for insurance policy j under the counterfactual prices, and α is the marginal utility of income (Small and Rosen, 1981). Assuming the insurer sets premiums at an actuarially fair value plus a competitive markup and is indifferent between the two policies, the net change in welfare before accounting for the externality is $\Delta W = \sum_i CV_i$.

I use average premiums for uninsured motorists insurance (UM) to estimate a lower bound on the monetized value of the externality from insurance. The UM premium represents a lower bound on the externality for two reasons. First, individuals would not voluntarily buy UM if there were no risk from other drivers being uninsured, and buying UM reduces consumption of other goods. Second, although UM in principle eliminates the financial risk from an accident caused by an uninsured driver, it typically does not provide complete coverage, particularly of non-pecuniary loss, because the liability limits are relatively low.[24] I estimate the average UM premium following a change in the share of uninsured motorists with the following expression:

$$\bar{P}_2^{UM} = \frac{TC_2^{UM}}{N_2^{insured}} = \frac{TC_1^{UM}(s_2/s_1)}{(1-s_2)N_2} = \frac{\bar{P}_1^{UM}(1-s_1)N_1(s_2/s_1)}{(1-s_2)N_2} \tag{2.14}$$

where TC_t^{UM} represents the total cost of covering uninsured motorists claims in period t, N_t is the total number of drivers in period t, and s_t is the share of uninsured drivers in period t. Assuming the total number of drivers does not change, $N_1 = N_2$,

[24] Most states now have mandatory UM insurance. Compelling drivers to buy UM involves a welfare loss in its own right. And persons compelled to buy this coverage are likely to purchase at the state's minimum liability limits, which are low. In addition, to get much compensation for non-pecuniary loss, one probably has to threaten to go to trial and this is costly to a plaintiff. Although UM has some properties of third-party insurance, it is really first-party insurance. Claims about one's own pain and suffering, etc. could easily be exaggerated. Auto liability insurance payment per paid claim is far less than for other third-party lines of insurance such as medical malpractice. This is documented at length in Bovbjerg et al. (1991). The vast majority of cases are settled out of court. By inference, payment for non-pecuniary loss must be much less for auto than for other third-party lines. Finally, with UM, especially compulsory UM, drivers may take less care in their driving. If UM had high liability limits, well above medical expenditure, drivers would have less incentive to exercise caution. Such moral hazard in itself generates an externality.

Table 2.6: Effects of counterfactual pricing policies

		Counterfactuals		
	Actual	Risk-rating	Community rating	Experience rating
Distribution of price changes				
5th percentile	—	-$129	-$486	-$15
25th percentile	—	-$129	-$94	$0
50th percentile	—	-$129	$18	$5
75th percentile	—	$175	$103	$37
95th percentile	—	$175	$417	$360
Welfare ($/person-year)	$0.00	$11.68	$2.16	-$17.36
Uninsured drivers				
Share of total drivers	0.04	0.05	0.04	0.05
Mean risk score	1.39	1.39	1.08	1.39
Mean risk preference	1.97	2.01	1.95	1.97
Drinker-drivers				
Share of total drivers	0.18	0.18	0.18	0.17
Mean risk score	1.07	1.07	1.08	1.07
Mean risk preference	1.55	1.55	1.55	1.57

Notes: The table reports the results of three counterfactual policy simulations on insurance choice, the decision to drink and drive, and social welfare. Welfare under the status quo is normalized to zero. "Share of total drivers" is the mean choice probability across individuals of being uninsured and choosing to drink and drive, respectively. "Mean risk score" and "Mean risk preference" are the means of the individuals' risk scores and risk preference factors, weighted by their choice probabilities for being uninsured and for drinking and driving. An individual's risk score is a linear combination of his individual risk factors (accidents, speeding tickets, DWI arrests, miles driven, drives to work), weighted by the implied hedonic prices and normalized so that the mean risk score in the sample is 1 and the standard deviation is 1. The risk preference factor is increasing in risk tolerance and is described in Section 2.4.

the counterfactual premium is a function of only the baseline average UM premium, which I obtain from National Association of Insurance Commissioners (2014), and the baseline and counterfactual shares of uninsured drivers from the data and estimation.

2.7.3 Counterfactual results

When premiums increase for drinker-drivers under risk-rating, the share of uninsured drivers increases from 4% to 5% (Table 2.6). The drivers who drop coverage are just as high-risk as the previously uninsured (mean observable risk score=1.39), but are more risk-tolerant, as evidenced by the increase in mean risk preference of uninsured drivers from 1.97 to 2.01.[25]

The simulated welfare gain from risk-rating is \$11.68 per person per year, or about 1% of the mean annual premium. However, this does not account for the externality that uninsured drivers impose on society. Assuming insured drivers internalize this externality by purchasing uninsured motorists' insurance (UM), the welfare gain from risk-rating disappears. If the mean UM premium per policy for a year is \$154 under the status quo (National Association of Insurance Commissioners, 2014), then the mean UM premium with 7% more drivers uninsured (the percentage increase before rounding the uninsured driver shares) is \$165, an increase of \$11.02, almost completely offsetting any welfare gain from risk-rating. One could argue that uninsured drivers will drive less; however, the newly uninsured drivers in the counterfactual scenario are just as high-risk, but even *more* risk tolerant than the previously uninsured, suggesting that they are at least as likely to drive without a license. Research also shows that a high percentage of drivers continue to drive even with license suspensions (Voas, Tippetts and McKnight, 2010; McCartt, Geary and Berning, 2003).

The most striking effect of community rating is the significant decrease in the average risk of the uninsured, from 1.39 to 1.08, where the mean risk score in the entire sample is 1. High-risk drivers buy insurance because their premiums are lower, and low-risk drivers, who now subsidize high-risk drivers, tend to drop coverage. How-

[25] An individual's risk score is a linear combination of the individual-level risk factors in the hedonic regression (accidents, speeding tickets, DWI arrests, miles driven, drives to work), weighted by the implied hedonic prices and normalized so that the mean risk score in the sample is 1 and the standard deviation is 1. A higher risk score implies a greater probability of an accident. An individual's risk preference factor is derived from a factor analysis on a set of observed variables related to risk preference and normalized so that the mean and standard deviation are both 1. Higher risk preference values indicate that the individual is more risk tolerant and lower values that he is more risk-averse (see Section 2.4).

ever, the low-risk drivers who drop coverage are similar to the previously uninsured drivers in terms of risk preferences. Very risk-averse drivers, even if they are low-risk, do not drop coverage. This keeps the market from unraveling. In addition, there is no detectable increase the share of drinker-drivers, despite increases in insurance among high-risk drivers (Tennyson, 2010). Overall, community rating leads to a mean welfare gain of $2.16 per person per year, without accounting for social welfare gains from having fewer high-risk, uninsured drivers.

Experience rating has a small effect on the probability of drinking and driving: the share of drinker-drivers decreases from 18% to 17%, with slightly more risk tolerant individuals choosing to continue drinking and driving. This modest effect is driven by the small marginal effect of pecuniary consequences on the probability of drinking and driving. Subjective premium increases following DWI arrest are also higher among drinker-drivers than non-drinker-drivers, which further attenuates any deterrent effect from experience rating.[26]

Although the mean welfare effects of these policies are modest, they have substantial distributional implications for consumer welfare (Figure 2.3). The individual welfare gains and losses are as high as 20% of mean annual premium under risk-rating and 80% of mean annual premium under community rating. High-risk drivers are worse off when premiums are based on private information and better off under community rating. Low-risk drivers are better off when private information is used and worse off under community rating.

[26] Hansen (2015) provides quasi-experimental evidence that drinking and driving recidivism is sensitive to the blood-alcohol content (BAC) level that the individual had at the time of the arrest. He argues that this is most likely due to a deterrent effect from more severe sanctions and punishments at higher BAC levels. However, he cannot disentangle the types of punishment that deter future drinking and driving. Pecuniary consequences likely play a role, but given the importance of preferences and non-pecuniary loss for drinking and driving behavior in my model, another plausible explanation for his finding is that an individual with a higher BAC level experiences more severe criminal punishment and updates his belief about how much a DWI arrest would affect his life.

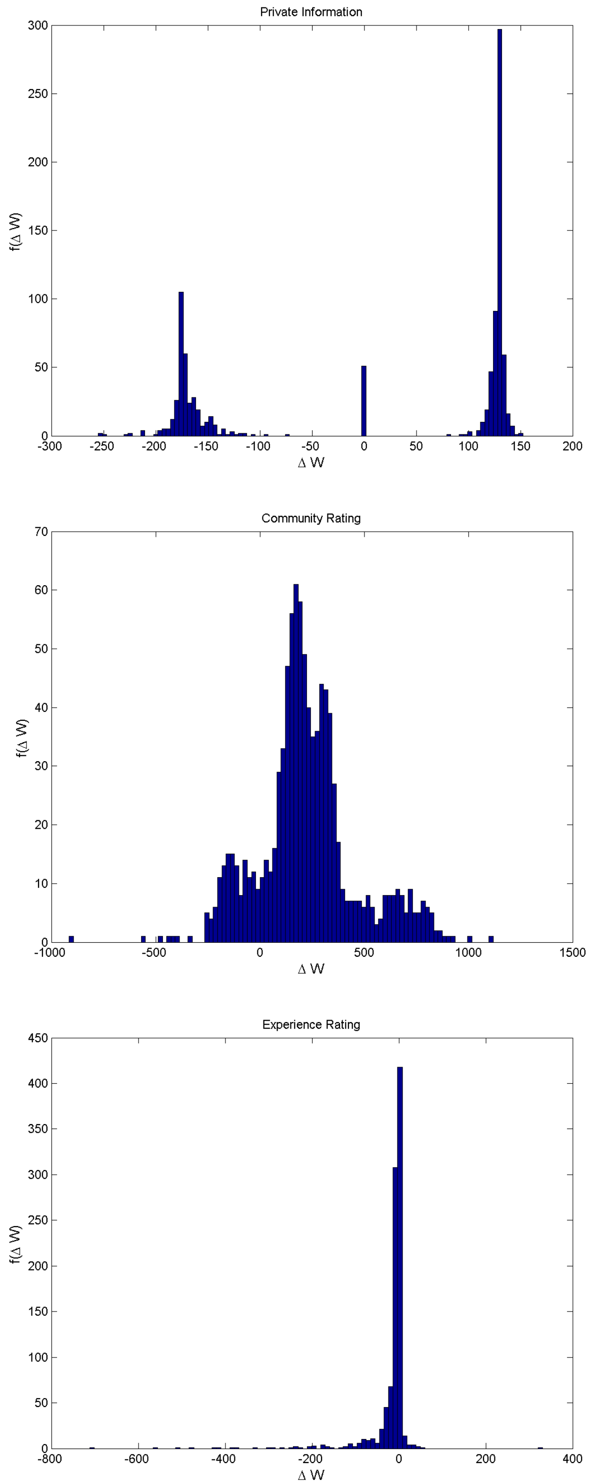

FIGURE 2.3: Distribution of welfare under counterfactual pricing policies

2.8 Conclusion

The theoretical importance of asymmetric information in insurance markets is widely recognized (e.g., Rothschild and Stiglitz, 1976). However, the degree of inefficiency caused by asymmetric information and mechanisms for overcoming it, particularly when individuals have multiple dimensions of private information, are less clear (e.g., Finkelstein and McGarry, 2006). This study shows that pricing policies that change the degree of asymmetric information between insurers and consumers can have surprising effects in a market with advantageous selection. In particular, improving information about risk is not a catch-all solution to efficiency in insurance markets.

Adjusting automobile insurance premiums based on private information about drinking and driving behavior induces a 7% increase in the share of uninsured drivers, which creates a negative externality on society that offsets the small welfare gain from eliminating residual adverse selection. This occurs because heterogeneity in risk preference induces advantageous selection: risk tolerant drivers are less likely to purchase insurance and more likely to drink and drive. In contrast, the market does not unravel in a substantial way under community rating, even though it increases the asymmetry of information between the insurer and the policyholder. Low-risk drivers remain in the market even when they face large price increases because they are risk averse and have a high valuation of insurance.

These pricing policies also have large distributional effects, which are potentially important to policymakers. Risky drivers experience substantial welfare gains and cautious drivers experience substantial welfare losses under community rating. Rewarding reckless driving and asking cautious drivers to pay for such behavior may be widely viewed as unfair. On the other hand, when risky drivers have to pay for reckless behavior, they are more likely to go uninsured, which imposes a substantial negative externality on other drivers.

Any decision to allow or restrict insurers from using information about individual risk in pricing has substantial effects on the market and the distribution of welfare, and, crucially, the effects depend on the sources of private information in the market. These findings are especially relevant as new technology becomes available to track very detailed driving behaviors, such as time of day, slamming on the breaks, etc., and policymakers are faced with important decisions about how to regulate the use of such information by insurers.

3

Physician Agency Under Multiple Payers

One of the most interesting features of the U.S. healthcare system is that a physician can be paid different amounts for providing the exact same service to two different patients. This happens because prices are set independently between payer and provider pairs—through bilateral negotiations, take-it-or-leave-it offers, or unilateral decisions by the government—which leads to different prices for patients who appear identical but are insured by different payers. How do individual physicians respond to the financial incentive that arises from these price differences across insurers?

This question is extremely relevant in the current healthcare environment. In a 2013 report, the Institute of Medicine noted that there is not enough evidence for which payment strategies will best improve quality of care at a lower cost (Institute of Medicine, 2013). Implicit in this statement is the idea that physicians respond to financial incentives. This view is widespread among private and public payers alike, as evidenced by continual efforts to implement payment strategies that change provider incentives. There is empirical evidence that physicians respond to prices on the extensive margin and in other dimensions (e.g., Bisgaier and Rhodes, 2011), but the existing research has been constrained by data from a single, often public, payer.

In order to effectively reform how we pay doctors in the U.S., we need a realistic understanding of how financial incentives work in the multi-payer healthcare system.

I estimate the effect of prices on treatment choice using the Massachusetts All-Payer Claims Database (MA APCD), which contains medical claims for the near-universe of patients under 65 in Massachusetts from 2009-2012. The key features of the MA APCD that make it ideal for this study are: (1) it compiles claims across the many private and public payers in the state, (2) it includes the actual prices paid for each service, and (3) it includes unique identifiers for physicians, payers, and patients. Prior studies have largely relied on aggregate, hospital-level data and price variation induced by changes in payments from a single—usually public—payer (e.g., Tai-Seale, Rice and Stearns, 1998; Gruber, Kim and Mayzlin, 1999; Clemens and Gottlieb, 2014). Some early studies also surveyed physicians to obtain prices for multiple payers (e.g., Sloan, Mitchell and Cromwell, 1978); however, these data did not include prices for the entire market, and they were subject to usual selection and reporting biases of surveys. The MA APCD allows me to study responses to fee changes at the physician level with complete information about all the prices in the market.

