

Adverse Selection and Moral Hazard in a Dynamic Model of Auto Insurance*

Przemyslaw Jeziorski[†] Elena Krasnokutskaya[‡] Olivia Ceccarini[§]

July 8, 2016

Abstract

We measure risk-related private information and investigate its importance in a setting where individuals are able to modify risk ex-ante through costly effort. Our analysis is based on a model of endogenous risk production and contract choice. It exploits data from multiple years of contract choices and claims by customers of a major Portuguese auto insurance company. We additionally use our framework to investigate the relative effectiveness of dynamic versus static contract features in incentivizing effort and inducing sorting on private risks, as well as to assess the welfare costs of mandatory liability insurance.

Keywords: insurance, adverse selection, moral hazard, dynamic demand

JEL Classification: D82, G22

*We would like to thank Jaap Abbring, Phil Haile, Steve Berry, Ariel Pakes, Robin Lee, Hanming Fang, and Robert Miller for helpful discussions. We are also grateful to seminar participants at Carnegie-Mellon University, Harvard University, Boston College University, Rice University, and Northwestern University, as well as participants of 2015 Washington IO Day Conference, 2015 Cowles Foundation Conference for Applied Microeconomics, 2016 Barcelona Summer Workshop for Applied Industrial Organization, and 2016 Society for Economic Dynamics Annual Meeting.

[†]Haas School of Business, email: przemekj@haas.berkeley.edu

[‡]Johns Hopkins University, email: ekrasno1@jhu.edu

[§]Marie Curie Fellow, Porto Business School, Portugal

1 Introduction

Economic theory has long argued that private information about idiosyncratic risks may adversely affect the functioning of insurance markets.¹ As a result, measuring private information about idiosyncratic risks has been at the center of empirical analysis. Recently, however, researchers have become increasingly concerned that, in many insurance markets, individuals are also able to modify risk in response to incentives (moral hazard). On the one hand, this possibility must then be taken into account when assessing the importance of private information. On the other, it implies that various contract features chosen by the industry may affect both sorting on risks as well as the choice of risk. The extent to which they do so in practice is an open question. Moreover, measuring the ability of alternative contract designs to affect these two margins is essential for assessing available policy options.

We consider these issues by studying the market for auto insurance that has repeatedly been the focus of empirical studies inquiring into the importance and implications of private information.² An important feature of the auto insurance market – which is one of the most well established and prominent insurance industries – is that it tends to employ broadly similar sets of contracts and is subject to fairly similar regulations in most countries around the world. While even this industry continues to evolve, the overall uniformity of business practices suggests that it has been able to reach a consensus on some fundamentally important issues it faces. In particular, a common feature of contracts in this market is that they combine variation in coverage with experience rating which ties contract premium to the recent realizations of an individual’s risk. While experience rating may help screen individuals on risk, it also suggests a possibility that insurers may aim to incentivize reduction in risk, i.e., they may believe that risk is modifiable. Building on this insight and exploiting various features of country-specific settings, several studies have proposed tests for moral hazard in auto insurance markets and indeed confirmed its presence.³

The potential for the presence of moral hazard in this market is also quite intuitive. While

¹The extensive literature exploring equilibrium effects associated with asymmetric information about idiosyncratic risks started with seminal work by Rothschild and Stiglitz (1976) and Wilson (1977). For a recent survey of this literature see Mimra and Wambach (2014).

²An incomplete list of papers testing for the importance of private information in the context of this market include Pueltz and Snow (1994), Chiappori and Salanie (2000), and Cohen (2005). Cohen and Einav (2007) measure the importance of private information using a structural model of contract choice which allows for private information about risk and risk aversion. There exists an extensive literature inquiring in these issues in other markets: for a review see Einav, Finkelstein, and Levin (2010).

³For details see Abbring, Chiappori, and Piquet (2003), Dionne, Pinquet, Maurice, and Vanasse (2011) and Israel (2004).

driving ability may differ across individuals, the probability of having an accident is directly affected by the individual's decision of how much and in which conditions to drive. A good driver who drives often might have the same probability of an accident as a bad driver who drives rarely or drives only when the driving conditions are good. Both of these drivers can even achieve zero probability of accidents by parking their car and, say, relying on public transportation. The cost of doing so is likely idiosyncratic and only privately known. In contrast to the frequency of accidents, the size of damages conditional on the accident is likely to depend on the circumstances rather than on individual's choices about his driving. The industry practice of pricing on the number of the accidents rather than their severity indicates that practitioners also believe that the size of the damages is not specific to individual's ability or effort at minimizing risk.

Motivated by the intuitive appeal and formal econometric evidence that confirms the presence of moral hazard in this market, we conduct our analysis using a framework, which includes the ability of agents to adjust their risk ex-ante through costly effort while allowing for agents to be heterogeneous in their cost of effort and the degree of risk aversion.⁴ We use this framework to measure private variability in risk-related attributes and to gain insight into the factors underlying design of insurance contracts in this market.

Specifically, we consider a dynamic finite horizon model of auto insurance with individuals entering the model upon obtaining the driving license. Subsequently, each period, an individual decides whether to purchase the basic (liability only) or extended coverage. He also chooses the level of risk, i.e., the probability of having an accident, which he controls through costly effort. We allow individuals to differ in their cost of effort and risk aversion. These factors may depend on demographics and car characteristics that are observed by the insurance company. In addition, we allow for the residual components which cannot be captured by the observables and remain private information of the driver. The model is dynamic because contracts are experience rated meaning that the price of the contract is increasing in the risk class and as a result in the number of recent accidents caused by the individual. The final horizon aspect of the model allows us to integrate the possibility of learning with driving experience, as well as age-related changes in the impact of incentives embedded in contract pricing.

We use data from a major Portuguese auto insurance company on a panel of all individuals

⁴We focus on ex-ante moral hazard, that is, an adjustment to the individual-specific frequency with which accidents realize. This is in contrast to ex-post moral hazard where individual decides whether to report the accident to the the insurance company once it occurred. Analysis of our data below indicates that ex-post moral hazard is unlikely to play an important role in the market we study.

it covered during the period between 2004 and 2010. The Portugal auto insurance industry relies on the risk rating system which assigns each driver to one of 18 risk classes (assignment is adjusted every year). Our data contain all the information collected by the insurance company on individuals' demographics, his car's characteristics, as well as individual history of contract choices (e.g., liability only or comprehensive, etc), and realized claims during the time the individual was enrolled with this insurance company. We also observe the risk rating for every year the individual appears in the sample. Furthermore, we observe the discount on the price of the contract given to an individual. The insurance company authorizes the agents selling its contracts to offer such discounts to potential clients in order to match their sales targets. These discounts could be as large as 25% (10% on average), but the probability of an individual receiving a discount and its size appear quite random.

An important objective of the insurance literature has been to infer the variability of risk and risk aversion. It is important to recognize, however, that if risk is modifiable, its variability is tied to the existing set of incentives. To understand the response of risk to alternative configurations of incentives, it is necessary to recover deeper primitives that contribute to the production of risk, such as cost of effort and risk aversion. This is the objective of our empirical analysis. To achieve this objective, our identification strategy exploits additional sources of variation in the data relative to the literature.

Traditionally, the literature has relied on cross-sectional data and variation induced by the heterogeneity of contracts with respect to the degree of coverage. Assuming that individual risk types are fixed, to the extent such contracts sort drivers on risk, the variation of realized risk across contracts is indicative of the underlying variability of risk in the population. Moreover, sorting across contracts among individuals with the same risk is indicative of the heterogeneity in risk aversion. Such an identification strategy may, however, lead to erroneous conclusions if risk is, in fact, modifiable. The reason is that the realized risk is the function of the contract the individual is in. Thus, the variation in realized risk across contracts may not be informative of the true heterogeneity in population, while individuals with the same risk in different contracts may differ not only in their risk aversion but also in their cost of modifying risk. Instead, as has been pointed out by Abbring, Chiappori, Heckman, and Piquet (2003), to distinguish moral hazard from adverse selection it is necessary to observe the same driver (or the same population of drivers) in several settings with different incentives for adjustment of risk in place. Our data allow us to use additional sources of variation to achieve identification. First, even with cross-

sectional data, we can achieve identification by comparing identical populations of drivers who were exogenously given different discounts and who chose the same contract. Second, the panel structure of our data allows to identify the model by exploiting within individual variation in risk as the individual progresses across risk classes (or the variation in the distribution of risks within the cohort of drivers and across several years which reflects changes in individuals' sorting across risk classes and contracts over time). We use both sources of identification in estimation for efficiency, but we find that using them separately leads to the same estimates.

The estimated model is successful in rationalizing observed risk production along multiple dimensions of incentives, including those associated with dynamic incentives related to movement across risk classes, as well as incentives embedded into contracts with different degrees of coverage (traditional moral hazard). It also accounts for the large response of risk to the size of the discount. The estimated model is quite parsimonious. It allows for variable risk aversion and two dimensions of the cost of effort, i.e., the heterogeneity in the level of the cost (in the extreme, the idiosyncratic cost of parking the car for a month depends on the availability and costs of using an alternative mode of transportation) and the heterogeneity in the responsiveness of the probability of an accident to individual effort (this naturally reflects the heterogeneity in individual ability, but it also captures the heterogeneity in how flexible the individual is with respect to timing of trips and the ability to avoid traffic congestion).

Our estimation results reveal significant private variation in parameters governing the cost of effort and risk preferences. This translates into non-trivial private variation in idiosyncratic risk. Both cost and preference parameters are important in generating variation in risk, although the effect of the variation in costs is three times more important.⁵ It is important to keep in mind that this variation in risk is conditional on the incentives present in the market. The estimated primitives allow us, however, to gain, through a series of counterfactual experiments, insights into considerations shaping the current menu of contracts as well as pricing strategies used in this market.

We find that experience rating in its current design does not result in effective separation of risk types: in simulations, even after 40 years of driving, the realized risk class explains less

⁵The importance of private information about the cost of effort seems also corroborated by the direction of recent changes in contract design in this market, which appear to put more emphasis on eliciting the cost of effort rather than the risk aversion. Innovative contracts offered by, e.g., the U.S. auto insurer *Progressive*, among others, tailor pricing to individual driving patterns. Using a device installed in the insured car, the company records timing of driving, traffic conditions, total amount of time and distance traveled, etc. This reveals some of the private information about the cost of effort and allows to not only better tailor prices to risks but also to provide incentives to modify those risks.

than 5% of variation in the cost and preference parameters. The reason is that car accidents are inherently random with significant movement of individuals across risk classes due to pure chance. This suggests tight limits on the ability of insurance companies to separate drivers based on unobserved risk attributes through experience rating. Despite this, the existing pricing based on risk rating system is very effective at incentivizing risk provision and holding the overall risk in the system low. This perhaps explains why risk rating systems feature so prominently in the design of auto insurance markets across time and across countries.