Cesarean sections offer a natural case study for a few reasons. First, childbirth is the most common reason for inpatient admission among privately insured, Medicaid, and uninsured patients (Pfuntner, Wier and Stocks, 2013). While the volume of cases alone makes this a relevant topic, it is even more important to understand the effect of cross-payer financial incentives given the high variation in insurance coverage of birth patients. Second, the C-section rate in the U.S. rose quickly over time to the current rate of 32 percent, which has raised questions about whether the marginal C-sections being performed are medically necessary (e.g., Caughey et al., 2014; MacDorman, Menacker and Declercq, 2008). The WHO estimated that over half a million C-sections performed in the U.S. in 2008 were clinically unnecessary

(Gibbons et al., 2010), and Molina et al. (2015) find that a C-section rate of approximately 19 percent is associated with improved health outcomes. This study does not attempt to measure over-utilization; rather, understanding the effect of prices on physician decisions is simply a critical step in designing incentives to lower the C-section rate without adversely affecting outcomes. Third, C-sections have received a lot of attention in the health economics literature (e.g., Gruber and Owings, 1996; Gruber, Kim and Mayzlin, 1999; Grant, 2009). Focusing on C-sections allows me to compare my findings—using new, individual-level, multi-payer data—to existing literature on financial incentives in healthcare.

My identification strategy relies on variation in the timing of contract renewals between physicians and insurers. The details of contracts between insurers and physicians are not publicly available, so I deduce the contract changes from the observed price data. The contract structures determined by my algorithm align with industry standards—in particular, all the physician contracts for a given payer tend to be renewed annually in the same month each year. In addition, I find that the three largest payers in Massachusetts renew their contracts with physicians in different quarters of the year. This variation across payers in the month of the renewal allows me to isolate the effect of one payer's price change, holding the others' constant.

I use a difference-in-differences approach to analyze the causal effect of a relative price change of C-sections compared to vaginal deliveries on the probability of C-section delivery, using five major contract change dates across three different large payers. I also test for heterogeneous effects across different patient risk groups and explore the effect on coding intensity, which provide insight into the mechanisms behind how physicians respond to price changes.

Physicians treat patients insured by different payers differentially following a relative price change. In particular, a patient is less likely to have a C-section when the relative price the doctor receives from her insurer for a C-section decreases. This effect is concentrated among patients classified as medium-risk based on factors observed before the delivery. In other words, physicians respond to prices when they are treating patients who are not clearly indicated to receive one type of treatment or another. This result is intuitive and aligns with the fact that obstetricians face high medical malpractice risk compared to other specialities (e.g., Studdert et al., 2016), and therefore seem unlikely to substantially deviate from the clinically accepted treatment path (when such a path is clearly indicated) in response to price.

The results of this study are particularly relevant for public payers, such as state Medicaid programs, that unilaterally set prices for their insured population. It is important that prices not be too high, given pressures on public budgets. Yet, this paper provides evidence that setting prices too low can result not just in barriers to access for low income patients, as shown by Sloan, Mitchell and Cromwell (1978); Zuckerman et al. (2004); Decker (2012); Sonchak (2015); Polsky et al. (2015) and others, but also differential treatment. This paper also provides new evidence that such differential treatment generalizes to patients insured by private payers. One potential policy intervention is all-payer rate setting, which could reduce treatment disparities and alleviate the administrative burden of pricing for payers and providers.

Section 3.1 provides background on the different pricing mechanisms for physician services in the U.S., and how they vary across payers. Section 3.2 introduces a conceptual framework for analyzing physician responses to prices in a multi-payer setting. Section 3.3 describes the data and sample that I use to estimate the empirical model outlined in Section 3.4. Section 3.5 presents the estimation results and Section 3.6 concludes.

3.1 Background

3.1.1 Multiple payers and physician payments

Prices paid by public and private payers for physician services are determined separately but are often closely linked. Both public and private payers have, over time, reimbursed physicians under a range of payment structures, from "reasonable and customary" charges determined by physicians to prospective payment agreed upon by both parties, and both are currently experimenting with new payment arrangements, such as value-based purchasing and bundled payments. However, there are important differences between payers and institutional features that lead to variation in prices across payers.

Public payers—government Medicare and Medicaid plans, in particular—set prices unilaterally. Since 1989, the Centers for Medicare and Medicaid Services (CMS) has used a physician fee schedule to pay providers (Centers for Medicare & Medicaid Services, 2013). Under the fee schedule, prices are determined by the product of resource-based relative value scale (RBRVS) weights for each procedure, a conversion factor, and geographic adjustment factors. The conversion factor converts the weights into dollar amounts and is typically updated annually by CMS, while the RBRVS weights are updated periodically to reflect technological change (American Medical Association, 2016). The fee schedule was introduced to reduce the administrative complexity of payment based on charges and to align the financial incentives of physicians with the resource burden required for different treatments (Hsiao et al., 1988).

Private payers negotiate with providers over prices or make take-it-or-leave-it offers.[1] In both cases, prices are set independently between each payer-provider pair, often on an annual basis. In negotiations over physician payments, a provider is

[1] This includes privately managed Medicare and Medicaid plans.

typically a physician group. The details of private payers' contracts with providers are proprietary, but both industry experts and recent research suggests that private payers often use CMS's price-setting mechanisms—in particular, Clemens, Gottlieb and Molnar (2015) find that 75% of private prices are based on the Medicare professional fee schedule. Variation in prices still arises, however, because payers negotiate or set different conversion factors.

Despite the use of similar pricing mechanisms, there can be wide dispersion in price levels across payers. First, government prices tend to be lower than private prices (e.g., Miller, Zuckerman and Gates, 1993; Krause, Ukhanova and Revere, 2016). This is sometimes called "cost shifting" to reflect the theory that providers have to raise private prices to make up for losses in government and uninsured segments (Fox and Pickering, 2008; Coughlin et al., 2014); however, it is possible that the differences simply reflect price discrimination and bargaining power. Second, the bargaining power of payers and physician groups affects prices in the private segment of the market. For smaller physician groups, a payer is more likely to offer a standardized take-it-or-leave-it contract. For larger groups, prices vary with the relative bargaining power of the payer and the physician group—we see higher prices in areas with more concentrated physician groups, and lower prices when insurers have a higher market share (e.g., Austin and Baker, 2015; Roberts, Chernew and McWilliams, 2017).

The central question of this paper is, are physicians' decisions about how to treat patients affected by this price dispersion—and if so, how does the physicians' behavior change with policy-relevant variables, including relative price changes and patient characteristics? The next section provides a conceptual framework for thinking about this question.

3.2 Conceptual Framework

3.2.1 The physician as the decision-maker

I assume the physician is the sole decision-maker for determining methods of diagnosis and treatment. This has been the norm in the health economics literature to date (Chandra, Cutler and Song, 2011), and it simplifies my analysis and interpretation of the results. Although the patient-empowerment revolution beginning in the late 20$^{\text{th}}$ century popularized the idea that patients should be involved in medical decisions (Ubel, 2012), recent work continues to provide evidence of physician paternalism rather than joint decision-making for common medical decisions (e.g., Fowler, Gerstein and Barry, 2013). In the case of deliveries, a woman likely searches for an obstetrician and has an opportunity to express her delivery preferences *ex ante*; however, at the time of the delivery, the obstetrician is ultimately responsible for making decisions.

Economic theory provides further motivation for this assumption because medical care has been used as an example of a credence good (e.g., Dulleck and Kerschbamer, 2006). A credence good is a good for which the consumer does not know the level of quality that she needs, so she pays an expert to provide both the diagnosis and treatment. As a result, it may be difficult or impossible for the consumer to ever know whether she received the appropriate treatment. For example, at the time of delivery if a vaginal birth is not progressing (dystocia), then the obstetrician may use his expertise to determine that a Cesarean section is medically indicated and perform the surgery. However, if the patient did not need a Cesarean section, she may never know the difference.[2]

In theory, the supplier of a credence good will report truthfully and provide the optimal level of treatment if the following four assumptions are satisfied: (1) homoge-

[2] If she needed but did not receive a Cesarean section, she may immediately experience adverse outcomes. In this state, the appropriate course of treatment is arguably observable *ex post*.

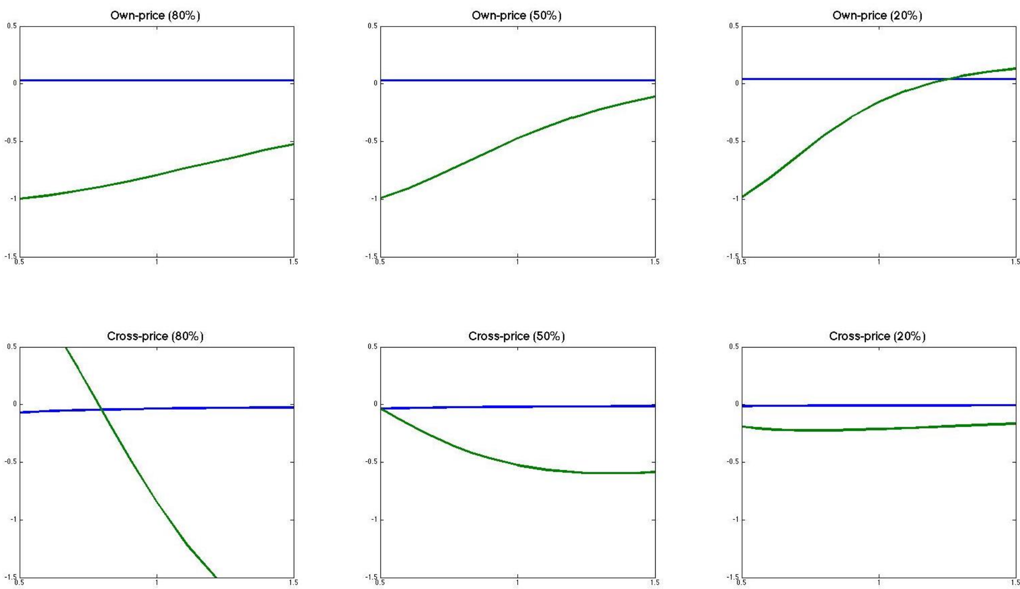

FIGURE 3.1: Own- and cross-price elasticities by payer market share

nous consumers; (2) commitment to receive the prescribed treatment conditional on diagnosis; (3) verifiability of treatment quality; and (4) observability of and liability for the outcome from treatment (Dulleck and Kerschbamer, 2006). In practice, however, the first assumption is always violated, and the last two are likely to be violated in many clinical situations, as in the example above. Thus, an obstetrician may perform medically unnecessary C-sections if he has a financial incentive to do so.

3.2.2 Physician response to prices under multiple payers

McGuire and Pauly (1991) developed a model of physician behavior in response to different prices from multiple payers. Their model unified a previously divisive literature on physician supply by showing that a physician's optimization problem does not have to involve targeting a certain level of income in order to rationalize "inducing" patient demand for treatment.

The key insights from the model are that the relative prices and market shares of different payers affect how physician behavior—the quantity of services provided to patients of each payer—changes when one payer's price changes. In particular,

The output from the McGuire and Pauly (1991) model is a set of elasticities that describe how the quantity of healthcare services (X) provided to patients with payer i respond to changes in the net revenue or margin (m) from payer j, where $i, j \in \{1, 2\}$. There are four elasticities: an own- and a cross-price elasticity for each of the two scenarios of profit maximization (no income effects) and literal target income. Each of these elasticities is a non-linear function of seven parameters, all of which have economic interpretations, and some of which are observable in the data. The two parameters that are directly observable are the ratio of the payers' market shares (σ) and the ratio of the payer's margins (ρ). These are key parameters of interest for policy.

Figure 3.1 replicates plots from McGuire and Pauly (1991) of the own- and cross-price elasticities against the relative market shares when the base payer has three different market shares: 80%, 50%, and 20%. The blue lines show the elasticities under profit maximization, and the green lines show the elasticities under a literal target income. Actual behavior likely lies somewhere in between and depends on the strength of the income and substitution effects—when income effects completely dominate, the physician pursues a target income; when substitution effects completely dominate, the physician maximizes profit.

A practical concern with this model is that physicians have to know something about the relative prices across patients in order to respond to them. While finding robust empirical evidence of a response to price would allow us to deduce that they have such information, we can also point to a few reasons why this is especially likely to be true for the case of C-sections. First, obstetricians have many interactions with patients before the delivery, providing ample time to learn about the patient's insurance carrier. Second, general price trends in obstetrics are well documented—C-sections have historically been reimbursed above vaginal deliveries, and private payers reimburse well above Medicaid. Third, the high volume of births gives providers an incentive to pay attention to prices for deliveries. Fourth, deliveries may be priced as a separate "carve-out" from the Medicare fee schedule due to their high volume and relevance only for the under-65 population, making it more likely that physicians are aware of negotiated prices for deliveries. Finally, there is general evidence that physicians are aware of prices and are in control of treatment choices. In a survey of physicians, Tilburt et al. (2013) find that a majority of physicians are "aware of the costs of the tests/treatments [they] recommend" (76%) and agree that they "need to take a more prominent role in limiting use of unnecessary tests" (89%) and "should adhere to clinical guidelines that discourage the use of marginally beneficial care" (79%).

3.3 Data

3.3.1 Data Description

The empirical analysis requires data on prices and procedures performed by physicians on patients with different payers. The Massachusetts All-Payer Claims Database (MA APCD) is an ideal data set for this purpose because it compiles administrative healthcare claims data for the under-65 population from all the different payers in the state of Massachusetts—something that had not been done previously, at least not in a way that was accessible to researchers.[3] The MA APCD was developed to support research in the public interest and to streamline state data submission requirements for payers, and it is maintained by the Center for Health Information and Analysis in Massachusetts.[4]

[3] Excluded are worker's compensation, TRICARE and the Veterans Health Administration, Federal Employees Health Benefit Plan, private insurers with fewer than 1,000 lives, and the uninsured, unless they are enrolled in the Commonwealth's Health Safety Net.

[4] Further details are available at http://www.chiamass.gov/ma-apcd/.

My analysis uses physician claim lines from the medical claims file of the MA APCD for the period 2009-2012. An observation is one "line" of a physician claim, which contains information about a specific service provided during a patient's visit to a healthcare provider. The claim line includes identifiers for the patient, payer(s), and provider(s) involved in the service; Current Procedural Terminology (CPT) codes identifying the service provided; International Classification of Diseases (ICD-9) codes that indicate the reason(s) for the service; and the dollar amounts billed, allowed, and paid to the physician for the service.

Price is measured by the allowed amount on the physician claim line. This is the rate that is specified in the contract between the payer and physician, and it is the amount that the physician can expect to receive for a service or procedure. A portion of the allowed amount may be paid by the patient if the patient's insurance plan includes cost-sharing (e.g., deductible, copay, or coinsurance); the remaining portion is then paid by the insurer. This breakdown is observable on the claim line; however, I assume the physician responds only to the overall allowed amount. In addition, the allowed amount on the physician claim line does not include payments to the hospital for use of the facility. By excluding these payments, I assume the physician is not responsive to facility fees. I also assume that physicians are not allowed to balance-bill patients for charges above the payer's allowed amount, which is akin to assuming all patients visit in-network providers.

In addition to medical claim lines, the MA APCD includes provider, product, and member eligibility files, which can be linked to the medical claims file. The provider file includes demographic and practice information about the servicing provider: name, gender, date of birth, specialty, National Provider Identifier (NPI), practice site address and phone number, and some practice characteristics. The member file contains demographic information about the patient: zip code, birth month and year, relation to insurance plan subscriber (e.g., spouse, child). The product file contains

Table 3.1: Sample selection procedure

All claim lines with delivery CPT code	338,463
Standardization	
Keep lines with no modifiers	282,741
Keep lines delivered in an inpatient hospital (site=21)	252,612
Keep lines with allowed amount > $500	245,129
Keep lines covered by primary insurance	238,943
Keep lines covered by top 15 payers by frequency	235,844
Estimation	
Keep lines not missing Physician NPI	210,159
Total number of physician NPIs	*3,604*
Keep lines with physician contract prices	176,666
Physician NPIs with a 12-month run & 3 births/mo	*651*

information about the insurance plan: product type (e.g., HMO, POS), contract type (single, family), coverage type (self-funded, individual, small group), and plan characteristics (e.g., annual deductible).