The design of the experience rating system appears targeted at providing particularly strong incentives to young drivers. By law, all new drivers are initially assigned to risk class 10, considerably above the average risk class in the population. The pricing schedule increases dramatically in a highly convex fashion right after class 10. This exposes young drivers to the threat of a severe monetary penalty for causing an accident. A natural way to interpret learning in the context of our model is as a decline in the cost of effort needed to achieve a certain risk level. Our estimates indeed reveal important differences in the costs of young and experienced drivers. That is, we find that young drivers are least able to modify risk and need significant incentives in order to do so.

On the other hand, differential coverage appears to be more effective as a sorting device. However, it has a limited potential for providing incentives because the large majority of drivers purchase only liability coverage and the insurers are legally prohibited from offering partial coverage on these contracts. The reason for the low prevalence of extended coverage contracts appears rooted in insurers' concern about the potential for upward risk expansion, which explains observed pricing of the extended contracts in a way that effectively excludes drivers with cheap cars (who according to our estimates have higher tolerance for risk), as well as drivers with low experience who have higher costs of controlling risk and thus are inherently riskier drivers.

The ability of drivers to modify risk and the propensity of contracts with differential coverage to induce sorting as well as to impact incentives for risk production suggest a possibly significant cost of legal restrictions on mandatory liability coverage prevalent in European car insurance markets. We illustrate this point by considering an alternative menu which includes a contract with partial liability coverage. We find that such a menu is capable of inducing a substantial reduction in the total number of accidents. Moreover, partial coverage not only incentivizes reductions in risk but also increases welfare of "safe" drivers by allowing them to reduce their insurance premium expenditure. Thus, our analysis uncovers substantial social costs of these regulations which should be taken into account when selecting among regulatory designs that may achieve the same ultimate

policy objectives.⁶

The paper is organized as follows. Section 2 summarizes the model and the relevant theoretical concepts. Section 3 describes specific features of our data and Section 4 explores descriptive evidence of moral hazard. Section 5 explains the estimation methodology. Section 6 reports estimation results while Section 7 presents our analysis of the mandatory liability insurance. Section 8 concludes.

2 Model

We begin by developing a formal model which rationalizes choices made by drivers while participating in the car insurance market. This framework defines the objects and processes that we study and serves as an important ingredient for the econometric model which we take to the data.

We consider an individual who upon obtaining the driving license at some time t_1 becomes affiliated with insurance company A. We follow this driver over time until the age of $T = 90$, which is the legal limit on the age of driving in Portugal. Driving a car exposes an individual to the risk of ‘at fault’ accidents and specifically to the risk of expense associated with own damages (the damages to the third party from “at fault” accidents are covered by mandatory own liability insurance). At the beginning of each period an individual decides whether to enroll in the basic liability coverage only or to purchase additional comprehensive coverage that (up to a small deductible) protects him from the risk of damages to his own car. The individual further decides on the level of risk (the distribution of the number of “at fault accidents”) he would like to maintain in a given year which is summarized by the parameter λ_t . The individual’s decisions reflect his risk aversion and his cost of achieving a given level of risk summarized by parameters $(\gamma; \theta)$ respectively. The individual may also leave company A with a fixed probability ρ in a given period.⁷

⁶Much of the discussion about the mandatory auto insurance laws with uniform coverage centers on the mechanisms needed to ensure that a victim in a car accident receives compensation from the guilty party. The significant increase in the number of accidents due to the current regulatory restrictions should be considered when assessing alternative policies such as certifying sufficient wealth, posting a bond, etc that might achieve the goal of covering the victims of an accident while providing better incentives to the drivers. Versions of such policies are implemented in some US states, e.g. New Hampshire.

⁷There are a number of reasons for an individual to exit a market, such as disease, death or loss of a car. The individual may also leave company A to sign up with a competing insurance company. In Portugal, drivers are effectively required to carry their risk class with them as they move from company to company. In addition, the insurance companies use the same scaling coefficients associated with movement across risk classes. Anecdotal evidence suggests that individuals usually move because they have been offered a better price discount by a competitor of A. Since the discounts cannot depend on individual’s private factors, such attrition does not result in a selected

Risk Exposure. An individual’s risk exposure depends on the contract he chooses, Y_t , his idiosyncratic risk, λ_t , and the distribution of damages to his car under an “at fault” accident, F_L . In order to characterize this object we introduce some additional notation.

Let us denote the number of “at fault” accidents in a given period by R_t and the associated vector of monetary damages to own car incurred in these accidents by \bar{L}_t with $L_{r,t}$ reflecting the damage from accident r . We assume that the number of accidents follows a Poisson distribution with parameter λ_t chosen by the individual.⁸ Following previous studies we assume that the distribution of $L_{r,t}$ is independent of λ . We use the function $C(R, \bar{L}; Y, \lambda)$ to summarize the individual’s risk exposure if he chooses contract Y and the level of risk λ . Specifically,

$$C(R, \bar{L}; Y, \lambda) = \begin{cases} R\bar{C} + \sum_{r=1}^R L_r & \text{if } Y = y^L \\ R\bar{C} + \sum_{r=1}^R \min \{L_r, D\} & \text{if } Y = y^C \end{cases}$$

where \bar{C} summarizes accident costs that are not included in damages assessed by the insurance company, such as monetarized health deterioration, convenience or psychic costs, D denotes the deductible specified in the comprehensive contract.

Cost of Effort. An individual is able to maintain the level of risk at λ by paying cost $\Gamma(\lambda; \theta)$ such that $\Gamma(\lambda; \theta) \geq 0$, and $\Gamma'(\lambda; \theta) \leq 0$.

Specifically, we assume that

$$\Gamma(\lambda; \theta) = g_0 + \frac{\theta_1}{1 + \theta_2 \lambda}$$

with $\theta_1 > 0$ and $\theta_2 > 0$. Parameters θ_1 and θ_2 jointly determine the slope and the curvature of the cost function (or alternatively the level and the slope of the marginal cost of decreasing risk).

Notice that in our specification it is possible to achieve $\lambda = 0$ at potentially a high cost. Such a situation would arise if the individual uses his car very rarely (for example, only in an emergency), possibly because of steep incentives at high risk classes.

Our model does not nest the case of ‘no moral hazard’ in a sense that adjustment of risk is possible at all non-zero risk levels. However, the model is capable of characterizing environments where risk adjustments in response to incentives are very small (so that they are not empirically

sample in the environment without switching costs. In this market insurance companies actively solicit customers (in contrast to the situation where individuals search for a better deal) so the absence of (or small) switching costs is not implausible.

⁸ In estimation, we distinguish between three types of accidents: (a) type 1: damage to own car, no third party involved; (b) type 2: accidents involving third party with damage to own car; and, (c) type 3: accidents involving third party without damage to own car. We assume that the type of accident is exogenously determined and the distribution of losses may depend on the type of accident.

relevant). Such outcomes arise, for example, when the curvature of the cost of effort function ($\Gamma''(\lambda; \theta)$) is sufficiently large.⁹

Contract Pricing. Insurance contract pricing is based on experience rating. An individual is assigned to a liability and comprehensive risk class for every period that he stays in the market. We summarize the individual's risk classification by vector $M_t = (K_t^L, K_t^C)$ such that $M_1 = (10, 10)$. The risk classes evolve according to the deterministic functions of the total number of related accidents (the number of "at fault" accidents with damage to the third party, $\tilde{R}_t^{(1)}$, for the liability component and the number of "at fault" accidents with positive damages to own car, $\tilde{R}_t^{(2)} = \sum_r^R 1(L_{r,t} > 0)$ for the comprehensive component if the individual is enrolled in comprehensive contract).

Baseline contract price is computed as a function of the individual's demographic characteristics and characteristics of his car, Z_t . It is then multiplied by a risk-class-specific scaling coefficient which is increasing in the risk class. The individual therefore anticipates that as his risk class changes so does the price he has to pay for contract Y_t . We denote price of contract Y_t by $p(Y_t, Z_t, M_t)$ to recognize this dependence.

Payoffs. An individual's preferences are summarized by the within-period utility function

$$U(\bar{w} + \pi; \gamma) = (\bar{w} + \pi) - \gamma(\bar{w} + \pi)^2,$$

where \bar{w} is a constant and π represents all expenses associated with the car insurance market. The expenses in a given period are depend on the realized accidents, the associated losses as well as the contract and risk level chosen by the individual, and of his risk classification. Specifically,

$$\pi(R, \bar{L}; Y, \lambda; Z, M) = -p(Y, Z, M) - C(R, \bar{L}; Y, \lambda) - \Gamma(\lambda; \theta).$$

The attractive property of this specification is that it represents a re-parameterized version of the utility function which explicitly accounts for individual's wealth.¹⁰ Under such re-parameterization

⁹In practice, even moderate values of $\Gamma''(\lambda; \theta)$ may generate negligible risk adjustments.

¹⁰The specification we use is a re-parametrized version of the following within-period utility function

$$U(x; w_i, \tilde{\gamma}_i) = (w_i + x) - \tilde{\gamma}_i(w_i + x)^2,$$

where w_i denotes the individual's wealth so that $\gamma_i = \frac{\tilde{\gamma}_i}{1 - 2(w_i - \bar{w})\tilde{\gamma}_i}$. In our context, w_i should be interpreted as a wealth category which is perpetually constant during the individual's driving career. We normalize \bar{w} at a value such that $\bar{w} + L > 0$ for the range of losses, L , observed in our sample. Such normalization is without loss of generality since the coefficient of risk aversion does not depend on \bar{w} .

the relevant individual heterogeneity is summarized by coefficient γ , and the absolute coefficient of risk aversion remains unchanged. This specification is very convenient in empirical work since the data for an individual's wealth are rarely available. Since the insurance company also lacks information on an individual's wealth, coefficient γ correctly reflects an individual's private information about his risk aversion.

Optimization Problem and Bellman Equation. The state of the individual's decision problem is summarized by a vector $s = (\gamma, \theta, Z, M)$; the utility and cost function parameters are included because they may change over time.¹¹ We assume that components of s follow Markov processes:

$$M_{t+1} = f_M(\tilde{R}_t^{(1)}, \tilde{R}_t^{(2)}, M_t, Y_t)$$

$$(\gamma_{t+1}, \theta_{t+1}, Z_{t+1}) \sim F_{\gamma, \theta, Z}(\cdot | \gamma_t, \theta_t, Z_t).$$

An individual decides on a policy function which maps his state into a contract choice and risk levels $\mathbf{g}_t(s) = (\mathbf{y}_t(s), \boldsymbol{\lambda}_t(s))$ to maximize for all $t \in \{1, \dots, T\}$

$$V_t(s) = E_{\mathbf{g}} \left\{ \sum_{l=t}^{\min\{\tau-1, T\}} \beta^{l-t} [U(\bar{w} + \pi(R_l, L_l; Y_l, \lambda_l, Z_l, M_l))] \Big| s_t = s \right\}, \quad (1)$$

where τ is the stopping time, reflecting exogenous exit.¹²

The Bellman equation for the above problem is given by

$$V_t(s_t) = (1 - \rho) \max_{Y_t, \lambda_t} E_{R_t, L_t} \left[U(\bar{w} + \pi(R_t, L_t; Y_t, \lambda_t, Z_t, M_t)) + \beta V_{t+1}(s_{t+1}) \Big| Y_t, \lambda_t, s_t \right], \quad (2)$$

with a terminal condition $V_T = 0$.