The provider file is essential to my analysis because it allows me to link providers across payers using the National Provider Identifier (NPI). Payers may use different or even internally-developed IDs to identify providers, and as a result, the medical claims file does not contain a unique physician identifier that is reliable and consistent across payers. However, the file does contain identifiers that allow me to link claims to the physician file in order to lookup and use the NPI as a unique physician identifier across all payers (see Appendix F).

3.3.2 Analysis Sample

The full sample of MA APCD data from 2009-2012 includes 338,463 delivery claim lines.[5] These lines are identified by CPT codes beginning with 594–vaginal delivery, 595–Cesarean section delivery, and 596–attempted or successful vaginal delivery

[5] Claim lines flagged as not highest version are excluded.

following a previous Cesarean section delivery.

I apply a number of sample restrictions to remove duplicate lines and standardize the sample (Table 3.1). First, I exclude the 16 percent of claim lines that contain modifiers, which most often indicate that the claim is a secondary claim for an assistant-at-surgery. Including these lines would duplicate deliveries because the primary physician and the assistant-at-surgery bill separately for the delivery. Second, I exclude claim lines where the site of service is not an inpatient hospital because the risks, incentives, and patient engagement may be quite different from inpatient deliveries. Third, I exclude claim lines with allowed amounts less than $500, which likely include reversals or denied claims—84 percent of the excluded lines in this category have allowed amounts less than or equal to zero. Fourth, I retain only those lines where the payer who submitted the claim is the primary payer. Patients can have multiple sources of insurance coverage, and including claims submitted by a patient's secondary or tertiary payer would duplicate the delivery without providing additional information, since price is based on the primary payer's contract. Finally, I restrict the sample to the top fifteen payers by frequency of delivery claim lines. These payers submit 98.8 percent of all delivery claim lines; excluded payers are generally out-of-state.

I also restrict the sample to make estimation tractable. First, I exclude claim lines that are missing a reliable physician NPI. It is difficult to identify individual providers in healthcare claims data, and it is even harder to do so across claims submitted by different payers. I outline my procedure for identifying individual physicians in the MA APCD in Appendix F. Using this procedure, I obtain 3,604 unique physician NPIs. Second, I only include claim lines for which I can identify the contracted prices for vaginal and Cesarean section deliveries. To reliably determine contracted prices, I set a threshold for physician-level claim volume of at least three births per month for a consecutive twelve-month period. I chose this threshold by inspecting delivery

Table 3.2: Summary statistics on payer concentration

	Overall	By physician			
	Mean	Mean	Std. dev.	Min	Max
	(1)	(2)	(3)	(4)	(5)
Panel A. Market-level concentration					
Number of payers	14	6.97	2.57	1	13
CR(4)	0.63	0.86	0.11	0.57	1
HHI	1,459	2,944	1,458	1,230	10,000
Panel B. Individual payer shares					
Private Payer #1	0.26	0.27	0.14	0.01	1
MassHealth	0.19	0.25	0.19	0.01	1
Other Medicaid (N=2)	0.18	0.24	0.18	0.01	0.80
Private Payer #2	0.12	0.15	0.12	0.01	1
Private Payer #3	0.08	0.11	0.09	0.01	0.60
Other Private (N=8)	0.17	0.18	0.13	0.01	1
Observations	176,666	610			

volume patterns by NPI and finding the volume that maximized the number of NPIs included in the sample without losing continuity of observations within each NPI. The latter is necessary for determining contract changes and for my difference-in-differences estimation strategy. My complete procedure for determining contracted prices is explained in Section 3.4.1. The full sample for the empirical analysis contains 176,666 delivery lines provided by 651 unique physicians.

3.3.3 Descriptive statistics

A physician sees patients from seven different payers, on average, and 86 percent of deliveries are insured by the top four payers for any one physician (Table 3.2, Panel A). The mean HHI of payer concentration faced by any one physician is nearly 3,000—well above the 2,500 threshold for a high concentration. The fact that payer concentration is high makes it more likely that a physician will respond to prices

Table 3.3: Summary statistics on prices and C-section rates

	Overall	By physician			
	Mean	Mean	Std. dev.	Min	Max
	(1)	(2)	(3)	(4)	(5)
Panel A. Prices					
Vaginal birth ($)	3,011	2,881	563	2,086	3,925
Cesarean section ($)	3,236	3,127	567	2,355	4,134
Difference ($)	225	246	125	-170	617
Panel B. Deliveries					
C-section rate	0.36	0.37	0.09	0	1
Deliveries per month	3,681	36	28	5	320
Observations	176,666	610			

because it is easier for him to know the prices for individual patients than if he saw an equal number of patients from many different payers.

The overall shares of deliveries insured by the various payers in Massachusetts varies widely. Over one third of deliveries in the sample are covered by the state's Medicaid program called MassHealth (19 percent) or by privately-administered Medicaid managed care plans (18 percent). The top three private payers by frequency of delivery claims insure 46 percent of all deliveries, with the top private payer insuring 26 percent, while eight private payers cover the remaining 17 percent of delivery claims. There is also variation in payer shares at the physician level. Although the mean payer shares across the 610 physicians are similar to the overall shares in the state (col. 2 vs. col. 1), for each payer there are physicians who see very low or very high shares of that payer's patients (cols. 4-5).

There is substantial variation across payers in the price of a delivery (Table 3.3, Panel A). The mean reimbursement for vaginal delivery in the whole sample is $3,011, and the mean for C-section delivery is $3,236, a difference of only $225. However, the range of differences is much wider, from -$2,610 to $4,699. The range of reimbursement rates for both vaginal delivery and C-section delivery are similar, from about $1,100 to $7,000.

To get a sense of how much of this variation any individual physician actually faces, I compute the coefficients of variation in allowed amounts for the full sample and compare them to the coefficients of variation computed at the individual physician level and averaged across physicians. The average physician faces just under two-thirds of the total variation in allowed amounts for each procedure: 0.19 versus 0.29 for vaginal birth, and 0.17 versus 0.27 for C-section (not shown). In other words, there is significant within-physician variation in prices in addition to variation across the entire market, the latter of which has been well documented (e.g., Cooper et al., 2015).

The mean C-section rate in the entire sample is 0.34, which is close to the national average of 32.2 (Centers for Disease Control and Prevention, 2016). Among physicians, there is substantial variation in C-section rates: the mean is 0.31 and the standard deviation is 0.19. The average physician in the sample performs 36 deliveries per month, which is likely higher than the population average because I dropped physicians with low volume for the pricing algorithm. It is also possible that some of the Physician NPIs in my sample do not identify individual physicians. The most likely alternative is that an NPI is used by multiple physicians or an entire physician group. If this were the case, it would preclude me from controlling for individual physician heterogeneity, but it does not materially change my approach or interpretation of the results.

3.4 Empirical Strategy

I estimate the effect of relative prices on physician treatment choice using a difference-in-differences framework. This approach, together with the observed price variation in the data, allows me to isolate the effect of a single payer's price change on physician behavior. The estimation requires two steps: (1) identifying changes in the contracted prices between physicians payers, and (2) estimating the difference-in-differences model.

3.4.1 Identifying contracted prices

The negotiated contracts between payers and physicians are proprietary and therefore not directly observable, but I can deduce the timing and level of price changes based on the prices I observe in the data. The basic features of prices that I expect for a given physician-payer pair are: (1) the prices for both vaginal delivery and Cesarean section delivery change on the same date; (2) prices do not change more frequently than once per quarter; and (3) prices change on the same date(s) each year. With these features in mind, I wrote an algorithm to determine dates on which contracts changed, and the level of the price changes. The algorithm is described in detail in Appendix G.

The variable of interest is the change in the relative price of a C-section compared to a vaginal delivery for each payer-provider pair in each month. This relative price change is given by

$$\Delta RP_{jkt} = \left(P^{CS}_{jkt} - P^{VB}_{jkt}\right) - \left(P^{CS}_{jk(t-1)} - P^{VB}_{jk(t-1)}\right)$$

where j, k, and t index providers, payers, and months, and P^{CS} and P^{VB} denote the contracted physician allowed amount for C-section and vaginal delivery, respectively. The variable ΔRP_{jkt} should be zero except for months in which the contract changed.

A key insight from my algorithm for identifying contract changes is that the top three private payers in Massachusetts (by share of births) renew their contracts with physicians in different months of the year. This is evident in Figure 3.2, which plots non-zero values of ΔRP_{jkt} for the top three payers in the top, middle, and bottom panels, respectively. The estimated contract change months, indicated with red bars,

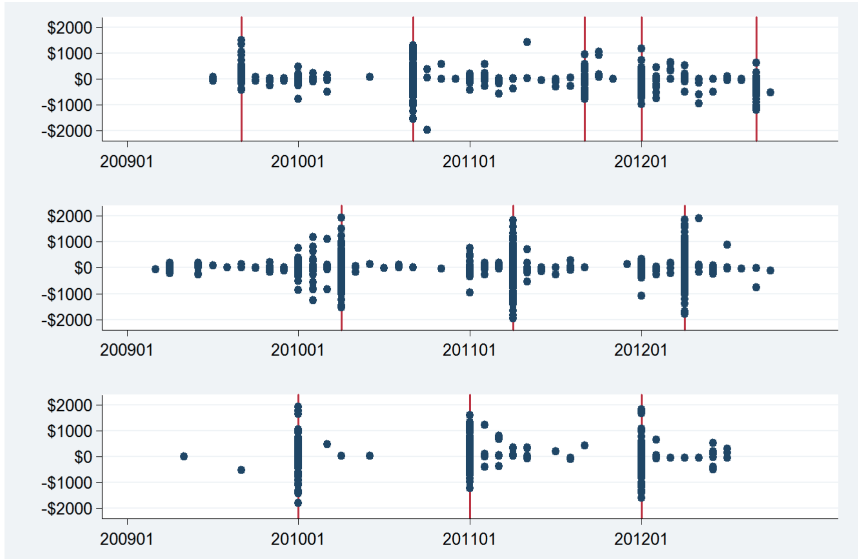

FIGURE 3.2: Contracted relative price changes over time for three large payers

are generally consistent over time for any given payer—i.e., within a panel—but vary across payers—i.e., across panels. The plots align with the prior expectation that a payer tends to renew most of its contracts once per year in the same month.

Another finding from identifying contract changes is that the prices paid by MassHealth for deliveries do not change over the sample period, 2009-2012. The MassHealth prices are also exactly the same for all physicians in the state. The price of a vaginal delivery is \$2,045 and the price of a C-section is \$2,310, which are generally the minimum prices physicians receive for these procedures in the sample. Because there is no empirical variation in these prices, I cannot do an event study on Medicaid price changes. However, the estimation results still have implications for Medicaid pricing if we assume physician behavior under private payer price changes can be generalized. In fact, we might expect the response to private price changes to be dampened by the difficulty of differentiating privately-insured patients from each other.

I use five of the observed contract changes in estimation, which are chosen based on two criteria. First, I only use events that occur after January 2010 because the

MA APCD started in 2009 and the data quality that year appears to be lower. Second, I do not use the four events that happened around January 2012 because multiple price changes occurred at once. This leaves five events between March 2010 and September 2012.

3.4.2 Difference-in-differences

To identify the effect of price on physician behavior, I estimate a difference-in-differences model of the delivery outcome for patient i:

$$\begin{aligned} Y_i = \beta_0 + \beta_1 PAYER_i * POST_i * DEC_i + \beta_2 PAYER_i * POST_i * INC_i \\ + \beta_3 PAYER_i + \beta_4 POST_i + \gamma_e + \alpha_j + \varepsilon_i \end{aligned} \tag{3.1}$$

where $Y_i = 1$ if patient i had a C-section delivery and 0 otherwise; DEC_i and INC_i are binary indicators that are 1 if the relative price change, ΔRP_{jkt}, was $\leqslant -\$200$ or $\geqslant \$200$, respectively; $PAYER_i$ and $POST_i$ are binary indicators that are 1 for the payer whose relative price changed and for claims incurred after the contract change date, respectively; γ_e is an event fixed effect; and α_j is a physician fixed effect. The coefficients of interest are β_1 and β_2, which give the difference-in-differences estimate of the effect of a payer's relative price decrease or increase on the probability of a C-section for a patient insured by that payer.

There are two key sources of identifying variation that allow me to estimate equation 3.1. First, there is variation in reimbursement across patients within physicians, which arises because physicians see patients with different insurers, and they each independently negotiate or offer take-it-or-leave-it contracts with the physician (or physician group) over allowed amounts, as discussed in Section 3.1. Second, there is variation in the timing of contract changes across payers. This variation arises because payers tend to renegotiate their contracts on different schedules, as illustrated in Figure 3.2.

In order to estimate equation 3.1, I modify the sample described in Section 3.3 in a few additional ways. First, I duplicate the sample for each of the five events, so the total number of observations becomes 883,330. Second, for each event, I delete all observations that lie outside of a three-month period on either side of the contract change date. I selected this bandwidth because it is the maximum time between any two events for the five events (the closest two occur on January 1 and April 1 in 2011). This reduces the sample size to 110,990, and removes most of the duplication of observations. Third, for each event, I remove all observations for physicians who experienced another contract change with any payer within the 6-month sample window. This reduces the sample size to 99,859, but only reduces the number of physicians from 610 to 606, which likely means that those four excluded physician NPIs, which were responsible for nearly 10% of the claims, were not good identifiers of individual physicians.

3.4.3 Heterogeneous effects

There are a number of variables that could affect the strength of the physician's response to a price change. In addition to the relative margins and the market shares discussed in Section 3.2, the elasticities that McGuire and Pauly (1991) derive also depend, naturally, on the disutility of inducing demand. One dimension across which this could vary is patient risk.

There are at least two reasons that the disutility of inducing demand (for either C-sections or vaginal deliveries) may be higher for patients who are especially high or low risk. First, if a patients' risk factors clearly indicate that one treatment path is best, then the moral cost of deviating from that is likely higher than if there is some clinical ambiguity that allows the physician to easily justify either treatment. Second, obstetricians face a high risk of medical malpractice lawsuits (e.g., Studdert et al., 2016). If the obstetrician deviates from the standard treatment path and the

Table 3.4: Risk regression

	Mean	Coefficient	Std. error	Marg. effect	Odds ratio
Prior C-section	0.14	6.49***	(0.084)	0.85	656.30
Over Age 35	0.04	0.09***	(0.035)	0.01	1.10
Malposition	0.06	3.26***	(0.032)	0.43	26.14
Disproportion	0.02	5.30***	(0.122)	0.69	200.20
Multiples	0.02	1.66***	(0.045)	0.22	5.24
Diabetes	0.00	1.29***	(0.159)	0.17	3.62
Obesity	0.00	0.92***	(0.172)	0.12	2.50
Constant	—	−1.37***	(0.008)	—	0.25
Observations	176,666				

patient has an adverse outcome, a lawsuit would likely be successful. Thus, there is both a moral cost and a financial cost of inducing demand, which is likely higher among patients for whom there is clear treatment path. Currie and MacLeod (2008) make a similar argument, and they find that C-sections are less responsive to medical malpractice laws when treating high-risk patients.

I assign patients to three risk categories—low, medium, and high—using the predicted probability of a C-section based on observed diagnosis codes. I first estimate the following model using a logistic regression:

$$P(Y_i = 1) = \Lambda(\beta_0 + \beta_1 \mathbf{X_i}) \tag{3.2}$$

where X_i is a vector of risk factors of the mother—prior C-section, over age 35, diabetes, and obsese—and risk factors of the fetus—malposition, disproportion, and multiples. I then use the estimated coefficients to predict the probability of a C-section based on these risk factors. Finally, I split patients into low, medium, and high risk categories based on observed breaks in the distribution of predicted probabilities.

All of the risk factors included in equation 3.2 are associated with an increase in the probability of a C-section (Table 3.4). Having a prior C-section is by far the most likely factor to lead to a C-section delivery, followed by disproportion and malposition of the fetus. The predicted probabilities range from 16 to 99 percent, with clear breaks in the distribution around 24 and 90 percent, which I use as the cutoff values for the three bins. About 78 percent of patients are low-risk, 6 percent are medium-risk, and 16 percent are high-risk.

I test for heterogeneous effects across risk level by including interactions between the difference-in-difference terms in equation 3.1 and the three risk categories. Given the discussion above, we expect physicians to be less responsive to price changes when treating low- and high-risk patients than medium-risk patients.

One potential problem with this approach is that diagnosis codes could be endogenous. Physicians may manipulate diagnosis codes in order to justify providing a certain type of treatment, and so the codes included in equation 3.2 might be correlated with omitted prices. To overcome this, I choose risk factors that are less ambiguous, are generally known before delivery (ICD-9 code families 650-659), and are therefore less likely to be used solely to justify treatment in response to price. In particular, I do not include diagnosis codes associated with long labor and other codes in the category of "complications occurring mainly in the course of delivery" (ICD-9 code families 660-669), which may be more subjective. One test of these assumptions is to perform the difference-in-differences analysis with each of the different coding families as the dependent variable. The results are generally consistent with my assumptions: long labor and abnormal forces are coded less frequently when prices go down, while the variables included in equation 3.2 do not respond in the same direction. I discuss these results in more detail in Section 3.5.

Another potential concern is that I use predicted probability of a C-section in a subsequent regression where C-section is the dependent variable. If there is no endogeneity problem in the risk regression, then the predicted value from that regression is an unbiased estimate of the probability of a C-section based only on the included

Table 3.5: C-section by price change and patient risk

	C-section
Relative price decrease	
Renewing payer × post-renewal	−0.011
Renewing payer × post-renewal × med-risk	−0.094***
Renewing payer × post-renewal × high-risk	0.010
Relative price increase	
Renewing payer × post-renewal	−0.037
Renewing payer × post-renewal × med-risk	0.029
Renewing payer × post-renewal × high-risk	−0.024
Controls	
Renewing payer	−0.005*
Post-renewal	−0.005**
Med-risk patient	0.496***
High-risk patient	0.801***
Event controls	Yes

risk factors. Including this predicted value directly in the subsequent regression may affect the standard errors, but should not lead to any bias in the estimated coefficients. In reality, I include a transformation of the predicted probability in the form of three large bins (low, medium, and high risk). To the extent that there is an endogeneity problem in equation 3.2, using these large bins should significantly reduce any correlation between prices and risk, assuming there is sufficient variation in predicted probabilities within each bin. I also perform a robustness check where I estimate the parameters of equation 3.2 on a sample of patients who are not included in the main analysis, and who deliver in months where no contract changed.

3.5 Results

3.5.1 Treatment choice

Physicians are less likely to perform a C-section on patients for whom the relative price of a C-section decreased, and this effect is concentrated among medium-risk

Table 3.6: Frequency of diagnoses following price change

		DD coefficients	
Dependent variable	Mean	Price decrease	Price increase
Count of diagnosis codes	1.935	−0.039***	0.151**
Long labor	0.007	−0.005***	−0.001
Malposition	0.073	0.013**	−0.018
Disproportion	0.029	0.008***	−0.017**
Any delivery complications	0.255	−0.033*	−0.001
Any pregnancy complications	0.172	0.003	0.020

patients (Table 3.5). The magnitude of the decrease in C-section utilization among medium-risk patients is 9.4 percentage points, or about 15%. The mean relative price decrease driving this result is -\$587, which is over twice the mean price difference across all physicians.

There is a slightly positive but not statistically significant increase in the C-section rate among medium-risk patients when the relative price increases. There are a few possible explanations for the asymmetry of the result. First, the mean price increase observed in the sample is 30% smaller in magnitude than the mean price decrease—\$416 versus \$587—which is a smaller incentive. Second, there are nearly fifty percent more occurrences of relative price decreases than increases in the estimation sample, providing more power to estimate the former coefficient. Third, there may not be much more potential to increase the C-section rate, given the already high rate. Finally, there could be other features of the contracts between physicians and insurers, such as utilization review, that we haven't accounted for but that are correlated with the price change and affect C-section rates. I explore this possibility in the next subsection.

3.5.2 Coding intensity

Using the difference-in-differences framework to test for changes in coding intensity across different diagnosis codes provides insight into a number of issues. First, as discussed in Section 3.4, it is a test for the endogeneity of coding. Second, it yields insight into exactly how providers respond to prices, particularly the certain types of deliveries that are on the margin. Third, it may give some clues about whether a simultaneous change in other contract dimensions could be driving the main results.

I estimate equation 3.1 for different diagnosis code families, where the dependent variable is whether or not any diagnosis code within that family was reported on the delivery claim line. The coefficients on $PAYER_i * POST_i * DEC_i$ and $PAYER_i * POST_i * INC_i$ represent the effect of a change in a payer's relative price decrease or increase on the probability that the relevant diagnosis is reported.

The probability of a diagnosis code could change for at least three reasons. First, there could be a true underlying change in patient risk. The difference-in-differences approach is intended to control for this. Second, if a certain diagnosis code is only reported when a C-section is performed—i.e., as justification for it—then when the probability of a C-section decreases, the probability of that code being reported will also decrease. Third, there could be some other change at the same time as the price change—such as increased scrutiny of C-sections through utilization review—that causes physicians to code more aggressively. This could also lower C-section utilization and affect coding through the second channel as well.

The key findings, reported in Table 3.6, are the following: the number of diagnosis codes reported on a delivery claim decreases in response to a relative price decrease and increases in response to a price increase; physicians are less likely to report any delivery complication following a price decrease, particularly long labor; malposition and disproportion are more likely to be coded following a price decrease and less likely following a price increase; and the coding of pregnancy complications is not responsive to delivery prices.

The first finding—that the number of codes responds to price changes—reflects that diagnosis codes are needed as justification for billing a C-section. The magnitude of the response is larger for a price increase than decrease, in contrast with the coefficients in Table 3.5. This provides additional support for the main finding that treatment choice responds to price, both increases and decreases, in the expected way. Furthermore, if other contract changes were driving these results, such changes would have to occur for physician-payer pairs with price increases and for those with price decreases. In other words, these types of non-price changes, such as value-based agreements and utilization review, would have to be highly correlated with price changes and widespread, which seems quite unlikely.

The probability of long labor being reported decreases by 70 percent on average when the relative price of a C-section decreases by \$587 (Table 3.6, row 2). Given the potentially subjective measurement of long labor (Thielking, 2015), it makes sense that these are the marginal types of C-sections that are responsive to price. Similarly, the probability that any delivery complications are reported decreases when the price decreases. This aligns with the assumption in Section 3.4 that this family of diagnosis codes is endogenous to price. In addition, there is no statistical effect on the coding of pregnancy complications, which further supports the use of risk factors known prior to delivery for estimating equation 3.2, and serves as a useful falsification test.

Malposition and disproportion respond to price changes in the opposite direction from long labor (Table 3.6, rows 3-4). When the relative price decreases, both are more likely to be coded; whereas when the relative price increases, they are less likely to be coded. The magnitudes of these effects relative to the mean and their statistical significance are smaller than for long labor; nevertheless, the results are a bit puzzling. We would expect malposition and disproportion not to respond to prices—or in this case, go down with the C-section rate—if they are reported objectively. It is hard to think of any reason why coding intensity of these complications would go up in response to lower prices, and vice versa. I turn again, then, to considering the possibility of other changes concurrent with the price change.

These coding patterns could reflect a response to increased utilization review by the payer—i.e., monitoring of claims and charts to identify over-utilization or mis-treatment. Assuming malposition and disproportion are more acceptable justifications for a C-section than long labor, the decrease in long labor and increase in the other two could be explained as physicians shifting their coding techniques. If this is true, then econometrically, we cannot disentangle the effect of utilization review from the effect of prices in the difference-in-differences estimates of the C-section probability.

However, there are arguments against the hypothesis that increased utilization review is driving the main results on treatment choice. First, the difference-in-differences estimates for malposition and disproportion are symmetric, at least in sign, for price increases and decreases. Yet, it is highly unlikely that price increases are associated with less intensive utilization review, which would have to be true for this hypothesis to hold. Second, the feedback mechanisms for physicians may not be as clear for utilization review as they are for prices, particularly if medical coders are the people who interact with utilization managers. A hospital in California reportedly took three years to lower its average C-section rate by 13 percent following "intense scrutiny" by a large insurer.[6] Third, utilization review is unlikely to be written directly into contracts and perfectly aligned with the contract change date.

Overall, these results show that physicians do respond to payer-specific prices when choosing whether or not to perform a C-section. The effect is concentrated

[6] https://www.theatlantic.com/health/archive/2015/05/how-one-hospital-reduced-unnecessary-c-sections/392924/

among patients who do not have a clear indication for treatment, likely those who are experiencing long labor.

3.6 Conclusion

This paper shows that treatment prices affect the allocation of healthcare services across patients in the multi-payer, U.S. healthcare system, at least in the case of obstetrics. In particular, a physician's choice of whether to perform a Cesarean section delivery depends on the prices paid by the patient's insurer, especially for medium-risk patients. This finding adds to the growing body of literature showing that physicians respond to financial incentives (e.g., Gruber, Kim and Mayzlin, 1999; Clemens and Gottlieb, 2014; Iizuka, 2012), and provides novel evidence that prices affect behavior toward privately insured patients and on the intensive (treatment) margin.

A central question still remains as to whether the variation in prices that we observe across payers leads to allocative efficiency. The argument for having multiple prices is that it allows for better matching of consumers to goods and services based on their preferences. In a normal market, people who have a high willingness-to-pay are more likely to get the good or service. In this market, however, the relative prices that drive physician behavior are set independently of patient preferences. The patient chooses an insurer, if she has a choice, taking into account high-level characteristics of the plan such as premium and breadth of the provider network. While these plan characteristics may be correlated with absolute prices, they are not necessarily correlated with the single price among thousands for a specific service, such as C-section, nor are they correlated with relative prices. In addition, these prices are generally hidden from consumers until after treatment, making it impossible for patients to directly respond to them. One way in which prices could steer a patient toward her preferred treatment is if she chooses a physician based on his history of providing more or less aggressive treatment and that history was driven by prices. However, the physician still faces different prices from different payers, is ultimately responsible for making the treatment decision, and—as the difference-in-differences results show—responds directly and immediately to price changes. For these reasons, negotiated prices are highly unlikely to reflect consumer preferences.

Despite not directly measuring welfare, this work still has important implications for policy. On the one hand, the results are encouraging for insurers and policymakers who aim to use prices as a policy instrument. If payers can negotiate lower physician prices for C-sections, the C-section rate will decrease, which is clearly an objective of some healthcare policy and medical organizations due to the increased cost and risk of complications from C-section delivery (e.g., Caughey et al., 2014; U.S. Department of Health and Human Services, 2010). On the other hand, this also implies that we cannot simply lower prices in order to lower healthcare costs without affecting the quantity of treatment.

Economists and policymakers have long recognized that changing prices will impact quantities, but new empirical evidence is still crucial to inform policy because the implications are very different depending on whether income or substitution effects dominate. CMS assumes that a reduction in physician fees will be offset with an increase in volume, and other work argues that fee reductions will be offset with more intense coding (Brunt, 2015). Both of these assumptions imply that income effects dominate. This paper shows that in the case of obstetrics, the substitution effect dominates—when the C-section price decreases, physicians substitute C-sections for vaginal delivery. Clemens and Gottlieb (2014) also find that the supply curve for physician and outpatient services is positively-sloped. The implications of these findings for policy will vary depending on the type of procedure. In the case of C-sections, lowering prices lowers treatment cost because vaginal delivery prices are lower, and arguably improves patient health, at least at the current margin.

These results are thus important for any policies that are likely to affect prices either directly or indirectly, including government pricing decisions, insurance rate regulation, and merger regulation, which can affect bargaining power and therefore prices. The latter is particularly salient as the trend toward provider consolidation and the push for increased insurer competition on the Health Insurance Exchanges continue (Dafny, 2014; Dafny, Gruber and Ody, 2015). The industrial organization literature has paid a lot of attention to the effects of insurer and provider consolidation on negotiated prices and premiums (e.g., Gowrisankaran, Nevo and Town, 2015; Dafny, Duggan and Ramanarayanan, 2012). However, that literature has said remarkably little about the effects of these negotiated prices on utilization. This work shows that privately negotiated prices can have a substantial impact on how patients are treated, which in turn will determine healthcare costs and outcomes.

Appendix A

Tables for the Moral Hazard Calculation

The mean subjective probability of speeding 15+ mph over the speed limit in the next year is 0.45 at Wave 1 and 0.40 at Wave 2 (Table A.1). The mean subjective probability of drinking and driving in the next year is 0.17 in both waves. The mean expected premium increase due to a speeding conviction is 28 percent at Wave 1 and 33 percent at Wave 2. The corresponding premium increases for a DWI conviction are 77 and 86 percent, respectively.

In Table A.2, the coefficients of interest are -0.332 (col. 2) and -0.089 (col. 5). These coefficients are reproduced in Table 1.6. Formulas for the expected premium increases if one speeds at least once during the year or if one drinks and drives at least once are shown in the footnote to Table A.2.

Table A.3 presents results of effects of changes in subjective probabilities on changes in actual driving behaviors and on the changes in actual accidents incurred during the year following Wave 1 resulting from changes in driving behavior.

Table A.1: Subjective probabilities and expected premium increases due to reckless driving

	Mean	Std. dev.
Speeding		
Subjective prob. of speeding 15mph+	0.45	0.40
Subjective prob. of speeding 15mph+ (2)	0.40	0.39
Subjective prob. of being pulled over | Speeding 15mph+	0.24	0.25
Subjective prob. of being pulled over | Speeding 15mph+ (2)	0.24	0.23
Subjective prob. of speeding conviction | Speeding 15mph+	0.51	0.37
Subjective prob. of speeding conviction | Speeding 15mph+ (2)	0.50	0.36
Expected premium increase | Speeding 15mph+ conviction	0.28	0.28
Expected premium increase | Speeding 15mph+ conviction (2)	0.33	0.34
Drinking and driving		
Subjective prob. of drinking and driving	0.17	0.30
Subjective prob. of drinking and driving (2)	0.17	0.30
Subjective prob. of being pulled over | Drinking and driving*	0.10	0.13
Subjective prob. of being pulled over | Drinking and driving* (2)	0.11	0.14
Subjective prob. of DWI conviction | Pulled over for drinking and driving*	0.60	0.33
Subjective prob. of DWI conviction | Pulled over for drinking and driving* (2)	0.61	0.32
Expected premium increase | DWI conviction	0.77	1.17
Expected premium increase | DWI conviction (2)	0.85	0.93
N	1136	

Notes: Excludes individuals without liability insurance.
(2) indicates a question elicited in Wave 2; otherwise, Wave 1.
* indicates a question related specifically to weekend drinking and driving.