Next, we describe the data used in this analysis.

3 Data Description

Our data are provided by a major Portuguese insurance company. For the reasons of confidentiality we cannot name this company; in subsequent exposition we will refer to it as company A.

¹¹In estimation we allow the costs of effort to evolve in a manner consistent with learning whereas an individual's risk aversion may change as a deterministic function of age.

¹²The stopping time τ is distributed as a Pascal distribution with parameter ρ , which indicates $\tau - 1$ consecutive failures and one success in the series of Bernoulli trials with a success probability ρ .

	Mean	Std. Dev.	5%	25%	50%	75%	95%
Male	0.746						
Age	35.42	4.83	26	29	33	42	45
Age of first-time drivers	29.67	5.04	19	28	33	40	45
Driving experience	10.83	3.54	4	8	10	11	13
Car value, €1,000	6.28	5.98	1.57	2.34	3.91	8.05	18.08
Car weight, 1,000kg	1.27	0.51	0.81	0.94	1.13	1.43	2.52
Car horse power	80.15	25.77	50	60	75	90	130.5
Observations	12,576						

Table 1: Data Summary Statistics.

The sample covers the period between 2004 and 2010. The data contain all the individual-level information used by the insurance company in pricing of the contract: consumer demographics (gender, age, years of driving experience, zip code) and car characteristics (car value, car horse power, car weight, car make and car age).¹³ For every driver we know the year when he joined company A for the first time and the year when his latest spell with this company began. Additionally, for every driver and for every year in the sample, we observe this driver’s risk class, his contract choice, and his insurance premium. We further have access to information on all claims filed by this driver during the sample years. For each claim we observe the date, the size, and whether the claim relates to the third-party or own losses.

For reasons that will be explained later we focus our attention on the subsample of individuals who started their participation in the car insurance market by signing a contract with company A upon obtaining their driving license and have continued their association with this company until and including part of the period covered in the data.¹⁴

Table 1 reports some basic statistics about our sample. As can be seen from the table our sample consists predominately of male drivers; an average driver is 35 years old and has close to 11 years of driving experience. Five percent of drivers in our sample have been driving less than five years. Generally, insureds obtain a driver’s license later in life relative to the US population (average age of first-time drivers is 30 and median age is 33). An average driver owns a car valued at €6,200 with the median car valued at €4,000 .

Risk classes, Contracts and Prices. Comparably to other auto insurance markets, Portuguese car insurance companies offer two types of contracts: basic insurance that covers damages

¹³The company uses several zip code bins to price its policies. We concentrate on the four largest bins comprising more than 98% of driver population. We drop the remaining 2%.

¹⁴We further restrict our sample to drivers with passenger cars who do not use their cars for commercial purpose, and exclude drivers whose cars are leased.

Risk class	Liability Component	Comprehensive Component	Risk class	Liability Component	Comprehensive Component
1	45%	45%	10	100%	100%
2	45%	45%	11	110%	110%
3	50%	45%	12	120%	120%
4	55%	45%	13	130%	130%
5	60%	60%	14	150%	150%
6	65%	65%	15	180%	150%
7	70%	70%	16	250%	150%
8	80%	80%	17	325%	150%
9	90%	90%	18	400%	150%

Table 2: Risk Class Adjustment Coefficient over Base Insurance. The table reports the coefficients which are used multiplicatively to adjust base premium for risk class. Base premium is computed on the basis of individual’s characteristics reported to insurance company.

to the third party (liability insurance) and extended coverage which in addition to liability insurance also includes a component which covers the damage to the own vehicle (comprehensive insurance). The liability insurance is mandatory in Portugal. Pricing of both types of contracts is experience rating based.¹⁵ Under this system each policyholder is placed in one out of 18 experience classes on the basis of their history of claims (a separate risk class is maintained for liability and comprehensive type of insurance). Beginner drivers start in class ten. Every year the experience class is updated: if the policyholder did not have any relevant claims (third party claims for liability part and own-vehicle claims for comprehensive part) in the previous year then his experience class is reduced by one. For every claim that he had in the previous year he is moved three classes up. The contract prices are increasing in the risk class. Specifically, the insurance company first constructs the base premium which reflects the driver’s characteristics reported to the insurance company. This is the premium which is charged to drivers in the reference class. For all other classes, the premium is adjusted multiplicatively so that it is scaled up/down for the risk classes above/below the reference class respectively. Any claim in which the policyholder is at least partially at fault, triggers upward transition. While the history of an individual’s claims is not necessarily public knowledge, a policyholder who switches insurance companies and is not providing his new insurer with his/her claims record gets automatically placed in a class 16 (that is in the class where he would end up if he had 2 accidents in his first year of driving). Table 2 summarizes the slope of the premium function with respect to the risk class. As this table indicates, the premium charged to drivers differs significantly across classes. Specifically, the premium

¹⁵Barros (1996) provides historical context of these regulations

schedule becomes significantly convex above class ten.

Table 3 summarizes the distribution of drivers in our data across the risk classes and contracts as well as reports the average annual premium paid by insureds for a given class and contract. Company A offers two types of comprehensive coverage: the first type of comprehensive contract imposes a 2% deductible whereas the second type offers full coverage. A very small fraction of insureds (less than 0.1%) choose the second type of contract. Therefore, we consider only the first type of contract in our analysis. The majority of observations (99%) are associated with risk classes one through ten (with the largest share corresponding to class one), and for every risk class most observations are for the individuals who choose to buy only liability coverage. Such risk-class distribution results from an overall low probability of accidents and is also observed in other auto insurance markets (for example, Dionne, Michaud, and Dahchour (2010) and Dionne, Michaud, and Dahchour (2013) document similar distributions of consumers across risk-classes in Quebec and France).

The third column of Table 3 reflects the base portion of the liability premium (set for class ten) for individuals who in our data are observed in various risk classes. It indicates that even an average base premium is roughly increasing in the risk class. This regularity is primarily driven by the fact that the insurance company charges higher premium to younger individuals and individuals with low driving experience who are necessarily located in higher risk classes. The disparity in premiums across classes is quite striking: an individual just entering the system on average has a base premium which is twice as high as the base premium paid by an individual in class one. Column four shows the average of the liability premiums after they are adjusted for the risk class. The difference in adjusted premiums is even more striking with the individuals in high classes paying up to four times more than individuals in risk class one.

Column seven summarizes the comprehensive part of the premium. In contrast to the premium for the liability portion of the contract, the premium for comprehensive coverage depends on the own car value. For individuals who purchase comprehensive coverage in our sample, comprehensive premium tends to be almost twice as high as the liability premium for the comparable risk class. Thus, individuals purchasing comprehensive coverage on average spend three times as much on car insurance relative to individuals purchasing just the liability portion. Not surprisingly, they tend to be wealthier as indicated by much higher values of cars owned by these individuals (columns five and eight).

As was mentioned earlier, in Portugal, insurance contracts are sold by agents who can provide

Risk Class	Liability				Comprehensive		
		Average Base	Average Adjusted	Average Car	Average Adjusted	Average Car	
	Obs	Premium	Premium	Value	Obs	Premium	Value
1	8739	536.6	241.5	4939.7	841	467.2	16029.4
2	1369	521.4	234.6	4847.3	151	477.9	16673.3
3	539	481.4	240.7	5153.8	73	425.8	14091.9
4	491	513.7	282.6	4987.4	62	429.4	13384.2
5	234	543.1	325.8	4944.1	29	655.3	14157.6
6	196	583.9	379.5	5211.9	29	676.7	13077.7
7	186	632.5	442.8	5956.4	31	792.9	13087.6
8	206	955.9	764.7	6117.9	34	1311.1	16395.9
9	321	1154.1	1038.6	5781.3	36	1649.1	17127.2
10	253	1169.7	1169.7	5674.9	26	1755.4	15777.3

Table 3: Risk Classes, Contracts and Premiums. This table summarizes the distribution of observations in our sample across risk classes and contracts, as well as reports the averages for various premium components.

discretionary discounts lowering the base premium levels for liability and comprehensive parts of the contract. The discount is usually given when an individual first signs the contract with the insurance company. It very rarely varies within person after that (for 99% of drivers in our sample the discount does not change during their tenure with the insurance company). We eliminate the drivers with a changing discount from our sample and assume that the discount given upon first signing the contract is permanent.¹⁶

Further, anecdotal evidence suggests that the majority of discounts are given out primarily in order to achieve sales targets (for example, higher discounts are given mostly at the end of the year). While not important for the analysis to follow, we do not find any statistically significant relationship between the size of the discount and individual’s demographics and/or car characteristics for the sample we use in estimation (consisting of drivers who sign their first insurance contract with company A). It is therefore even more unlikely that the size of the discount is related to drivers’ unobserved characteristics in this sample.

Importantly, the discounts are applied as a percentage of the base price, therefore, since the risk class adjustment is also multiplicative, the discounts change the slope of premium schedule across risk classes and with it the incentives for risk adjustment. Effectively, drivers with lower discounts

¹⁶We find no evidence that the discount is contingent on the type of the contract. In addition, such concerns are not very relevant for our estimation sample, which consists of the drivers whose first contract is with company A, since the overwhelming majority of young drivers purchase liability only contract. As we noted before, in the data individuals retain their original discount throughout their association with company A and in particular whenever they change a contract. This regularity holds both when an individual switches from liability to extended coverage and when he switches from the extended to liability only coverage.

face stronger incentives to drive safer on the margin. Thus, discounts provide exogenous cross-sectional price variation which maybe used to identify the magnitude of risk adjustments in this setting. Additionally, such risk adjustments would be clearly motivated by monetary incentives as opposed to fear or other factors.

Risk and Associated Expenses. Table 4 summarizes risk associated with “at fault” accidents. As the table indicates, an average driver has 0.037 “at fault” accidents which results in damage to the third party on average in any given year. Younger drivers face a higher risk of 0.061 accidents on average. Further, the variability of risk in the population of young drivers is higher relative to the general population. The drivers choosing only liability coverage appear to be slightly safer than the general population.

	Obs	Liability Claims		Comprehensive Claims	
		Mean	Std. Dev.	Mean	Std. Dev.
Number of Claims					
All Drivers	12,576	0.037	0.193		
Young Drivers (≤ 5 years)	629	0.061	0.245		
Liability Contract Only					
All Drivers	11,252	0.036	0.192		
Young Drivers (≤ 5 years)	503	0.062	0.249		
Comprehensive Contract					
All Drivers	1,324	0.043	0.203	0.076	0.282
Young Drivers (≤ 5 years)	126	0.045	0.210	0.121	0.328

Table 4: Number of Claims. This table summarizes the mean and the standard deviation for the number of claims filed by an individual in a given year by the type of the contract and the type of the claim. ‘Liability claim’ refers to a claim associated with the third-party damages that arise from an ‘at fault’ accident. ‘Comprehensive claim’ refers to a claim associated with the damage to own vehicle incurred during an ‘at fault’ accident.

The drivers who choose comprehensive coverage are associated with a higher number of liability claims relative to the general population. This regularity may arise either due to selection of inherently “riskier” drivers into the contract with higher coverage (adverse selection) or because relaxed incentives associated with higher coverage result in lower effort at risk reduction and thus higher risk (moral hazard). The number of claims associated with damage to own car filed by individuals enrolled in a comprehensive contract significantly exceeds the number of liability claims (0.076 vs. 0.043 or 0.121 vs. 0.045 for young drivers). This is, perhaps, not very surprising since comprehensive claims cover a single car as well as multiple car accidents whereas liability claims apply only to multiple car accidents.

Table 5 provides information on the losses associated with “at fault” accidents. The average liability claim is equal to €1,784 whereas a median claim is €879. The claims can be quite small

	Mean	Std. Dev.	5%	25%	50%	75%	95%
Liability claim (€1000)	1.784	6.474	0.238	0.723	0.879	1.200	4.313
Comprehensive claim (€1000)	2.418	3.418	0.258	0.652	1.236	2.600	9.142
Comprehensive claim (relative to the value of own car)	0.176	0.257	0.013	0.039	0.084	0.202	0.715

Table 5: Claims' Sizes. This table summarizes the distribution of sizes of claims filed by individuals in a given year by the type of the claim.

(€238 (at 5% quantile of the claims distribution) and also quite substantial (€4,313 at the 95% quantile of the claims distribution). While these numbers certainly appear non-trivial recall that the average annual rate of accidents is 0.037. Thus a risk exposure of a risk neutral individual would only be €66 on average (with 5% - 95% inter-quantile range given by €8 to €160). Of course, exposure could be six times this amount at the upper end of the risk distribution (for an individual with 0.075 accidents on average and if the claim is €4,313). Similarly, an average comprehensive claim is €2,418 which is close to 18% of an individual's car value (median claim is €1,236 or 8% of an individual's car value). Computations similar to those above indicate that the risk exposure of an average driver to the risk of own car damage (if he is risk neutral) would be €104 (with median exposure equal to €53). Thus, the expected risk in the system is not very large while high risk exposure is possible with relatively small probability. In general, the average risk exposure for the risk neutral individual appears to be quite low relative to the premium he pays for the insurance coverage.

4 Descriptive Evidence of Moral Hazard

In this section we report some descriptive evidence of moral hazard, as defined by our model, and of individual heterogeneity underlying risk production. This analysis helps to build intuition for our identification strategy. It is worthwhile to remember that the coefficients obtained here cannot be easily related to the objects of interest in this study such as the magnitude of risk adjustments in response to specific incentives or the variability in the cost of effort/the degree of risk aversion across individuals in the population. We turn to the model-based estimation next in order to recover the model primitives which can be used to make such assessments.

We begin by regressing the number of liability claims on individual characteristics, risk class and the type of contract chosen by this individual. The results are summarized in Table 6. The coefficient on the risk class variable reflects competing effects. On the one hand, higher risk

	(1)	(2)	(3)	(4)	(5)
Risk Class	0.00356*** (0.000761)	0.00283** (0.00117)	-0.0865*** (0.00280)	-0.0940*** (0.00463)	-0.00590** (0.00243)
Risk Class×Discount				0.000450** (0.000220)	0.000536** (0.000170)
Extended Contract	0.00543 (0.00561)	0.0152** (0.00711)	0.0378* (0.0226)	0.0363* (0.0216)	
Observables	No	Yes	Yes	Yes	Yes
Driver FE	No	No	Yes	Yes	Yes
Arellano-Bond	No	No	No	No	Yes
N	12,576	12,576	12,576	12,576	5,337

Table 6: Evidence of Moral Hazard. The table reports results of the regression analysis of the relationship between the number of liability claims and individual’s risk class. The last column reports estimates based on Arellano-Bond procedure which uses lags as instruments. This explains the lower number of observation. We omit the coefficient on lagged dependent variable.

classes attract riskier drivers (adverse selection) and this induces positive association between the risk class and the number of accidents. On the other hand, higher risk classes provide stronger incentives to reduce risk since monetary consequences of an accident are much higher relative to the lower classes. This should induce negative association between the number of accidents and the risk class. The coefficient in front of the indicator of whether an individual purchases extended coverage similarly captures the combined effect of sorting (selection) and the response to weakening of incentives for risk reduction.

Table 6 shows that in the absence of controls for individual characteristics, the coefficient in front of the risk class variable is small but positive, and the coefficient in front of the contract choice indicator is not statistically different from zero. As we control for driver and car characteristics and thus, in part, for possible selection, the magnitude of the coefficient in front of the risk class variable decreases while the coefficient in front of the contract choice indicator becomes statistically significant and positive.¹⁷ Once we control for unobserved heterogeneity using individual-specific fixed effects, the coefficient in front of the risk class variable becomes negative, and the coefficient in front of the contract indicator further increases. The former finding can only be explained by the presence of moral hazard, since sorting across risk classes would work in the opposite direction. The later finding is also consistent with moral hazard. The first effect appears to be larger than the second effect.

¹⁷We use age, years of driving, and zip code dummies, as well as flexible second order complete polynomial of car characteristics.

It is possible that risk adjustment (moral hazard) arises for reasons other than monetary incentives embedded into the contract pricing. The results above indicate that individuals change their driving as they move across risk classes which may occur due to the natural progression of time if they move downward or because they cause an accident. We control for the former effect through ‘the years of driving’ variable which is included among the observable demographics. The latter may cause individuals to change the probability of accidents due to fear or because of a health or vehicle damage. To directly test for the presence of moral hazard caused by contractual incentives, we include an interaction between risk class and discretionary discount. The results are presented in column four of the Table 6. We observe that individuals with higher discounts adjust their driving less in response to the increase in risk class. This is consistent with moral hazard in response to contractual incentives, since discounts mute incentives embedded in experience-rating-based pricing.

As demonstrated by Nickell (1981), standard estimation methods (such as the “within” estimator used above) of dynamic models with fixed-effects may produce biased estimates. This is because by construction, the risk class is a function of lagged accidents, which in turn depend on the lagged values of the error term. If this is the case then the risk class may not be strictly exogenous. Instead it would be predetermined, i.e., correlated with the lagged error term.¹⁸ We thus re-estimate the relationship between the number of accidents and the risk class using a two-step version of Arellano and Bond (1991) estimator (also see Holtz-Eakin, Newey, and Rosen (1988)). Importantly, such an estimator is consistent even if the risk class is predetermined. The estimator uses lagged dependent and predetermined variables as instruments, and this decreases the sample size. Despite the smaller sample, we find a statistically significant coefficient in front of the risk class variable, which indicates negative risk adjustment associated with the increase in risk class. We additionally find statistically significant coefficient in front of the interaction term between the risk class and discount variables, which indicates response to contractual incentives. Also, the magnitudes of the estimated coefficients appear to be more in line with those we obtain in our estimation section.

Distribution of damages. Next, we describe the determinants of the distribution of liability claim sizes. Specifically, we would like to investigate if individuals engage in ex-post moral hazard, that is, if they strategically settle “at-fault” accidents without reporting them to the insurance

¹⁸Arellano and Bond (1991) formally define predetermined variables and demonstrate that their presence may likely result in the bias of standard estimators of fixed-effects models.

	(1)	(2)	(3)
Risk class	38.89 (78.71)	-39.33 (149.2)	26.40 (56.40)
Driver FE	No	Yes	Yes
Extended sample	No	No	Yes
N	12,576	12,576	250,000

Table 7: Testing for ex-post moral hazard. The table shows the results of the regression analysis of the relationship between the size of liability claim and individual risk class (conditional on submitting a claim). The first two columns use the sample of drivers that started and never left company A. Last column uses a full sample of drivers.

company in order to avoid upward risk class adjustment. Our contract structure contains a convenient exclusion restriction that enables us to test for ex-post moral hazard. In particular, the risk class adjustment does not depend on the size of the claims but only on their number. If ex-post moral hazard is present, we should see that the average size of submitted claims increases with the current risk class, because individuals should be willing to make larger side payments in higher risk classes.

Table 7 reports the results of a regression analysis of the relationship between the size of submitted claims and the risk class. This analysis reveals that, on average, the size of submitted claims is not related to the risk class. In order to control for selection bias, we conduct the same analysis comparing only the sizes of the claims submitted by the same driver in different risk classes. We achieve this comparison by including driver fixed effects and present the result in the second column of Table 7. We also repeat the analysis with a full sample of 250,000 drivers.¹⁹ We find no evidence that drivers engage in ex-post moral hazard.²⁰

5 Estimation Methodology

We now discuss the estimation methodology that we use to recover the primitives of the model described in Section 2. First, we summarize our identification strategy which is based on the properties of the model and features of our data summarized in Sections 3 and 4. Next, we specify the details of the econometric model. In particular, we explain how the individual-level parameters

¹⁹We also investigate the possibility of ex-post moral hazard when submitting comprehensive claims. We obtain similar results by regressing the value of the submitted comprehensive claim on the comprehensive risk class for customers in the comprehensive contract.

²⁰Our institutional environment is quite different from that of Abbring, Chiappori, and Zavadil (2008), who find evidence of ex-post moral hazard. These authors study the Dutch market in which withdrawal of liability claims is possible. Claim withdrawal gives the opportunity for a “not-at-fault” driver to submit the claim, and subsequently withdraw it, if the side payment is made. Thus, the Dutch market is more likely to be subject to ex-post moral hazard than the Portuguese market.

characterizing the cost of effort function and the within-period utility function vary over time; and how the distribution of these parameters in the population is linked to the characteristics of drivers and their cars observed in the data. After that, we summarized the econometric procedure used to recover the parameters of the model.