Table A.2: Relationship between subjective probability of reckless driving and premium increase

		Subj. prob. of speeding 15mph+			Subj. prob. of drinking and driving		
	Mean (std. dev.)	Random pair	Random pair (bootstrap)	All responses	Random pair	Random pair (bootstrap)	All responses
Panel A. Speeding	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Subj. prob. of Speeding 15mph+	0.394						
	(0.395)						
Subj. prob. of Drinking and driving	0.119						
	(0.263)						
E[Premium increase | Speeding 15mph+]	0.053	−0.332***	−0.279***	−0.236***			
	(0.132)	(0.110)	(0.108)	(0.071)			
E[Premium increase | Drinking and driving]	0.077				−0.089***	−0.104***	−0.101***
	(0.197)				(0.024)	(0.028)	(0.021)
N	4,331	6,444	2,252	2,161	2,031	2,134	2,148
R^2		0.007	0.002	0.004	0.008	0.010	

Notes: Standard errors in parentheses. Excludes individuals without liability insurance. E[Premium increase | Speeding 15mph+] = E[Premium increase | Speeding Conviction]*P(Speeding Conviction | Pulled Over)*Prob(Pulled Over | Speeding 15mph+). E[Premium increase | Drinking and Driving] = E[Premium increase | DWI Conviction]*P(DWI Conviction | Pulled Over)*P(Pulled Over | Weekend Drinking and Driving). Regressions include individual and wave fixed effects.
* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table A.3: Moral hazard steps 2 and 3

	Drove after drinking in past year (wave 2)	Any accident in past year (wave 2)	Any accident (pooled)
	(1)	(2)	(3)
Subj. prob. of drinking and driving	0.693***		
	(0.036)		
Subj. prob. of speeding		0.046**	
		(0.021)	
Drove after drinking in past year (wave 2)			0.048***
			(0.015)
Wave 1			0.121***
			(0.014)
N	1,164	1,169	2,351

Notes: Regression in col. 3 includes individual fixed effects.
* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$

Appendix B

Survey of Alcohol and Driving

The model is estimated using individual-level survey data from the SAD, a panel survey conducted by Battelle Memorial Institute during 2010-12. The design of the SAD was guided by questions included in national surveys, such as the Behavioral Risk Factor Surveillance System.[1] Eligibility for the SAD required participants to have driven and consumed alcohol during the last month, to be 18 or older, and to reside within one of eight cities: Raleigh, NC; Hickory, NC; Philadelphia, PA; Wilkes-Barre, PA; Seattle, WA; Yakima, WA; Milwaukee, WI; or La Crosse, WI. These cities represent a broad geographic spread of large and small cities, and the four states in which they are located vary in alcohol consumption, DWI arrest rates, criminal laws pertaining to DWI, demographic composition, and insurance law and regulation. The sample was restricted to drinkers in order to get enough volume to study low-probability events, such as DWI arrests.

The SAD was administered in three surveys, depicted in Figure B.1. For the baseline variables, the SAD used a telephone screening followed by a computer-assisted

[1] Survey instruments for the SAD are at http://dialog.econ.duke.edu/dapstudy.

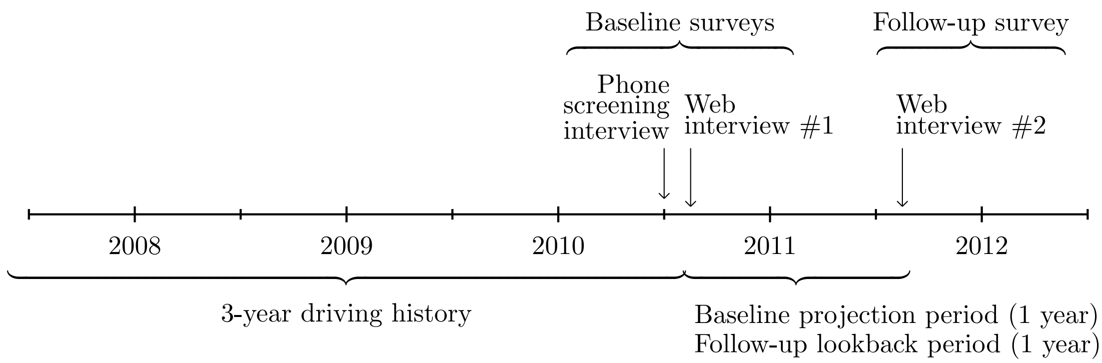

FIGURE B.1: Timing of administration of the Survey of Alcohol and Driving

survey approximately one month later to elicit respondents' three-year driving history, including accidents, speeding violations, and DWI arrests; household demographics, including wealth, number of adults and children, and number and types of cars; and the following automobile insurance policy characteristics: bodily injury liability limits, collision coverage (yes/no), collision insurance deductible (if applicable), premium, name of insurer, and number of insured drivers by age group (< 25 and $\geqslant 25$).[2] In addition, the SAD elicited whether the respondent had health insurance and the source (Medicare, Medicaid, military, private).

The SAD also elicited individual characteristics not available in administrative data. These include: risk preference based on hypothetical gambles over lifetime income (Barsky et al., 1997); risk preference based on hypothetical gambles over health[3]; a measure of impulsivity based on questions developed by Loewenstein et al.

[2] Respondents were paid a nominal amount for looking up their insurance policy in order to answer the insurance-related questions. In addition, the respondent is not necessarily the policyholder. For example, some respondents reported joint policies with a spouse, parents' policies that cover them, or, rarely, employer-provided auto insurance.

[3] The SAD asked: "We want you to keep imagining that you have gotten into an auto accident that leaves you paralyzed. Suppose that doctors could cure you of the paralysis by performing an operation. Without the operation, you would be paralyzed for the rest of your life. But if the operation went well, you would be completely cured of your paralysis. If the operation did not go well, you would die immediately without any pain. Would you choose to have the operation if the chance of dying was (X)?" The initial probability that would make the person indifferent between having and not having the operation was randomized. The computer program elicited responses until the change in subjective probabilities fell below a threshold.

(2001)[4]; non-pecuniary loss, measured as willingness to pay to avoid paralysis from an automobile accident (Sloan et al., 1998); and smoking and illicit drug use. The SAD elicited individual beliefs about: the importance of alcohol in one's social life, the increase in odds of an accident when driving after having four drinks versus no alcohol, and how much arrest for DWI would affect one's life. The SAD elicited subjective probabilities of: being involved in an accident, drinking and driving, and being arrested for DWI in the next year. Finally, the SAD elicited respondents' expectations about premium increases conditional on an accident or DWI conviction.

Twelve months after the baseline survey, the SAD used another computer-assisted survey to elicit whether the respondent ever drove after having 4-5 drinks, was involved in any accidents, or was arrested for DWI in the past 12 months. The total number of respondents in the baseline was 1,634. After accounting for attrition and missing values on key variables such as insurance coverage characteristics, I obtain a final sample size of 966.

[4] In psychology, impulsivity is an aspect of personality that describes the tendency to act on a whim, perhaps not considering future consequences of present actions.

Table B.1: Comparison of sample means

Variable	SAD Full Sample	National[a] Full Sample	National[a] Drinkers
Alcohol consumption			
Had any alcoholic beverage in past 30 days	1.00	0.55***	1.00
Any binge drinking (30 days) | drinking	0.41	—	0.33***
Avg alcoholic drinks per day | drinking	3.18	—	2.60***
Driving behavior			
Driven within the last 30 days	1.00	0.97***	0.99***
Drove after drinking too much (30 days) | drinking	0.03[b]	—	0.03[c]
Arrested for DUI (1 year)	0.007[d]	0.005[e]	—
Demographic characteristics			
Age	45.13	46.51**	44.61
Female	0.55	0.51**	0.46***
Married	0.49	0.50	0.51
Employed	0.77	0.55***	0.63***
Has health insurance	0.90	0.81***	0.82***
Adults in household	2.03	2.28***	2.29***
Children < 18 in household	0.67	0.75**	0.74**
Education			
Less than high school graduate	0.01	0.15	0.12
High school graduate or GED	0.10	0.29	0.27
Some college, technical school, or 2-year college	0.22	0.30	0.31
College graduate or graduate education	0.68	0.25	0.31
Chi-square critical value		(975.19)***	(676.63)***
Race			
White	0.87	0.66	0.69
Black	0.09	0.11	0.11
Other	0.04	0.23	0.21
Chi-square critical value		(212.00)***	(174.00)***
Income			
Less than $15k	0.07	0.14	0.11
$15-$35k	0.14	0.30	0.26
$35-$75k	0.34	0.29	0.30
>$75k	0.44	0.27	0.34
Chi-square critical value		(218.08)***	(99.47)***
State			
North Carolina	0.44	0.27	0.24
Pennsylvania	0.15	0.37	0.38
Washington	0.21	0.19	0.19
Wisconsin	0.19	0.16	0.19
Chi-square critical value		(174.90)***	(303.81)***

Notes: (a) Data from the 2011 Behavioral Risk Factor Surveillance System (BRFSS) except as noted; (b) estimate based on dividing the value from the past twelve months by twelve (0.40/12 = 0.03); (c) data from the 2012 BRFSS; (d) estimate based on dividing the value from the past three years by three (0.0217/3 = 0.007); (e) data from the Bureau of Justice Statistics Arrest Data Analysis Tool, 2011. Asterisks indicate statistical significance at the * 0.1 ** 0.05 and *** 0.01 levels from two-sided t-tests (binary and continuous variables) and Pearson's chi-squared tests (categorical variables) that the mean or distribution of the National data is equal to the SAD.

Table B.2: Distribution of insurance choices

	Collision insurance		
Liability limit	No	Yes	Total
uninsured	0.04	*b*	0.04
state min.	0.03	0.12	0.15
100,000	0.02	0.13	0.15
300,000	0.01	0.40	0.41
500,000	0.01	0.19	0.20
1,000,000	*a*	0.04	0.04
Obs. = 966	0.11	0.89	1.00

Notes: The table shows the distribution of observed insurance choices. "Liability limit" refers to the upper bodily injury liability limit. Responses are grouped to the nearest limit listed in the table, with upper cutoffs at 200k, 400k, and 750k, inclusive—e.g., a respondent who reports a 200k limit would be included in the 100k category in the table. Wisconsin residents who report a 100k limit are included in the "state min." category because this is the minimum in that state. Liability-collision combinations excluded from the choice set are: (a) A $1M liability limit without collision because it was only reported by two people in the sample, so they are attributed to the $500k liability limit, no collision policy; and (b) Collision insurance without liability insurance, because it cannot be purchased alone in any state.

Appendix C

Accident cost tables

Table C.1 reports the mean accident costs by severity level and cost category. Panels A and B are from Blincoe et al. (2014). Panel C is derived from Panels A and B to obtain a discrete distribution of per-accident costs in the four categories relevant for automobile insurance: injury to others, injury to self, property damage to others, and own property damage. Table C.2 reports mean expected accident costs by each combination of insurance plan and drinking and driving behavior. These expected costs are calculated using the accident cost distributions in Table C.1, Panel C, and the formulas in Table 2.2.

Table C.1: Distribution of accident costs by severity level, 2010

	Severity							
	PDO	MAIS 0	MAIS 1	MAIS 2	MAIS 3	MAIS 4	MAIS 5	Fatal
Panel A: Total accident costs (millions)								
Property damage	45,235	8,378	18,694	1,957	1,096	279	87	370
Medical	0	0	13,148	8,292	7,143	3,364	2,539	373
EMS	518	96	308	66	42	14	5	30
Market productivity	0	0	9,886	13,426	11,090	2,985	2,025	30,797
Household productivity	1,111	206	3,255	4,038	3,375	762	616	9,567
Workplace cost	1,148	211	1,180	896	582	109	64	389
Legal cost	0	0	4,877	2,283	1,979	603	524	3,514
QALYs	0	0	80,395	115,464	81,166	34,812	26,322	255,646
Panel B: Total vehicles/people in accidents								
Vehicles	18,508,632	—	—	—	—	—	—	—
People	—	4,583,265	3,459,200	338,730	100,740	17,086	5,749	32,999
Panel C: Cost per accident ($)								
Injury to others	113	207	60,318	787,163	1,930,633	4,607,064	10,303,881	16,797,068
Injury to self	150	112	32,681	426,490	1,046,029	2,496,137	5,582,710	9,100,761
Property damage to others	1,837	3,374	9,974	10,663	20,080	30,138	27,931	20,695
Own property damage	2,444	1,828	5,404	5,777	10,879	16,329	15,133	11,212

Notes: The table reports total accident costs, number of vehicles or people involved, and mean costs per accident at each of seven severity levels in the U.S. in 2010 based on Blincoe et al. (2014), Table 1-8. PDO stands for "property damage only" and MAIS stands for "Maximum Abbreviated Injury Scale", which is a measure of the severity of injuries, where zero is least severe and five is most severe. EMS refers to medical, police, and fire emergency service costs. QALYs refers to quality-adjusted life years and is an estimate of the value of life lost due to injuries and deaths. The values derived in Panel C are based on the average number of vehicles and people involved in a crash of 1.75 and 2.85 from Blincoe et al. (2014), Table 1-3. Injury to others and self includes all accident costs in Panel A except property damage.

Table C.2: Expected pecuniary cost of driving conditional on insurance policy and drinking behavior

Policy-behavior combination				Mean expected cost							
Liability	Collision	Ded.	Drink	Bodily injury		Property damage		Premium increase		Fine	Total
				Others	Self	Others	Own	Accident	DWI	DWI	
uninsured	—	—	no	1,111	35	261	169	0	—	—	1,576
state min.	no	—	no	497	35	10	169	53	—	—	764
100,000	no	—	no	435	35	0	169	54	—	—	694
300,000	no	—	no	401	35	0	169	58	—	—	664
500,000	no	—	no	355	35	0	169	64	—	—	623
state min.	yes	500	no	497	35	10	29	86	—	—	657
100,000	yes	500	no	435	35	0	29	87	—	—	587
300,000	yes	500	no	401	35	0	29	92	—	—	558
500,000	yes	500	no	355	35	0	29	97	—	—	516
1,000,000	yes	500	no	151	35	0	29	109	—	—	325
uninsured	—	—	yes	1,797	52	419	271	0	0	99	2,638
state min.	no	—	yes	811	52	17	271	86	111	99	1,447
100,000	no	—	yes	704	52	1	271	88	113	99	1,328
300,000	no	—	yes	650	52	0	271	96	123	99	1,290
500,000	no	—	yes	574	52	0	271	103	133	99	1,233
state min.	yes	500	yes	811	52	17	47	140	181	99	1,347
100,000	yes	500	yes	704	52	1	47	142	183	99	1,228
300,000	yes	500	yes	650	52	0	47	149	194	99	1,190
500,000	yes	500	yes	574	52	0	47	157	203	99	1,132
1,000,000	yes	500	yes	245	52	0	47	176	229	99	848

Notes: The table reports the means of expected driving costs conditional on each combination of insurance policy and drinking and driving behavior. The first four columns of the table describe the insurance policy-drinking behavior combination. Liability is the upper limit (per accident) on bodily injury liability to others; Collision is an indicator for whether the plan covers damage to the policyholder's property; Ded. is the collision insurance deductible, assumed to be \$500; and Drink is an indicator for whether or not the individual hypothetically chooses to drink and drive in the second period. The expected cost of the seven pecuniary consequences are calculated according to Table 2.2. All consequences are reported in dollars. The Total column is the sum of the individual consequences and corresponds to the mean of equation (2.3) over all individuals for policy j and drinking behavior d.