5.1 Identification

The object of the analysis in this paper is to recover individual-specific parameters (θ_i, γ_i) summarizing the cost of effort function and individual preferences for risk (or their distribution in the population) from the data on the history of accidents, risk classes, and contract choices. We begin by explaining how this can be achieved in the context of ideal data which are characterized by an infinite number of observations per individual or continuously varying incentives to which an individual might be exposed. We next turn to the realistic data where we might observe only a few observation per individual or a finite (small) number of settings with different incentives. Our point of departure is identification of models with fixed individual-specific risk (and homogeneous or individual-specific risk aversion). This helps us pinpoint the need for additional sources of variation to identify the model with endogenous risk and allows us to transparently illustrate how individual heterogeneity may be mis-measured if fixed risk models are used to interpret the data generated by the model with modifiable risk.

Ideal Data. A challenge specific to the auto insurance market is that idiosyncratic level of risk is not directly observed in the data. Instead, the researcher observes only discrete realizations of risk (accidents). When the risk is fixed and individuals are heterogeneous with respect to risk but share the same risk aversion, one of two sources of variation could be used. First, individual-level rate of risk is observed as the number of observation per individual goes to infinity. Then, the distribution of risk in the population can be recovered by nonparametrically aggregating individual-specific risk levels. The second approach exploits the existence of multiple contracts which differ in the degree of coverage. Individuals sort into these contracts on risk under appropriate pricing. In the limit (as the number of contracts goes to infinity) perfect sorting may be achieved. Then the chosen degree of coverage given the prices could be used to infer individual risk, and individual-level risks could be aggregated to recover the distribution of risk in the population. To identify a model where individual risk is still fixed but both risk and risk aversion may vary across individuals in population, in an ideal dataset would combine these two sources of variation. Observing an individual over many periods identifies idiosyncratic risk while observing an individual choosing

from infinitely many contracts, given the risk level, identifies the individual's risk aversion.

Consider now the setting where individuals are able to modify their risks, and they are potentially heterogeneous in the cost of effort and risk aversion. The task is to identify the distribution of these factors using the levels of risk individuals choose in response to various incentives. In the ideal data an individual would be repeatedly observed over many periods facing several different incentive schemes so that the risk level he chooses can be recovered for each of them. This could be achieved in multiple ways: (a) individuals may be observed visiting several different risk classes many times each; (b) exogenously facing several different premium levels (due to discounts or because of company experimentation), with many observations per price level; (c) or having many observations in contracts with different coverage. Any combination of these incentives could be used as well. Identification of the individual's parameters exploits the first order necessary conditions for the optimality of his risk choices. Specifically, these parameters can be recovered if the individual's chosen level of risk is known for three different incentive schemes. In other words, the relationship between these levels of risk and the parameters underlying the individual's decisions can be uniquely inverted in the same way as the relationship between the chosen degree of coverage and the idiosyncratic risk (or risk aversion) had been inverted in the context of the models with fixed risk.²¹

Realistic Data. In practice, a researcher may only observe an individual for a few periods, the number of contracts he may choose from is small, and the number of observations where an individual faces different incentives is limited. As a result, a researcher may only be able to identify a member of the finite-parameter family of distributions which best describes the data. Specifically, in the setting with fixed individual-specific risks and homogeneous risk aversion, a researcher would exploit the difference in the average number of accidents across subpopulations sorted into different contracts. In the setting with fixed risk where individuals are heterogeneous in risk and risk aversion, a researcher relies on two measures of sorting: (a) the difference in risk realizations across subpopulations associated with different contracts and (b) shares of different

²¹Notice that in the setting with ex-ante moral hazard the degree of coverage and the level of risk chosen by an individual depend both on his cost of effort and his risk aversion. Thus, it is impossible to design an identification strategy that would isolate and recover these factors one at a time. It is necessary to simultaneously disentangle their influence from the observable outcomes. This is in contrast to the settings with ex-post moral hazard where the observed action is taken after the risk is realized, e.g. an individual may decide how much health care to consume after getting sick or whether to submit a claim after an accident took place. The action, therefore, does not reflect an individual's risk aversion but only his cost of transforming the risk realization into an action. This allows separate identification of risk aversion from the identification of other risk-related parameters. Einav, Finkelstein, Ryan, Schrimpf, and Cullen (2013) propose an identification strategy based on this reasoning when investigating ex-post moral hazard in the market for health insurance.

contracts in the population.²²

For the model with modifiable risk, a researcher has to exploit the distributions of accidents for the same population but under different incentive schemes. We achieve this by following over a period of several years a cohort of drivers who obtain their licenses, and thus enter the market, in the same year. The mix of risk classes within such a cohort changes across the years which means that each year a cohort is observed to be subjected to a different mix of incentives. We supplement these moments with the moments associated with sub-populations which are randomly assigned to different discount levels when entering the market (and thus different price schedules associated with movement across risk classes).

Potential for Bias. The discussion above highlights the bias arising from applying the mechanism used in the context of a model with heterogeneous agents but fixed risk to the data generated by a model where risk is modifiable. The identification under fixed risk models relies on the differences in realized risk across individuals who choose contracts with different levels of coverage. To see how such an identification mechanism may lead to mis-measurement of the variation in individual-specific factors (causing the mis-measurement of the variation in private information), consider a case where an individual's risk aversion is negatively correlated with his cost of effort.²³

Suppose, first, that the selection into higher coverage is primarily driven by risk aversion considerations. In the model with modifiable risk, the extended coverage attracts highly risk averse individuals who have low costs of effort on average (due to the negative correlation in these factors assumed above) and thus choose relatively low levels of risk. They might appear to be similar to individuals who have higher costs of effort and have higher tolerance for risk but are subject to stronger incentives since they choose the liability-only contract. Then, a model which assumes fixed risk will underestimate the heterogeneity in risk-related factors and thus in the idiosyncratic risks that may realize in the environment with a different set of incentives. The heterogeneity in risk aversion would, however, be overestimated. Assuming ideal data, consider two individuals who are revealed to have the same level of risk but who choose contracts with different coverage. If risk is fixed, the difference in coverage is informative about the difference in their risk aversion.

²²One of the restrictions related to sorting is occasionally replaced in the literature by the parametric assumption on the process generating multiple accidents for a given individual within a period (e.g., Poisson distribution). This allows one to estimate the distribution of risks in the population from the shares of individuals having a different number of accidents in a given period.

²³The bias arises even if factors are not correlated. We use an example with the negative correlation because the effects are extreme and therefore very obvious in this case.

However, if risk is, in fact, modifiable, the individual who choose higher coverage (individual one) would maintain lower risk level if he is placed in the lower coverage contract chosen by the other individual (individual two). Individual one's choice of higher coverage (associated with risk level equal to the risk level of individual two) then reveals that he must be less risk averse than we would have estimated under the assumption of the fixed risk since his choice is not only motivated by the increase in coverage at a given risk but also by the reduction in the effort (and thus cost associated with effort) since he is able to increase his risk level when under higher coverage.

Second, consider the environment where the selection into higher coverage is primarily driven by cost of effort considerations. Then the extended coverage attracts individuals who have high costs of effort and tend to have low risk aversion (due to negative correlation assumed above). Such individuals will substantially relax their efforts under extended coverage which would make them appear very different from the lower cost and higher risk aversion individuals who choose the liability only contract and are therefore subject to stronger incentives. In this case and under fixed risk assumption, the variation in individual risk that may be realized under an alternative set of incentives may be substantially overestimated. The estimate of variability in risk aversion is also likely to be biased.

Similar arguments also establish that the model which ignores heterogeneity in either the cost of effort or in risk aversion may yield biased estimates of the importance of individual heterogeneity even if it takes into account the possibility of risk modification. Indeed, in the example above if a researcher ignores one of the sources of heterogeneity (s)he would tend to misinterpret the degree of selection as indicative of low or high heterogeneity correspondingly. Thus, to assess the extent and nature of the variation in the individual-specific factors and to recover the primitives required to guide policy analysis, it is essential to allow, alongside the possibility of moral hazard, for the heterogeneity in individual risk aversion and the cost of effort functions.

5.2 Parametrization

We now discuss our econometric model which is based on the economic model of insurance coverage and risk level choices outlined in Section 2. In this section, we specify how primitives of the model vary both within and across individuals in our setting. We use this specification to match patterns of risk and coverage choices observed in the data.

An individual in our setting is characterized by a triplet $(\gamma_i, \theta_{1,i}, \theta_{2,i})$. We assume that the parameter $\theta_{2,i}$ is fixed for each individual. The parameter $\theta_{1,i}$ may evolve over time within-individual

in a manner consistent with learning. In particular, we allow $\theta_{1,i}$ to take two values: $\theta_{1,i}^{high}$ and $\theta_{1,i}^{low}$. On obtaining a license, all individuals start with the high level of θ_1 and then stochastically transition to the low level over time; the probability of transition in any give period, p_{low} , is a parameter of the model; low level of θ_1 is an absorbing state. In the interest of tractability we assume the two levels of $\theta_{1,i}$ are proportional so that $\theta_{1,i}^{high} = \phi\theta_{1,i}^{low}$ where ϕ is a parameter of the model which is constant across individuals.

Next, let $z_{i,t}$ denote characteristics of an individual i that are observable in the data with \bar{z}_i representing the within-individual average of these characteristics. Then, we assume that $(\tilde{\gamma}_i, \theta_{1,i}^{low}, \theta_{2,i})$ are jointly distributed according to the truncated normal distribution (truncated at zero) such that

$$\begin{pmatrix} \gamma_i \\ \theta_{1,i}^{low} \\ \theta_{2,i} \end{pmatrix} \propto TN \left(\begin{pmatrix} \bar{z}_{1,i}\beta_{\gamma,1} + z_{2,i,t}\beta_{\gamma,2} \\ \bar{z}_{1,i}\beta_{\theta_1} \\ \bar{\theta}_2 \end{pmatrix}, \begin{pmatrix} \sigma_{\gamma,i}^2 & \sigma_{\theta_1,\gamma} & \sigma_{\theta_2,\gamma} \\ \sigma_{\theta_1,\gamma} & \sigma_{\theta_1}^2 & \sigma_{\theta_1,\theta_2} \\ \sigma_{\theta_2,\gamma} & \sigma_{\theta_1,\theta_2} & \sigma_{\theta_2}^2 \end{pmatrix}; \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \frac{1}{2\bar{w}} \\ \infty \\ \infty \end{pmatrix} \right),$$

We further assume that the distribution of γ_i is truncated from above in order to ensure the monotonicity of the utility function. We include in $\bar{z}_{1,i}$ the gender of individual, zip code dummy, and dummy corresponding to within individual average car value.²⁴ The parameter γ responsible for the risk attitude of the driver is allowed to change with age in a deterministic manner. Specifically, the mean of γ_i may be different for drivers under 25 years of age ($z_{2,i,t}$ is a dummy variable capturing this effect). Further, the variance of this parameter in population, $\sigma_{\gamma,i}^2$, may differ across car values (low vs. medium vs. high).

We thus estimate mean parameters $(\beta_{1,\gamma}, \beta_{2,\gamma}, \beta_{\theta_1}, \bar{\theta}_2)$, variance-covariance parameters $(\sigma_{\gamma}^2, \sigma_{\theta_1}^2, \sigma_{\theta_2}^2, \sigma_{\gamma,\theta_1}, \sigma_{\gamma,\theta_2}, \sigma_{\theta_1,\theta_2})$, and parameters characterizing learning process (ϕ, p_{low}) . We calibrate parameter \bar{C} to values implied by the studies of the value of life and set it to €500.²⁵ The utility parameter \bar{w} is normalized to €50,000. This normalization is without a loss of generality as explained in Section 2.

²⁴We very rarely see individuals change location in the data. When this happens we exclude such individual from our data set.

²⁵We obtain this number by multiplying the average value of life associated with car accidents, €500,000 by the probability that an accident results in a fatality or serious injury estimated in these studies to be around 0.001. We borrow these numbers from the studies of the value of life summarized in Viscusi and Aldy (2003).

5.3 GMM Estimation

We focus our attention on the subsample of individuals who started their participation in the car insurance market by signing a contract with the company A immediately after obtaining their driver's license and have continued their association with this company until and including part of the period covered in the data. There are two main reasons for using this sample. First, focusing on drivers that start their driving career with company A solves the initial condition problem since we do not need to account for the correlation between the risk-class and unobservable characteristics of entering drivers. Second, such a sample design maintains of the exogeneity of discretionary discounts since they do not depend on the risk class of the driver.

We implement a Simulated Method of Moments (see Pakes and Pollard, 1989)²⁶ within the context of the full solution nested fixed-point approach. We target the following moments.²⁷

1. The shares of drivers in various risk classes and with a given level of driving experience: $\mathbf{E}[\mathbf{1}\{K_{i,t}^L = K\}\mathbf{1}\{\text{years driving} \in \mathcal{E}\}]$, for $\mathcal{E} = [0, 2], (2, 5], (5, 10], [11, \dots)$.
2. The first and the second moment for the number of liability claims joint with the risk class and driving experience: $\mathbf{E}[R_{i,t}^L \mathbf{1}\{K_{i,t}^L = K\}\mathbf{1}\{\text{years driving} \in \mathcal{E}\}]$, $\mathbf{E}[(R_{i,t}^L)^2 \mathbf{1}\{K_{i,t}^L = K\}\mathbf{1}\{\text{years driving} \in \mathcal{E}\}]$, for all feasible K .
3. The share of drivers purchasing comprehensive coverage; the first and the second moment for the number of liability claims conditional on contract choice: $\mathbf{E}[\mathbf{1}\{Y_{i,t} = Y^C\}]$, $\mathbf{E}[R_{i,t}^L \mathbf{1}\{Y_{i,t} = Y^C\}]$, $\mathbf{E}[(R_{i,t}^L)^2 \mathbf{1}\{Y_{i,t} = Y^C\}]$
4. Within-person liability claim dynamics: $\mathbf{E}[(R_{i,t} - R_{i,t-1})\mathbf{1}\{K_{i,t-1}^L = 1\}]$,
5. The expected number of liability claims and the market share of the comprehensive contract conditional on the price discount: $\mathbf{E}[\mathbf{1}\{Y_{i,t} = Y^C\}\mathbf{1}\{D_{i,t} = D\}]$, $\mathbf{E}[R_{i,t}^L \mathbf{1}\{D_{i,t} = D\}]$, for $D = 2.5\%, 7.5\%, \dots$
6. The expected number of liability claims and the market share of the comprehensive contract conditional on the covariates – gender, zip code, car value (large, medium and small): $\mathbf{E}[\mathbf{1}\{Y_{i,t} = Y\}\mathbf{1}\{X_{i,t_i}^L = X\}]$, $\mathbf{E}[R_{i,t}^L \mathbf{1}\{X_{i,t_i}^L = X\}]$, $\mathbf{E}[(R_{i,t}^L)^2 \mathbf{1}\{X_{i,t_i}^L = X\}]$

²⁶Our choice of estimation technique is motivated by the necessity to resort to simulation moments. Since simulated maximum likelihood estimation calls for using a large number of simulation draws we choose the Simulated Methods of Moments in the interest of computational feasibility.

²⁷After removing degenerate and co-linear moments we end up with 275 moment conditions.

The moments are clustered at the level of an individual insured and GMM estimation is based on efficient weighting matrix.

Prior to implementing the Simulated Method of Moments procedure we estimate a number of auxiliary objects. Specifically, we distinguish between three types of accidents depending on the positive damage to third party only, own car only, or both. We impose that, conditional on the ‘at-fault’ accident, the type of the accident is exogenous, and is distributed according to a multinomial distribution with parameters p_1 , p_2 and $p_3 = 1 - p_1 - p_2$. These parameters are estimated from the data available for the individuals who have chosen comprehensive contracts. Following the literature, we similarly assume that the distribution of damages to own car $L_{t,r}$ depends only on observable car characteristics (specifically, it does not depend on the individual’s unobserved type). We estimate this distribution nonparametrically using a full sample of more than 30,000 claims for consumers that purchased the comprehensive coverage.

6 Results

In this section we summarize our estimation results.

6.1 Parameter Estimates

The estimated parameters are reported in Table 8. They are of reasonable magnitudes, have the expected signs and, are statistically significant. The estimates indicate regularities that have also been documented in other studies. For example, we find that women tend to be more risk averse and that the cost of effort varies across locations and is increasing in wealth (as proxied by the car value).²⁸

We estimate that the cost of effort for inexperienced drivers is substantially higher than the cost of effort for those who had been driving for a while. Our estimates indicate that an inexperienced driver has a 47% chance of becoming experienced in a year. This estimate appears reasonable since not all individuals have an opportunity to drive intensively and thus to learn fast. The estimated rate of learning implies that 95% of drivers become experienced within 5 years.

Table 9 reports some implied magnitudes associated with risk production and attitude towards risk. The estimates imply that an average accident rate equals to 0.082 with standard deviation

²⁸This is consistent with the perception in the literature that the cost of effort is in part the cost of inconvenience (in most cases the cost of inefficiently spent time) which tends to be increasing in individual income.

	Estimates	Std. Errors
Cost of Effort, Scaling Parameter (θ_1^{low}):		
Constant	1.08***	0.159
Medium car value	0.270	0.318
Large car value	0.57***	0.139
Zip code 1 and 2	-0.40***	0.130
Zip code 3	-0.32**	0.153
Female	0.15*	0.083
Learning:		
Cost multiplier (ϕ)	2.42***	0.16
Probability of learning (p_{low})	0.47***	0.018
Cost of Effort, Reciprocal Parameter (θ_2)		
Risk aversion, (γ):		
Constant	8.45***	0.006
Car value (linear term)	1.86***	0.003
Zip code 1 and 2	-0.87***	0.307
Zip code 3	0.26***	0.002
Female	0.36***	0.083
Under 25 multiplier (ψ)	-0.020	4.211
Higher Order Parameters:		
σ_{θ_1}	0.20***	0.072
σ_{θ_2}	1.98***	0.251
$\sigma_{\gamma,LM}$	1.23***	0.005
$\sigma_{\gamma,H}$	0.97***	0.190
ρ_{θ_1,θ_2}	0.000	0.000
$\rho_{\theta_1,\gamma}$	-0.55**	0.239
$\rho_{\theta_2,\gamma}$	0.65***	0.003

Table 8: Parameter Estimates. This table reports the estimated parameters of the model. Designation ‘Young driver’ applies to drivers who are less than 21 years old. The variance of the distribution of risk aversion parameter is estimated for low or medium car value ($\sigma_{\gamma,LM}^2$) and for high car value ($\sigma_{\gamma,H}^2$) separately.

0.034. Note that these numbers characterize all ‘at fault’ claims (i.e., both liability and comprehensive). These estimates indicate that idiosyncratic risk varies importantly across individuals in this environment. We will assess the role played by private information in generating such variability in risk in the Section 6.4.

Further, we compute a measure that, for each individual, reflects the marginal cost of changing the rate of accidents from the population average by one percentage point. The distribution of this measure is thus informative of the variability of the cost of effort in the population. Additionally, for each individual we compute a risk premium this individual would be willing to pay in excess of the expected loss to avoid the risk to his car associated with an average rate of accidents. This variable provides a monetized measure of risk aversion; and the distribution of this variable

	Unconditional Statistics			Unobserved Components	
	Mean	Standard Deviation	Coefficient of Variation	Standard Deviation	Coefficient of Variation
Accident rate	0.082	0.034	0.415	0.025	0.302
Marginal cost	33.042	15.989	0.484	7.679	0.232
Risk premium	73.950	122.541	1.657	83.024	1.123

Table 9: Some Implied Measures. This table reports implied values for the variables determining individual choice of risk level. ‘Marginal Cost’ reflects individual-specific cost of changing the rate of accidents from the population average (0.082) by one percentage point. ‘Risk Premium’ summarizes the amount an individual would be willing to pay in excess of the expected loss to avoid all risk to his car associated with an average rate of accidents.

informs us about the variability of risk aversion in the population. In the interest of brevity, we refer to the measures described above as ‘marginal cost’ and ‘risk premium.’

The results indicate that the cost of reducing the risk by one percentage point on average is equal to €33 with standard deviation of €16. It is comparable in magnitude with the €24 average risk exposure (average comprehensive claim \times reduction in risk = €2418 \times 0.01) associated with 0.01 chance of accident. Drivers in this population are willing to pay risk premium of €73 on average to avoid risk to their car associated with the average accident rate. The ‘back of the envelope’ calculation indicates that an expected risk exposure is about €2418 \times 0.08 = €192. An average individual is thus willing to pay 38% extra in order to avoid such risk. This reflects a substantial degree of risk aversion. Further, the risk premium has a standard deviation of €123 in the population indicating that an important fraction of drivers are very risk averse.

6.2 Model Fit

Our results indicate that the model fits the data quite well.²⁹ Table 10 compares several measures reflecting consumer contract and effort choices computed from the model to those computed from the data. As can be seen from the table, the model fits the contracts’ market shares (overall and conditional on covariates) within one percentage point. It is also capable of reproducing the average accident rate conditional on the contract, conditional on the car value (that is, within the groups with different risk aversion), and across risk classes.

We are somewhat less successful in fitting the variance of the number of accidents conditional on comprehensive contract. This is most likely explained by the fact that we do not observe many

²⁹Notice that the results that we present in the paper are based on the full set of moments which includes the moments associated with discount levels. However, the estimates that are obtained without the discount moments are very similar. In particular, all the values reported in Table 10 remain the same, including the values related to the market shares conditional on the various discount levels.