Appendix D

Optimists and pessimists

One of the key assumptions that Chiappori et al. (2006) and Fang and Wu (2016) make is that individuals have "realistic expectations"—i.e., their subjective assessment of risk is equal to the objective risk. Under this assumption, both papers show that there cannot be a negative correlation between insurance choice and accident risk in a competitive insurance market unless insurance loading factors are quite high—which is not likely in a perfectly competitive market, as Fang and Wu (2016) point out. However, Spinnewijn (2013) shows that violations of the realistic expectations assumption can break the positive correlation between insurance coverage and accident risk. In particular, a negative correlation can arise when a group of individuals are both overly optimistic about their true risk ("baseline optimistic") and pessimistic about how much their behaviors affect that risk ("control pessimistic"). Here I use subjective beliefs observed in the SAD together with other features of the survey to show that risk tolerant individuals are baseline optimistic and control pessimistic.

D.1 Baseline Optimism

I test for baseline optimism in the SAD in two ways. First, I take advantage of the fact that the SAD contains subjective probabilities about future events in the next year *and* realizations of the same events collected in the follow-up survey a year later. To test whether there is heterogeneity in the accuracy of individuals' risk perceptions by risk preference, I run the following regression for three outcomes (y)—accidents, drinking and driving, and speeding:

$$y_{it+1} = \beta_0 + \beta_1 SubjProb_{it}(y) + \beta_2 RiskPrefFactor_i + \varepsilon_i$$

The coefficient of interest is β_2. If $\beta_2 > 0$, risk tolerant individuals underestimate their objective risk relative to risk averse individuals because the deviation of their actual risk from their subjective probability is higher. If $\beta_2 < 0$, risk tolerant individuals overestimate their objective risk relative to risk averse individuals.

For all three outcomes, risk tolerant individuals have higher accident risk conditional on their subjective beliefs, and are therefore baseline optimistic (Table D.1), although the coefficient on risk preference is not statistically significant for speeding. In fact, because the risk preference factor is mean-zero and therefore takes negative values for risk averse individuals, we can conclude that risk tolerant individuals are baseline optimistic and risk averse individuals are baseline pessimistic.

I perform a second, simple test of baseline optimism in the SAD sample by finding the correlation between risk preference and the "optimism index" developed by Sloan et al. (2013). The optimism index counts the number of times each SAD respondent underestimate the probability of an adverse event for ten events. The correlation between this index, which is increasing in optimism, and the risk preference factor is 0.22 (p-value=0.000). In other words, risk tolerant individuals are much more optimistic in general.

Table D.1: Accuracy of subjective beliefs by risk preference

	Wave 2 Realization		
	Any Accident	Drunk Driving	Speeding > 15 mph
Subjective probability (Wave 1)	0.069	0.478***	0.064***
	(0.065)	(0.036)	(0.023)
Risk preference factor	0.023**	0.075***	0.007
	(0.009)	(0.011)	(0.009)
Constant	0.080	0.087	0.057
	(0.012)	(0.013)	(0.013)
Observations	965	966	964

Notes:

Finally, there is theoretical motivation for these results from the idea in prospect theory that people overweight small probabilities Kahneman and Tversky (1979). The "s-curve" implies that people tend to overweight small probabilities. Even though accident probabilities are relatively small on average, we can expect low-risk drivers to be relatively pessimistic compared to high-risk drivers as long at the low-risk drivers' accident probability is not too close to zero. Because the risk averse individuals are low risk, we should expect them to be relatively pessimistic and the risk tolerant drivers—who are higher risk—to be relatively optimistic.

D.2 Control Pessimism

In order to break the positive correlation, baseline optimistic drivers must also be control pessimistic—in other words, they must believe their their actions do not have a significant effect on their accident risk. The SAD makes it is straightforward to test this with respect to drinking and driving behavior. One of the questions asked is "If you drank 4 drinks and then drove home, what would be the odds (compared with not having had any alcohol at all) of getting into accident?" There is a set of five discrete response options ranging from "No increase in odds" to "Odds increased more than 100%". The correlation between the increase in accident odds and the risk preference factor is -0.093 (p-value=0.004), implying that risk tolerant drivers believe drinking and driving is less likely to affect the odds of an accident than risk averse drivers do, and are therefore relatively control-pessimistic.

It is possible that risk tolerant drivers are also more frequent drinkers and therefore have (or believe they have) a higher tolerance for alcohol so don't think four drinks is enough to substantially affect their driving. I run a regression to control for the self-reported number of drinks that it would take for an individual to exceed the BAC limit, and still find a negative and statistically significant coefficient on the risk preference factor. In a similar regression controlling for how well a person believes they can handle alcohol compared to others, the coefficient on risk preference is still negative, though not statistically significant (p-value=0.213).

Appendix E

Counterfactual premiums under risk-rating

I first recover the residual variation in the expected accident probabilities for drinker-drivers and non-drinker-drivers from the perspective of the insurer using a logit model:

$$P\left(accident_{it+1}\right) = \Lambda\left(\beta_1\, observed\ risk_{it} + \beta_2\, drink\ drive_{it}\right) \tag{E.1}$$

where $accident_{it+1}$ is an indicator for whether the individual was involved in any accident during the follow-up year; $observed\ risk_{it}$ is the vector of individual risk factors in the hedonic regression, which are observed and used by the insurer under the status quo; $drink\ drive_{it}$ is an indicator for whether the individual reports ever having driven after having too much to drink in the past twelve months; and $\Lambda(x) = \frac{\exp(x)}{1+\exp(x)}$ is the logistic function. I also include a full set of insurer dummy variables to account for variation in the risk of the insured population at the insurer level. The measure of interest from equation E.1 is the relative risk of an accident for a drinker-driver compared to a non-drinker-driver. Relative risk is approximated by $\exp(\beta_2)$, the odds ratio corresponding to the coefficient on $drink\ drive_{it}$.[1]

[1] This approximation is reasonable for low-probability events, such as accidents (Deeks, 1998).

Having recovered an estimate of the relative risk of an accident for drinker-drivers compared to non-drinker-drivers conditional on other observed characteristics, I calculate the change in premium for a drinker-driver:

$$\begin{aligned}
\Delta P^{DD} &= E[cost|DD] - E[cost] \\
&= P(accident|DD) \times (cost|DD) - P(accident) \times cost \\
&= \Big[P(accident|DD) - P(accident)\Big] \times cost \\
&= \Big[RR^{DD} \times P(accident|NDD) - P(accident)\Big] \times cost \\
&= \left[RR^{DD} \times \frac{P(accident)}{P(NDD) + RR^{DD} \times P(DD)} - P(accident)\right] \times cost \\
&= \left[\frac{RR^{DD}}{P(NDD) + RR^{DD} \times P(DD)} - 1\right] \times E[cost]
\end{aligned}$$

where $cost$ is the cost of an accident, DD is a drinker-driver, NDD is a non-drinker-driver, and RR^{DD} is the relative risk of a drinker-driver compared to a non-drinker-driver. The third line follows by assuming that accident cost does not vary with drinking and driving behavior; the fourth line follows from the definition of relative risk; and the fifth line follows from the law of total probability.[2] Because $P(NDD)+P(DD) = 1$ and assuming RR^{DD} is strictly greater than one, $P(NDD)+ RR^{DD} \times P(DD) < RR^{DD}$, so the premium change for a drinker-driver is positive. The premium change for a non-drinker-driver is:

$$\Delta P^{NDD} = \left[\frac{1}{P(NDD) + RR^{DD} \times P(DD)} - 1\right] \times E[cost]$$

[2] By the law of total probability (line 1) and the definition of relative risk (line 2):

$$\begin{aligned}
P(accident) &= P(accident|NDD) \times P(NDD) + P(accident|DD) \times P(DD) \\
P(accident) &= P(accident|NDD) \times P(NDD) + RR^{DD} \times P(accident|NDD) \times P(DD) \\
P(accident|NDD) &= \frac{P(accident)}{P(NDD) + RR^{DD} \times P(DD)}
\end{aligned}$$

which is negative because $P(NDD) + RR^{DD} \times P(DD) > 1$.

I use the SAD to estimate the probabilities of having an accident, being a drinker-driver, and being a non-drinker driver. I use information from Table C.1 to estimate accident costs. Finally, I use these estimates, along with the estimated relative risk from the logit model, to calculate average premium adjustments for drinker-drivers and non-drinker-drivers.

Appendix F

Identifying individual physicians based on NPI

I use the Provider files from the MA APCD (Release 2.0) to identify individual physicians and their National Provider Identifiers (NPIs). My objective is to obtain the unique NPI of the physician who was responsible for each claim line in the Medical Claims file. The problem is that the Medical Claims file does not contain NPIs for the physician who performed each claim line. The Medical Claims file does contain a Linking Provider ID that is intended to link the Medical Claims file to the Provider file, which contains rich demographic and practice information about providers, including NPI. Unfortunately, the Linking Provider ID is not a unique field in the Provider file, making a simple left join infeasible. In fact, I was unable to identify any set of variables in the Provider file and also in the Medical Claims file that would yield a unique key. I take the following steps to overcome this, which are also illustrated in Figure F.1:

- Stack Commercial and Medicaid provider files.
- Collapse the combined file by removing duplicates based only on Release ID, provider affiliations (linkage IDs and dates), and begin and end dates, yielding

a 57% reduction in observations, from 42 to 18 million.

- Count number of unique Provider ID Codes for each Linking Provider ID. Provider ID Code indicates the type of entity (person, facility, group, etc.).
- Separate Linking Provider IDs into the following groups:
  - Persons — Provider ID Code = 1 (Person) for all observations with a single Linking Provider ID
  - Non-Persons — Provider ID Code $\neq$ 1 for all observations with a single Linking Provider ID
  - Unknown/Multiple — Multiple Provider ID Codes for a single Linking Provider ID
- Among Linking Provider IDs associated as Person, collapse on only needed variables, then separate into the following groups:
  - One Line — exactly one observation per Linking Provider ID
  - Two Lines — exactly two observations per Linking Provider ID
  - Multiple Lines — more than two observations per Linking Provider ID
- Create final Physician Map from Linking Provider ID to NPI using One Line Persons as is; collapsing Two Line Persons down based on logic, where possible; and discarding Multiple Line Persons and Non-Persons.

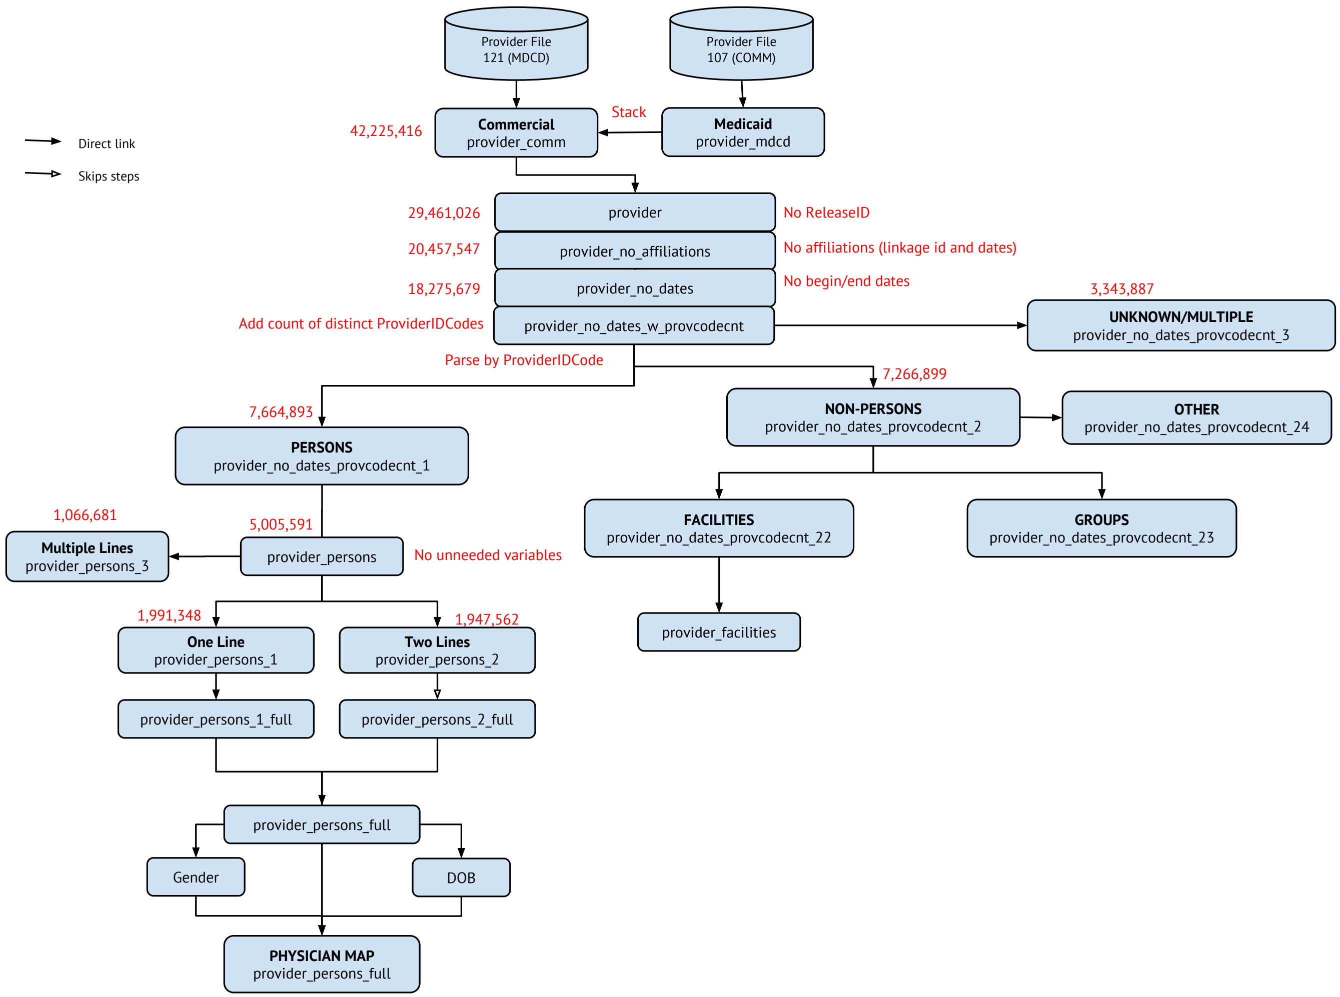

Figure F.1: Schematic of code for identifying individual physician NPIs in MA APCD

Appendix G

Identifying contract and price changes

To identify contract and price changes, I first extract the payer ID, provider ID (NPI), incurred date, procedure, and allowed amount for a subset of the sample of birth claims. I restrict the sample in two dimensions.