	Model	Data
Market Share of Comprehensive Contract		
overall	0.123	0.113
conditional on		
low car value	0.009	0.008
medium car value	0.061	0.052
high car value	0.413	0.396
Average Number of Liability Claims Conditional on		
liability contract	0.041	0.040
comprehensive contract	0.049	0.044
Std. Dev. Number of Liability Claims Conditional on		
liability contract	0.203	0.203
comprehensive contract	0.222	0.206

Table 10: Model Fit. This table summarizes fit of the model to the data. It reports the actual and predicted market shares of the comprehensive contract in the sample and conditional on covariate values as well as actual and predicted average number of liability accidents conditional on the type of the contract and other covariate values.

individuals selecting this contract especially for some values of covariates (such as low car value).




Figure 1: Distribution Over Risk Classes. This figure plots the empirical and simulated distribution of individuals across liability risk classes.

Years Driving	Avg. λ	Baseline		Homogenous marginal cost		Homogenous risk aversion	
		Variance	(Coef of Var) ²	Variance	(Coef of Var) ²	Variance	(Coef of Var) ²
1	0.118	0.0012	0.083	0.00008 (-92%)	0.006	0.00106 (-9%)	0.076
3	0.081	0.0009	0.129	0.00012 (-86%)	0.019	0.00072 (-15%)	0.109
5	0.070	0.0008	0.165	0.00015 (-80%)	0.030	0.00063 (-21%)	0.128
10	0.070	0.0008	0.163	0.00016 (-77%)	0.033	0.00059 (-24%)	0.122
20	0.070	0.0008	0.161	0.00017 (-77%)	0.034	0.00060 (-24%)	0.120
40	0.071	0.0008	0.161	0.00017 (-77%)	0.034	0.00060 (-24%)	0.120

Table 11: Risk Decomposition. This table reports the results of the simulation study which analyzes contributions of heterogeneity in marginal costs and in risk aversion to the variability of idiosyncratic risk.

The dynamic fit of our model is summarized in Figure 1 which plots the distribution of individuals across liability risk classes in the data together with the distribution of these individuals across classes simulated from the model. The two graphs are quite close which indicates that the model captures the dynamic evolution of individuals' histories in the population quite well.

6.3 Determinants of Risk: Cost versus Risk Aversion

Table 11 investigates relative importance of the cost of effort and risk aversion in risk production. The table reports statistics characterizing the distribution of risks in the population under three scenarios: (a) baseline model; (b) an environment where all individuals have cost parameters set to the mean of the unconditional distribution while preserving individuals' heterogeneity in risk aversion; (c) an environment where all individuals have their risk aversion coefficient set to the mean of the unconditional distribution while the heterogeneity in costs is preserved. We perform this analysis taking driving experience into account in order to capture the effects of learning and sorting which occur over time. In simulations we maintain independence of costs and utility parameters in order to achieve orthogonal decomposition.³⁰

The results reported in the table indicate that variability of risk is significantly reduced when variability of one of the factors is shut down. However, the risk appears to be more responsive to the heterogeneity in the cost parameters. Over time the importance of cost heterogeneity diminishes (as it is reduced through learning) and the importance of risk aversion increases.

³⁰The fractions of variance explained by costs and risk aversion do not always sum up to one in the table due to the simulation error.

	Mean	Std. Dev.	$\rho(\cdot, \theta_1)$	$\rho(\cdot, \theta_2)$	$\rho(\cdot, \gamma)$
Unconditional Distribution					
θ_1	1.11 (0.079)	0.40 (0.074)	1.00 (0.000)	-0.02 (0.049)	0.25 (0.098)
θ_2	6.59 (0.792)	1.93 (0.233)	-0.02 (0.049)	1.00 (0.000)	0.25 (0.052)
γ	8.54 (0.069)	1.08 (0.042)	0.25 (0.098)	0.25 (0.052)	1.00 (0.000)
Purged of Observable Variation					
θ_1	0.00 (-)	0.19 (0.051)	1.00 (-)	0.02 (0.074)	-0.43 (0.202)
θ_2	0.00 (-)	1.90 (0.233)	0.02 (0.074)	1.00 (-)	0.56 (0.012)
γ	0.00 (-)	0.92 (0.019)	-0.43 (0.202)	0.56 (0.012)	1.00 (-)

Table 12: Estimated Distribution of Individual-level Parameters. This table characterizes the estimated distribution of individual cost and utility parameters.

6.4 Importance of Asymmetric Information

Table 12 reports the estimated distribution of individual cost and utility parameters. The results indicate that all three parameters exhibit non-trivial variation both unconditionally and after the influence of observable factors have been purged away.

Specifically, the unobserved variation accounts for 23% and 73% of overall variation in θ_1 and γ respectively. The covariates used in estimation mirror information available to the insurance company. Hence, these findings indicate that significant informational asymmetries related to costs and risk aversion might be present in this market.

The estimates also show that cost and utility factors are correlated to an important degree. Interestingly, the utility coefficient reflecting individual risk aversion has an overall positive correlation with cost factors, but when the influence of observable factors is purged away the correlation with θ_1 is negative. Thus, overall individuals with low tolerance for risk tend to have high marginal costs and high sensitivity of marginal costs to adjustments in risk whereas conditional on observables low risk tolerance is associated with low marginal costs but still high sensitivity to risk adjustments.

A second look at Table 9 indicates that unobserved variation in costs and risk aversion translates into significant unobserved variation in idiosyncratic risk. Specifically, close to 50% of the variation in risk across individuals comes from the factors unobserved by the insurance company. Thus, in contrast to findings of the earlier literature, we find substantial variation in unobserved risk (the coefficient of variation is 0.3).

6.5 Impact of Risk Provision Incentives

In this section we assess the magnitudes of risk adjustments prompted by dynamic considerations associated with contract pricing across risk classes and static considerations related to the variation in the degree of coverage.

We begin with incentives embedded into experience rating. To abstract away from additional complications related to sorting into contracts we focus on the drivers with low car value who predominantly (with probability 99%) choose a liability only contract. At the end of our discussion we comment on the differences between the low and medium or high car value drivers. The dashed




Figure 2: Risk Provision Across Classes. This figure shows the average of accident rates chosen by agents with low car value under exogenous and endogenous allocation into risk classes. The endogenous allocation graph reflects sorting into risk classes in the long run (after 40 years). We focus on low car value to isolate behavior under liability contract since agents with low car value very rarely choose comprehensive coverage.

line in Figure 2 displays the levels of risk optimally chosen by individuals with low car values when they are associated with the risk class at random (exogenously) rather than on the basis of their driving history. In this exercise we assume that after being randomly allocated into the risk class an individual correctly takes into account the probabilities of transition across risk classes upon having an accident and the monetary implications associated with such transitions. We disregard learning and evaluate an individual's behavior on the basis of his low (long-run) marginal cost of effort. The figure shows that the levels of risk chosen by individuals decline significantly with the risk class. This reduction in risk is extreme in classes above class 10. In this region the incentives are so strong they reduce the expected level of risk to 0.02. The chosen levels of risk increase after risk class 15. The price incentives disappear at this point since an individual is always guaranteed the placement in class 16.

Risk Class	Allocation					
	Exogenous			Endogenous		
	Mean Accident Rate	Relative to Class One	Std. Dev. Accident Rate	Mean Accident Rate	Relative to Class One	Std. Dev. Accident Rate
1	0.076		0.025	0.063		0.039
2	0.075	-0.013	0.026	0.083	0.317	0.023
3	0.074	-0.026	0.026	0.074	0.175	0.03
4	0.074	-0.026	0.026	0.068	0.079	0.033
5	0.073	-0.039	0.026	0.073	0.159	0.03
6	0.071	-0.066	0.027	0.069	0.095	0.032
7	0.07	-0.079	0.027	0.063	0.000	0.033
8	0.069	-0.092	0.027	0.068	0.079	0.031
9	0.067	-0.118	0.028	0.066	0.048	0.031
10	0.066	-0.132	0.028	0.064	0.016	0.03
11	0.064	-0.158	0.028	0.067	0.063	0.029
12	0.06	-0.211	0.028	0.065	0.032	0.028
13	0.054	-0.289	0.028	0.061	-0.032	0.027
14	0.046	-0.395	0.027	0.057	-0.095	0.026
15	0.037	-0.513	0.025	0.051	-0.190	0.024
16	0.043	-0.434	0.027	0.061	-0.032	0.022
17	0.053	-0.303	0.028	0.072	0.143	0.019
18	0.066	-0.132	0.028	0.082	0.302	0.015

Table 13: Risk Provision Across Classes. This table quantifies individual responses in terms of chosen accident rates to the incentives embedded in experience rating. Numbers reflect decisions of individuals with low car values. The endogenous allocation graph reflects sorting into risk classes in the long run (after 40 years).

Table 13 quantifies these responses. It indicates that on average individuals are quite responsive to the incentives. Specifically, an average individual reduces his risk by 4% when he is moved from class one to class five, and by 13% when he is moved from class one to class ten. As we indicated earlier the incentives really kick in after class ten so that the accident rate chosen in class one is reduced almost by 30% once the individual is placed into the class 13. The results in the table also indicate that individuals are quite heterogeneous in their responses. The standard deviation of the risk rates range between 0.026 to 0.028 which constitutes 30% to 50% of the average risk.

The choices of agents with medium or high car values resemble the choices made by agents with low car values. The main difference is that individuals with higher car values have higher risk exposure under liability contracts (since risk is proportional to car value). They also tend to be more risk averse. As a result such individuals generally choose lower levels of risk on average. Additionally, their chosen risk levels vary across risk classes to a lesser degree (from 0.049 in class one to 0.03% in class ten).

Having documented the magnitudes of risk responses that arise under exogenous allocations we turn to the analysis of the risk levels that arise under real-life operation of this system which in addition to risk provision also facilitates sorting of individuals across risk classes. The impact and importance of sorting could be already seen in Figure ?? by comparing the levels of risk chosen under exogenous (dashed line) and endogenous (solid line) allocation. The latter is based on the long-run (hypothetical) positions of individuals who progressed through the risk classes according to the rules of experience rating. As can be easily seen in the figure, the average risk profile under endogenous allocation is much flatter and is actually upward sloping for some of the lower risk classes. Only in the classes above ten the risk profile is downward sloping. It closely tracks average risk under exogenous allocation while remaining always above by two percentage points on average.