First, I only include deliveries involving routine obstetric care, including antepartum and postpartum care. These are the most frequently observed codes, and restricting the sample to these codes allows me to focus on only two prices per provider-patient pair: one for vaginal delivery and one for C-section delivery. This restriction reduces the sample by 17 percent, with 2 percent associated with vaginal births after a previous C-section (VBAC), 7 percent associated with delivery-only lines, and the remaining 8 percent associated with deliveries that include postpartum but not antepartum care. Attempting VBAC is relatively rare and involves additional complications, while claim lines that indicate delivery-only may mean that the patient switched physicians, or that the pregnancy was terminated (by miscarriage or abortion). Claims without antepartum care may indicate that the physician changed or that the provider bills for antepartum care separately.

Second, I restrict the sample based on volume of deliveries over time. To do this, I count the number of consecutive months that each physician (NPI) delivers and the average number of deliveries in those months. I sort and visually inspect volume per month at the physician level to determine a threshold level. There is a tradeoff between including more data and losing the density of deliveries over time that is necessary to identify exactly when a contract likely changed. I restrict the sample to include only NPIs with at least twelve consecutive months with a delivery and at least three deliveries per month.

Before describing the algorithm to identify the contract periods for this subset of physicians, I illustrate the concept with example.

G.1 Example

The logic of the pricing algorithm is best understood by visualizing an example, such as the one shown in Figure G.1. The figure shows the count of deliveries for a given physician-payer pair by price and incurred date, with all the vaginal deliveries given in the top panel (CPT code 59400) and the C-sections in the bottom panel (CPT code 59510). The yellow and orange lines indicate likely contract change dates, which are based on when an entirely new "menu" of prices shows up. For example, before 9/1/2010, there are four prices observed, two for each procedure: $2,894.29, $4,035.96, $3277.40, and $4570.18. Beginning on 9/1/2010, none of these prices are ever observed again, and there are four new consistently observed prices ($3,168.51, $4,168.41, $3,588.05, and $4720.34) plus one rogue price ($4077.83). This set of prices is active until 9/5/2011, when yet another completely new menu of prices appears. This aligns with what we would expect from physician contract changes: the same prices are in place for an entire contract period, and in general, the contract changes annually. The goal of the algorithm, then, is to identify dates at which the entire menu of prices changes.

Count of Billing

Row Labels	5/2/2010	5/4/2010	5/5/2010	6/1/2010	6/2/2010	6/3/2010	6/4/2010	6/5/2010	7/3/2010	7/4/2010	7/5/2010	8/1/2010	8/2/2010	8/3/2010	8/5/2010	8/6/2010	9/1/2010	9/5/2010	10/3/2010	10/4/2010	10/5/2010	10/7/2010
59400																						
2894.29		1		1							1	1	2	1		1						
3168.51																			1		
3359.67																						
3476.2																						
3614.27																						
4035.96					1		1	1	1													
4168.41																				1	1	
4276.07																						
4352.34																						
4579.98																						
59510																						
3277.4			1		1	1									1							
3588.05																						1
3614.27																						
3738.81																						
3868.5																						
4077.83																						
4570.18	1					1				1				1								
4579.98																						
4720.34																	1	1	1			
4758.63																						
4843.51																						

Row Labels	11/1/2010	11/2/2010	11/5/2010	11/6/2010	12/1/2010	12/4/2010	12/7/2010	1/5/2011	2/1/2011	2/4/2011	2/5/2011	3/4/2011	3/6/2011	3/7/2011	4/1/2011	4/2/2011	4/3/2011	4/6/2011	5/4/2011	5/6/2011	6/4/2011	6/6/2011
59400																						
2894.29																						
3168.51				1	1		1			1		1	1	1		1		1	1			1
3359.67																						
3476.2																						
3614.27																						
4035.96																						
4168.41		1	1					1	1							1	1			1		
4276.07																						
4352.34																						
4579.98																						
59510																						
3277.4																						
3588.05				1		1	1		1													
3614.27																						
3738.81																						
3868.5																						
4077.83												1										
4570.18																						
4579.98																						
4720.34	1								1		1				1					1	1	
4758.63																						
4843.51																						

Row Labels	7/3/2011	8/2/2011	8/3/2011	8/4/2011	8/5/2011	9/5/2011	9/6/2011	9/7/2011	10/2/2011	10/3/2011	10/5/2011	11/1/2011	11/2/2011	11/3/2011	11/4/2011	11/6/2011	11/7/2011	12/3/2011	1/5/2012	3/1/2012	4/2/2012	4/4/2012	4/5/2012
59400																							
2894.29																							
3168.51			1																				
3359.67																				1	1		
3476.2						1		1	1		1			1	1	1							
3614.27																							
4035.96																							
4168.41		1																					
4276.07																						1	1
4352.34						1				1	1												
4579.98																							
59510																							
3277.4																							
3588.05																							
3614.27																							
3738.81																							
3868.5							1	1					1					1					
4077.83																							
4570.18																							
4579.98																							
4720.34	2			1	1																		
4758.63																			1				
4843.51												1					1						

FIGURE G.1: Example identification of contract change dates for a payer-provider pair

G.2 Algorithm

- Create a table of minimum and maximum incurred dates for each physician-payer-allowed amount triplet.
- Sort the table by physician, payer, minimum date, then maximum date.
- For each physician-payer pair:
  - For the first observation (earliest observed allowed amount):
    - Set the contract start date equal to the minimum date.
    - Set the contract end date equal to the maximum date.
  - For all subsequent observations:
    - Retain the previous contract start date unless the minimum date is greater than the previous contract end date.
    - Retain the previous contract end date unless the maximum date is greater than the previous contract end date.
- Collapse this table by taking the maximum contract end date for each physician-payer-contract start date triplet.
- Convert contract dates (year-month-day) to year-months.
- Merge in mean prices during each contract period.
- Drop contract periods <60 days and with zero vaginal deliveries or C-sections.

After applying this algorithm, I summarize the results by payer and find, consistent with what we would expect, that for any given payer, most contracts change in the same month within a 12-month period and in the same month each year. In addition, different payers change their contracts in different months, providing the identifying variation needed for the difference-in-differences.

Bibliography

Abbring, Jaap H., James J. Heckman, Pierre-André Chiappori, and Jean Pinquet. 2003. "Adverse Selection and Moral Hazard in Insurance: Can Dynamic Data Help to Distinguish?" *Journal of the European Economic Association*, 1(2/3): 512–21.

Abbring, Jaap H., Pierre-André Chiappori, and Jean Pinquet. 2003. "Moral Hazard and Dynamic Insurance Data." *Journal of the European Economic Association*, 1(4): 767–820.

Agency for Healthcare Research and Quality. 2011. "Medical Expenditure Panel Survey Summary Data Tables I.F.2-3(2011)." http://meps.ahrq.gov/mepsweb/data_stats/quick_tables.jsp (accessed November 6, 2014).

American Medical Association. 2016. *RVS Update Process Booklet 2017.* Chicago, IL: American Medical Association. https://www.ama-assn.org/sites/default/files/media-browser/public/rbrvs/ruc-update-booklet_0.pdf (accessed February 22, 2017).

Austin, Daniel R., and Laurence C. Baker. 2015. "Less Physician Practice Competition is Associated With Higher Prices Paid for Common Procedures." *Health Affairs*, 34(10): 1753–60.

Bajari, Patrick, Han Hong, and Ahmed Khwaja. 2014. "Moral Hazard, Adverse Selection, and Health Expenditures: A Semiparametric Analysis." *RAND Journal of Economics*, 45(4): 747–63.

Baker, Tom, and Rick Swedloff. 2013. "Regulation by Liability Insurance: From Auto to Lawyers Professional Liability." *UCLA Law Review*, 60(6): 1412–50.

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro. 1997. "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study." *Quarterly Journal of Economics*, 112(2): 537–79.

Beitel, George A., Michael C. Sharp, and William D. Glauz. 2000. "Probability of Arrest While Driving Under the Influence of Alcohol." *Injury Prevention*, 6(2): 158–61.

Bisgaier, Joanna, and Karin Rhodes. 2011. "Access to Specialty Care for Children With Public Insurance." *New England Journal of Medicine*, 365(11): 2324–33.

Blincoe, Lawrence, Ted R. Miller, Eduard Zaloshnja, and Bruce A. Lawrence. 2014. "The Economic and Societal Impact Of Motor Vehicle Crashes, 2010." National Highway Traffic Safety Administration Report No. DOT HS 812 013, Washington, D.C. https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812013 (accessed May 6, 2015).

Bovbjerg, Randall R., Frank A. Sloan, Avi Dor Dor, and Chee Ruey Hsieh. 1991. "Juries and Justice: Are Malpractice and Other Personal Injuries Created Equal?" *Law and Contemporary Problems*, 54(1): 5–42.

Brobeck, Stephen, Michael Best, and Tom Feltner. 2013. "Uninsured Drivers: A Societal Dilemma in Need of a Solution." Consumer Federation of America, Washington, D.C. http://www.consumerfed.org/pdfs/140310_uninsureddriversasocialdilemma_cfa.pdf (accessed November 11, 2014).

Brunt, Christopher S. 2015. "Medicare Part B Intensity and Volume Offset." *Health Economics*, 24(8): 1009–26.

Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney. 2012. "Pricing and Welfare in Health Plan Choice." *American Economic Review*, 102(7): 3214–48.

Caughey, Aaron B., Alison G. Cahill, Jeanne Marie Guise, and Dwight J. Rouse. 2014. "Safe Prevention of the Primary Cesarean Delivery." *American Journal of Obstetrics & Gynecology*, 210(3): 179–93.

Centers for Disease Control and Prevention. 2016. "Births - Method of Delivery." https://www.cdc.gov/nchs/fastats/delivery.htm (accessed February 22, 2017).

Centers for Medicare & Medicaid Services. 2013. "Key Milestones in CMS Programs." http://www.cms.gov/About-CMS/Agency-Information/History/ (accessed December 8, 2013).

Chandra, Amitabh, David M. Cutler, and Zirui Song. 2011. "Who Ordered That? The Economics of Treatment Choices in Medical Care." In *Handbook of Health Economics*. Vol. 2, edited by Mark V. Pauly, Thomas G. McGuire and Pedro P. Barros, 397–432. Oxford: Elsevier.

Chiappori, Pierre-André, and Bernard Salanié. 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy*, 108(1): 56–78.

Chiappori, Pierre-André, Bruno Jullien, Bernard Salanié, and François Salanié. 2006. "Asymmetric Information in Insurance: General Testable Implications." *RAND Journal of Economics*, 37(4): 783–98.

Clemens, Jeffrey, and Joshua D Gottlieb. 2014. "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?" *American Economic Review*, 104(4): 1320–49.

Clemens, Jeffrey, Joshua D. Gottlieb, and Timea Laura Molnar. 2015. "The Anatomy of Physician Payments: Contracting Subject to Complexity." National Bureau of Economic Research Working Paper 21642.

Cohen, Alma. 2005. "Asymmetric Information and Learning: Evidence From the Automobile Insurance Market." *Review of Economics and Statistics*, 87(2): 197–207.

Cohen, Alma, and Peter Siegelman. 2010. "Testing for Adverse Selection in Insurance Markets." *Journal of Risk and Insurance*, 77(1): 39–84.

Consumer Federation of America. 2014. "Penalties for Driving without Auto Insurance by State as of January 2014." Consumer Federation of America, Washington, D.C. http://www.consumerfed.org/pdfs/140310_penaltiesfordrivingwithoutautoinsurance_cfa.pdf (accessed October 1, 2015).

Cooper, Zack, Stuart Craig, Martin Gaynor, and John Van Reenen. 2015. "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured." National Bureau of Economic Research Working Paper 21815.

Coughlin, Teresa A., John Holahan, Kyle Caswell, and Megan McGrath. 2014. "Uncompensated Care for Uninsured in 2013: A Detailed Examination." The Henry J. Kaiser Family Foundation, Washington, D.C. https://kaiserfamilyfoundation.files.wordpress.com/2014/05/8596-uncompensated-care-for-the-uninsured-in-2013.pdf (accessed February 22, 2017).

Currie, Janet, and W. Bentley MacLeod. 2008. "First Do No Harm? Tort Reform and Birth Outcomes." *Quarterly Journal of Economics*, 123(2): 795–830.

Cutler, David M, and Sarah J. Reber. 1998. "Paying For Health Insurance: The Trade-off Between Competition and Adverse Selection." *Quarterly Journal of Economics*, 113(2): 433–66.

Dafny, Leemore. 2014. "Hospital Industry Consolidation—Still More to Come?" *New England Journal of Medicine*, 370(3): 198–9.

Dafny, Leemore, Jonathan Gruber, and Christopher Ody. 2015. "More Insurers Lower Premiums: Evidence From Initial Pricing in the Health Insurance Marketplaces." *American Journal of Health Economics*, 1(1): 53–81.

Dafny, Leemore, Mark Duggan, and Subramaniam Ramanarayanan. 2012. "Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry." *American Economic Review*, 102(2): 1161–85.

de Meza, David, and David C. Webb. 2001. “Advantageous Selection in Insurance Markets.” *RAND Journal of Economics*, 32(2): 249–62.

Decker, Sandra L. 2012. “In 2011 Nearly One-Third of Physicians Said They Would Not Accept New Medicaid Patients, but Rising Fees May Help.” *Health Affairs*, 31(8): 1673–9.

Deeks, Jon. 1998. “When Can Odds Ratios Mislead? Odds Ratios Should be Used Only in Case-Control Studies and Logistic Regression Analyses.” *BMJ*, 317(7166): 1155–6.

Dionne, Georges, Christian Gouriéroux, and Charles Vanasse. 2001. “Testing for Evidence of Adverse Selection in the Automobile Insurance Market: A Comment.” *Journal of Political Economy*, 109(2): 444–53.

Dulleck, Uwe, and Rudolf Kerschbamer. 2006. “On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods.” *Journal of Economic Literature*, 44(1): 5–42.

Ehrlich, Isaac, and Gary S. Becker. 1972. “Market Insurance, Self-Insurance, and Self-Protection.” *Journal of Political Economy*, 80(4): 623–48.

Einav, Liran, Amy Finkelstein, Iuliana Pascu, and Mark R. Cullen. 2012. “How General Are Risk Preferences? Choices Under Uncertainty in Different Domains.” *American Economic Review*, 102(6): 2606–38.

Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen. 2013. “Selection on Moral Hazard in Health Insurance.” *American Economic Review*, 103(1): 178–219.

Einav, Liran, and Amy Finkelstein. 2011. “Selection in Insurance Markets: Theory and Empirics in Pictures.” *Journal of Economic Perspectives*, 25(1): 115–38.

Fang, Hanming, and Zenan Wu. 2016. “Multidimensional Private Information, Market Structure and Insurance Markets.” National Bureau of Economic Research Working Paper 22773.

Fang, Hanming, Michael P. Keane, and Dan Silverman. 2008. “Sources of Advantageous Selection: Evidence From the Medigap Insurance Market.” *Journal of Political Economy*, 116(2): 303–50.

Finkelstein, Amy, and James Poterba. 2004. “Adverse Selection in Insurance Markets: Policyholder Evidence From the U.K. Annuity Market.” *Journal of Political Economy*, 112(1): 183–208.

Finkelstein, Amy, and Kathleen McGarry. 2006. “Multiple Dimensions of Private Information: Evidence From the Long-Term Care Insurance Market.” *American Economic Review*, 96(4): 938–58.

Fowler, Floyd J., Bethany S. Gerstein, and Michael J. Barry. 2013. “How Patient Centered Are Medical Decisions?” *JAMA Internal Medicine*, 173(13): 1215–21.

Fox, Will, and John Pickering. 2008. “Hospital and Physician Cost Shift: Payment Level Comparison of Medciare, Medicaid, and Commercial Payers.” Milliman Insight online publication. us.milliman.com/insight/research/health/pdfs/Hospital-and-physician-cost-shift (accessed February 22, 2017).

Gibbons, Luz, José M. Belizán, Jeremy A. Lauer, Ana P. Betrán, Mario Merialdi, and Fernando Althabe. 2010. “The Global Numbers and Costs of Additionally Needed and Unnecessary Caesarean Sections Performed per Year: Overuse as a Barrier to Universal Coverage.” World Health Report (2010) Background Paper, No 30.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town. 2015. “Mergers When Prices Are Negotiated: Evidence From the Hospital Industry.” *American Economic Review*, 105(1): 172–203.

Grant, Darren. 2009. “Physician Financial Incentives and Cesarean Delivery: New Conclusions From the Healthcare Cost and Utilization Project.” *Journal of Health Economics*, 28(1): 244–50.

Gruber, Jonathan, and Botond Köszegi. 2001. “Is Addiction ‘Rational’? Theory and Evidence.” *Quarterly Journal of Economics*, 116(4): 1261–303.

Gruber, Jonathan, and Maria Owings. 1996. “Physician Financial Incentives and Cesarean Section Delivery.” *RAND Journal of Economics*, 27(1): 99–123.

Gruber, Jonathan, John Kim, and Dina Mayzlin. 1999. “Physician Fees and Procedure Intensity: The Case of Cesarean Delivery.” *Journal of Health Economics*, 18(4): 473–90.

Guck, Thomas P., Gary N. Elsasser, Michael G. Kavan, and Eugene J. Barone. 2003. “Depression and Congestive Heart Failure.” *Congestive Heart Failure*, 9(3): 163–9.

Hackmann, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski. 2015. “Adverse Selection and an Individual Mandate: When Theory Meets Practice.” *American Economic Review*, 105(3): 1–56.

Hansen, Benjamin. 2015. “Punishment and Deterrence: Evidence From Drunk Driving.” *American Economic Review*, 105(4): 1581–617.

Hepburn, Silvia R., Thorsten Barnhofer, and J. Mark G. Williams. 2009. "The Future is Bright? Effects of Mood on Perception of the Future." *Journal of Happiness Studies*, 10(4): 483–96.

Hsiao, William C., Peter Braun, Yntema Douwe, and Edmund R. Becker. 1988. "Estimating Physicians' Work for a Resource-Based Relative-Value Scale." *New England Journal of Medicine*, 319(13): 835–41.

Iizuka, Toshiaki. 2012. "Physician Agency and Adoption of Generic Pharmaceuticals." *American Economic Review*, 102(6): 2826–58.

Institute of Medicine. 2013. *Best Care at Lower Cost: The Path to Continuously Learning Health Care in America.* Washington, D.C.: The National Academies Press. http://www.nap.edu/catalog.php?record_id=13444 (accessed December 3, 2013).

Insurance Information Institute. 2011. *The Insurance Fact Book 2011.* New York, NY: Insurance Information Institute.

Insurance Information Institute. 2012. "Competition in the U.S. Auto Insurance Market Has Driven Down Premium Costs." http://www.iii.org/press-release/competition-in-the-us-auto-insurance-market-has-driven-down-premium-costs-013012 (accessed April 5, 2017).

Israel, Mark. 2004. "Do We Drive More Safely When Accidents are More Expensive? Identifying Moral Hazard From Experience Rating Schemes." Unpublished.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica*, 47(2): 263–92.

Kaiser Family Foundation. 2014. "Insurance Market Competitiveness." http://kff.org/state-category/health-insurance-managed-care/insurance-market-competitiveness (accessed September 23, 2016).

Khwaja, Ahmed, Frank Sloan, and Yang Wang. 2009. "Do Smokers Value Their Health and Longevity Less?" *Journal of Law and Economics*, 52(1): 171–96.

Krause, Trudy Millard, Maria Ukhanova, and Frances Lee Revere. 2016. "Private Carriers' Physician Payment Rates Compared With Medicare and Medicaid." *Journal of Texas Medicine*, 112(6): 63.

Lemaire, Jean. 2012. *Bonus-Malus Systems in Automobile Insurance.* Dordrecht, Netherlands: Springer.

Levitt, Steven D., and Jack Porter. 2001. "How Dangerous are Drinking Drivers?" *Journal of Political Economy*, 109(6): 1198–237.

Lieberman, Joseph A. 2003. "The Differential Diagnosis of Fatigue and Executive Dysfunction in Primary Care." *Journal of Clinical Psychiatry*, 64(Suppl. 14): 40–43.

Lipsey, R. G., and Kelvin Lancaster. 1956. "The General Theory of Second Best." *Review of Economic Studies*, 24(1): 11–32.

Loewenstein, George, Roberto Weber, Janine Flory, Stephen Manuck, and Matthew Muldoon. 2001. "Dimensions of Time Discounting." Paper presented at Conference on Survey Research on Household Expectations and Preferences, Ann Arbor, MI.

MacDorman, Marian F., Fay Menacker, and Eugene Declercq. 2008. "Cesarean Birth in the United States: Epidemiology, Trends, and Outcomes." *Clinics in Perinatology*, 35(2): 293–307.

Madden, Gregory J, and Patrick S Johnson. 2010. "A Delay-Discounting Primer." In *Impulsivity: The Behavioral and Neurological Science of Discounting.* , edited by Gregory J. Madden and Warren K. Bickel Bickel, 11–37. Washington, D.C.: American Psychological Association.

Mahoney, Neale. 2015. "Bankruptcy as Implicit Health Insurance." *American Economic Review*, 105(2): 710–46.

Manski, Charles F. 2004. "Measuring Expecations." *Econometrica*, 72(5): 1329–76.

Maryland Insurance Administration. 2015. "2015 Report on the Effect of Competitive Rating on the Insurance Markets in Maryland." MARS# 995. http://insurance.maryland.gov/Consumer/AppealsandGrievancesReports/ Competitive-Rating-in-MD-Report-2015.pdf (accessed April 5, 2017).

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory.* New York, NY: Oxford University Press, Inc.

McCartt, A.T., L.L. Geary, and A. Berning. 2003. "Observational Study of the Extent of Driving While Suspended for Alcohol Impaired Driving." *Injury Prevention*, 9(2): 133–7.

McGuire, Thomas G., and Mark V. Pauly. 1991. "Physician Response to Fee Changes With Multiple Payers." *Journal of Health Economics*, 10(4): 385–410.

Miller, Mark E., Stephen Zuckerman, and Michael Gates. 1993. "How do Medicare Physician Fees Compare With Private Payers?" *Health Care Financing Review*, 14(3): 25–39.

Molina, George, Thomas G. Weiser, Stuart R. Lipsitz, Micaela M. Esquivel, Tarsicio Uribe-Leitz, Tej Azad, Neel Shah, Katherine Semrau, William R. Berry, Atul A. Gawande, and Alex B. Haynes. 2015. "Relationship Between Cesarean Delivery

Rate and Maternal and Neonatal Mortality." *Journal of the American Medical Association*, 314(21): 2263–70.

National Association of Insurance Commissioners. 2011. *A Consumer's Guide to Auto Insurance.* Washington, D.C.: National Association of Insurance Commissioners. http://www.naic.org/documents/consumer_guide_auto.pdf (accessed April 5, 2017).

National Association of Insurance Commissioners. 2012*a*. *2011 Competition Database Report.* Washington, D.C. http://www.naic.org/prod_serv/CLR-OPS-13.pdf (accessed April 5, 2017).

National Association of Insurance Commissioners. 2012*b*. *2011 Market Share Reports for Property/Casualty Groups and Companies.* Washington, D.C. http://www.naic.org/prod_serv/MSR-PB-12.pdf (accessed April 16, 2017).

National Association of Insurance Commissioners. 2014. *Auto Insurance Database Report 2011/2012.* Washington, D.C.: National Association of Insurance Commissioners. http://www.naic.org/prod_serv/AUT-PB-11_2014.pdf (accessed May 18, 2015).

North Carolina Department of Insurance. 2010. *A Consumer's Guide to Automobile Insurance.* Raleigh, NC: North Carolina Department of Insurance. http://www.ncdoi.com/_Publications/ConsumerGuidetoAutomobileInsurance_CAU1.pdf (accessed May 23, 2013).

Perreira, Krista M., and Frank A. Sloan. 2002. "Excess Alcohol Consumption and Health Outcomes: A 6-year Follow-up of Men Over Age 50 From the Health and Retirement Study." *Addiction*, 97(3): 301–10.

Pfuntner, Anne, Lauren Wier, and Carol Stocks. 2013. "Most Frequent Conditions in U.S. Hospitals, 2010." Agency for Healthcare Research and Quality, Healthcare Cost and Utilization Project (HCUP) Statistical Brief No. 162, Rockville, MD. https://www.hcup-us.ahrq.gov/reports/statbriefs/sb162.pdf (accessed March 13, 2015).

Polsky, Daniel, Michael Richards, Simon Basseyn, Douglas Wissoker, Genevieve M Kenney, Stephen Zuckerman, and Karin V Rhodes. 2015. "Appointment Availability After Increases in Medicaid Payments for Primary Care." *New England Journal of Medicine*, 372(6): 537–45.

Roberts, Eric T., Michael E. Chernew, and J. Michael McWilliams. 2017. "Market Share Matters: Evidence of Insurer and Provider Bargaining Over Prices." *Health Affairs*, 36(1): 141–8.

Robinson, Patricia A., Frank A. Sloan, and Lindsey M. Eldred. Forthcoming. "Advantageous Selection, Moral Hazard, and Insurer Sorting on Risk in the U.S. Automobile Insurance Market." *Journal of Risk and Insurance.*

Rothschild, Michael, and Joseph Stiglitz. 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics*, 90(4): 629–49.

Rust, John. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica*, 55(5): 999–1033.

Saito, K. 2006. "Asymmetric Information in Insurance: Some Testable Implications." *Journal of Risk and Insurance*, 73(2): 335–56.

Shavell, Steven. 2005. "Minimum Asset Requirements and Compulsory Liability Insurance as Solutions to the Judgment-Proof Problem." *RAND Journal of Economics*, 36(1): 63–77.

Sloan, Frank A., Lindsey M. Eldred, Tong Guo, and Yanzhi Xu. 2013. "Are People Overoptimistic About the Effects of Heavy Drinking?" *Journal of Risk and Uncertainty*, 47(1): 93–127.

Sloan, Frank A., W. Kip Viscusi, Harrell W. Chesson, Christopher J. Conover, and Kathryn Whetten-Goldstein. 1998. "Alternative Approaches to Valuing Intangible Health Losses: The Evidence for Multiple Sclerosis." *Journal of Health Economics*, 17(4): 475–97.

Sloan, Frank, Janet Mitchell, and Jerry Cromwell. 1978. "Physician Participation in State Medicaid Programs." *Journal of Human Resources*, 13(Supplement): 211–45.

Small, Kenneth A., and Harvey S. Rosen. 1981. "Applied Welfare Economics With Discrete Choice Models." *Econometrica*, 49(1): 105–30.

Sonchak, Lyudmyla. 2015. "Medicaid Reimbursement, Prenatal Care and Infant Health." *Journal of Health Economics*, 44: 10–24.

Spinnewijn, Johannes. 2013. "Insurance and Perceptions: How to Screen Optimists and Pessimists." *Economic Journal*, 123(569): 606–33.

Statista. 2015. "Advertising Spending of Selected Insurance Brands in the United States in 2015." https://www.statista.com/statistics/264968/ad-spend-of-selected-insurance-companies-in-the-us (accessed September 13, 2016).

Studdert, David M., Marie M. Bismark, Michelle M. Mello, Harnam Singh, and Matthew J. Spittal. 2016. "Prevalence and Characteristics of Physicians Prone to Malpractice Claims." *New England Journal of Medicine*, 374(4): 354–62.

Tai-Seale, Ming, Thomas H. Rice, and Sally C. Stearns. 1998. "Volume Responses to Medicare Payment Reductions With Multiple Payers: A Test of the McGuire-Pauly Model." *Health Economics*, 7(3): 199–219.

Tennyson, Sharon. 2010. "Incentive Effects of Community Rating in Insurance Markets: Evidence From Massachusetts Automobile Insurance." *Geneva Risk and Insurance Review*, 35(1): 19–46.

Thielking, Megan. 2015. "Sky-high C-section Rates in the US Don't Translate to Better Birth Outcomes." *STAT*, December 1. https://www.statnews.com/2015/12/01/cesarean-section-childbirth.

Tilburt, Jon C., Matthew K. Wynia, Robert D. Sheeler, Bjorg Thorsteinsdottir, Katherine M. James, Jason S. Egginton, Mark Liebow, Samia Hurst, Marion Danis, and Susan Dorr Goold. 2013. "Views of US Physicians About Controlling Health Care Costs." *Journal of the American Medical Association*, 310(4): 380–8.

Ubel, Peter A. 2012. *Critical Decisions: How You and Your Doctor Can Make the Right Medical Decisions Together.* New York, NY: HarperOne.

U.S. Department of Health and Human Services. 2010. "Maternal, Infant, and Child Health." https://www.healthypeople.gov/2020/topics-objectives/topic/maternal-infant-and-child-health/objectives (accessed February 24, 2017).

Voas, Robert B., A. Scott Tippetts, and A. Scott McKnight. 2010. "DUI Offenders Delay License Reinstatement: A Problem?" *Alcoholism: Clinical and Experimental Research*, 34(7): 1282–90.

Weisburd, Sarit. 2015. "Identifying Moral Hazard in Car Insurance Contracts." *Review of Economics and Statistics*, 97(2): 301–13.

Wise, Michael G, and James R. Rundell. 1994. *Concise Guide to Consultation Psychiatry.* 2nd ed. Washington, D.C.: American Psychiatry Press.

Zafar, Basit. 2011. "Can Subjective Expectations Data be Used in Choice Models? Evidence on Cognitive Biases." *Journal of Applied Econometrics*, 26(3): 520–44.

Zuckerman, Stephen, Joshua McFeeters, Peter Cunningham, and Len Nichols. 2004. "Changes in Medicaid Physician Fees, 1998-2003: Implications for Physician Participation." *Health Affairs*, 23(Web Exclusives): W4–374–84.

Biography

Patricia Alexander Robinson was born in Washington, D.C., on June 20, 1987, and was given the nickname "Alex" by her parents. She lived with her family in Arlington and Alexandria, VA, and Orinda, CA—where she attended Miramonte High School—before moving to Seattle, WA, where she graduated from Lakeside School in 2005. Alex graduated *magna cum laude* from Whitman College in Walla Walla, WA, in 2009 with a bachelor's degree in Economics-Mathematics. After working for three years as an actuarial analyst at Regence BlueShield in Seattle, she enrolled at Duke University, where she earned her Ph.D. in Economics in 2017 under the supervision of Frank Sloan.