The discrepancy in the levels of risk under exogenous and endogenous allocation is indicative of sorting. Indeed, the individuals for whom low levels of risk are economically justified progress downward across risk classes and are more likely to find themselves in a lower rather than higher risk class. Individuals who remain in higher classes either have higher costs of adjustment or higher tolerance for risk and are thus endogenously high risk drivers. Specifically, their chosen levels of risk substantially exceed the average levels of risk that would have been chosen in those classes by the non-selected population. Interestingly, the worst drivers would be stuck in classes above class ten where experience-based pricing punishes them with extreme incentives so that in equilibrium their risk is substantially below the risk of more flexible individuals. We will study the importance of sorting in more detail in the next section.

Allocation	Liability Contract	Comprehensive Contract
Exogenous	0.037	0.089
Endogenous	0.045	0.067

Table 14: Risk Provision Across Contracts. This table summarizes the risk levels chosen by agents with high car values under different levels of insurance coverage.

Next, we take a look at the risk choices across different levels of coverage. Table 14 summarizes risk levels for the population of drivers with the high car values (who choose comprehensive coverage with probability 0.42). The table indicates that the difference in coverage may induce a strong response in the choice of risk levels. Indeed, comprehensive coverage results in more than twice as high levels of risk relative to the liability only coverage under exogenous allocation.

However, under endogenous allocation the difference is less striking since individuals who select into comprehensive coverage tend to have lower tolerance for risk and thus are likely to avoid increasing accident rates.

We investigate the magnitudes of sorting induced by experience rating next.

6.6 Sorting

The experience rating system facilitates sorting of individuals into risk classes on the basis of their history of accidents. We first investigate the sorting that could be achieved in the long run and then turn our focus to a short-run sorting.

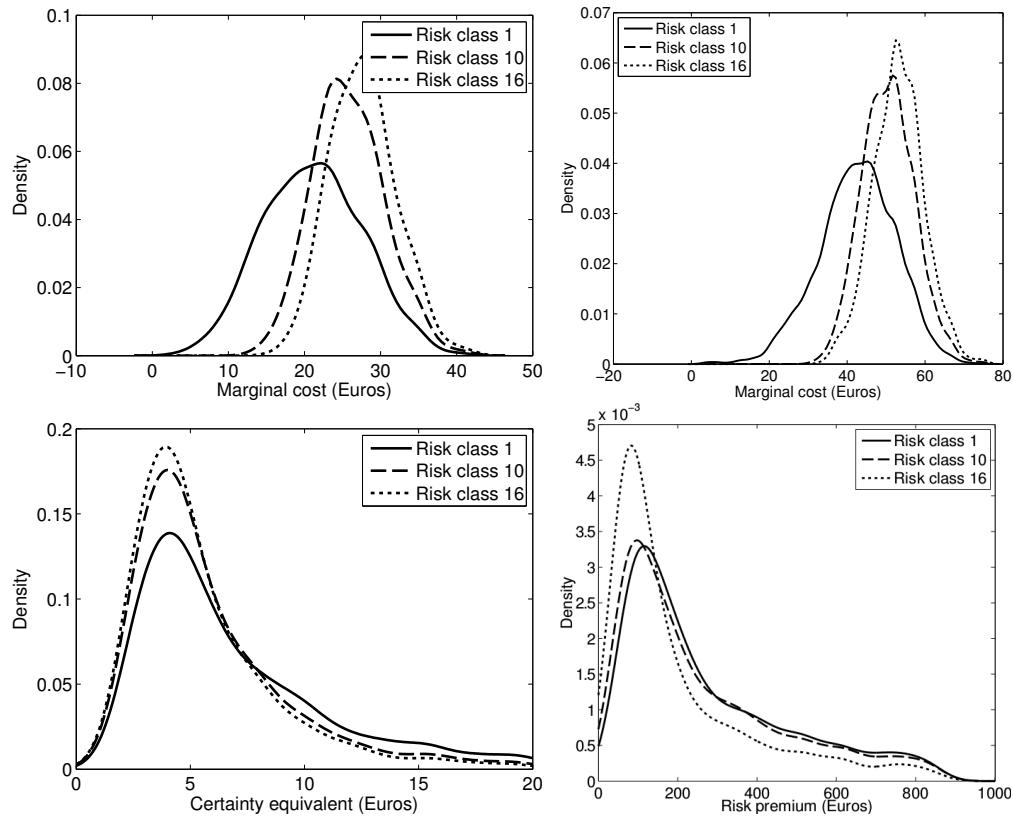


Figure 3: Sorting across Risk Classes. This figure demonstrates sorting into risk classes on the basis of marginal cost of effort and risk premium. It plots the distributions of these factors for the individuals who endogenously reached classes 1, 10 or 16 in the long run (the left panel shows the distributions for individuals with low car values and the right panel shows the distribution for individuals with high car values).

In Figure 3 we plot the distribution of marginal costs and risk premiums defined as above for the individuals who endogenously end up in classes one, ten and sixteen in the long run. We show these graphs for the individuals with low and high car values separately. Let us first consider individuals with low car value. The figure indicates that sorting is indeed present. That is, in the long run the individuals who end up in class one are characterized by lower marginal costs of




Figure 4: Sorting across Contracts. This figure demonstrates sorting across contracts on the basis of marginal costs of effort and risk premium. It shows the distribution of these factors among individuals who choose liability and comprehensive contracts respectively.

adjusting risk and higher risk aversion relative to the individuals who end up in class 10 or 16. In the case of marginal costs this difference is mostly manifested by the difference in means, whereas in the case of risk premium the distributions differ by the mass they allocate to the upper tail (that is, by the probability of an individual to be very risk averse). Nevertheless, the supports of these distributions remain largely overlapping across classes. Since classes one and 16 are drastically different in their past risk realizations it appears that experience rating is not very effective in sorting individuals on their risk attributes. This is mostly because accidents are quite rare and some individuals that are endogenously high risk manage to reach lower classes undetected.

Further confirmation of this regularity can be found in Table 15. To generate this table we have chosen an individual who is average on the basis of his observables (he is a man of 30 years old who lives in zip-code of type 1 and owns a low value car). For this individual we generate an artificial sample consisting of individuals who are his identical replicas in terms of observables but have different realizations of unobserved factors. We then document the risk class assignment of each of these individuals over the course of 40 years after obtaining a driver’s license. Table 15 reports an explanatory power (R^2) of the regression relating an individual’s marginal cost, his risk premium, and his “inherent riskiness” to the liability risk class he finds himself in after a certain number of years and to the indicator of whether he chooses comprehensive coverage. We measure inherent riskiness as the level of risk he would choose in class one under exogenous allocation. This measure can be seen as a one-dimensional index of individual primitives that reflects his potential riskiness but abstracts from endogeneity of incentives. As can be seen from Table 15 the

Driving Years	Included Variables	Marginal Cost	Risk Premium	Risk in Class 1
Low Car value				
1 year	Risk class (RK)	0.006	0.001	0.005
3 years	Risk class (RK)	0.015	0.002	0.016
5 years	Risk class (RK)	0.024	0.003	0.026
10 years	Risk class (RK)	0.048	0.006	0.057
40 years	Risk class (RK)	0.017	0.003	0.022
High Car value				
1 year	Risk class (RK)	0.002	0.002	0.001
	RK + Compr. contract	0.002	0.002	0.001
3 years	Risk class (RK)	0.004	0.005	0.004
	RK + Compr. contract	0.005	0.006	0.004
5 years	Risk class (RK)	0.007	0.008	0.008
	RK + Compr. contract	0.027	0.079	0.016
10 years	Risk class (RK)	0.025	0.010	0.029
	RK + Compr. contract	0.118	0.296	0.203
40 years	Risk class (RK)	0.016	0.000	0.026
	RK + Compr. contract	0.112	0.309	0.213

Table 15: Measuring Sorting. This table reports R^2 of the regression relating an individual's marginal cost of effort, the risk premium and the level of risk an individual would choose if he were exogenously placed in class one to the liability risk class he sorts into after a certain number of years and to the indicator of whether he chooses comprehensive coverage.

placement in a given risk class provides very little information about an individual's unobserved cost, risk aversion or inherent riskiness.

Let us now consider individuals with high car values. A sizable fraction of these individuals enroll in the contracts with comprehensive coverage once they reach lower risk classes (the comprehensive contract is too expensive for most individuals when they are in high risk classes). The right panel of Figure 3 indicates that sorting on the marginal cost is somewhat less pronounced than in the case of the low car values. The distribution of marginal costs appears to be almost identical for classes 10 and 16. The sorting on risk aversion across risk classes appears to be somewhat more important. Figure 4 documents sorting across contracts for individuals with high car values which appears to be more significant. Turning again to Table 15 we can see that in the long run the contract choice has significantly higher predictive power than the risk class allocation. This regularity holds for all three measures. Thus, differential coverage is more effective in sorting individuals on their unobserved factors.

To understand these results, consider the informational content of R^2 measure. Intuitively, the

variance of the individual factor, e.g. the coefficient of risk aversion, (and therefore the variance of the error term in the regression summarized in Table 15) within the group associated with a given contract (or risk class) is smaller than the variance of the coefficient of risk aversion in the population. The R^2 is determined by the largest variance of the coefficient of risk aversion across such groups since it reflects the remaining error in the regression analysis (specifically, higher values of such variance lead to lower values of R^2). Since selection into contracts is associated with the partitioning of the support it results in a substantial reduction of variance. In contrast, experience rating does not lead to support partitioning. Further, most of individuals in a given cohort are always located in the lowest class. Thus, the variance of the regression error is as large as the variance of the coefficient of risk aversion in the lowest class which is very close to the variance of the coefficient of risk aversion in the population.

We also investigate sorting in the short run: after individuals have been driving for one, three, five or ten years. Table 15 illustrates the information acquisition within this time frame. As can be seen from the table more information is revealed with time. However, even after five or ten years very little sorting across classes has occurred. Not surprising the sorting on risk classes provides no information in the early years both for the drivers with low and high car value. After 5-10 years of driving the risk class does contain some information about the type but this information is extremely limited. As before, selection into a comprehensive contract (which becomes economically viable after 5 years) is much more informative about the agent's type than his placement into the risk class.

7 Illustrative Exploration of Contract Design

Our findings about moral hazard and private variation in the cost of effort- and risk aversion-related factors have implications for contract design. Specifically, when deciding on contract terms an insurance company may emphasize dynamic pricing based on experience rating or offer a more extensive menu of static contracts. The first strategy will incentivize risk reduction whereas the second strategy will be effective in screening customers on risk and thus will allow to price risk more precisely. It is worth noting that static contracts may also provide incentives for risk reduction when they offer reduced coverage and thus expose customers to more risk on the margin.

Historically, European companies have been legally prevented from incentivizing risk reduction through static contracts due to the prevailing law on minimal coverage. As a result European

insurance practitioners might place larger emphasis on the dynamic risk provision incentives delivered through experience-based pricing. In this section we investigate potential welfare costs of this law. We consider a counterfactual world where a partial liability contract is offered in addition to the full liability and comprehensive coverage present in the data. We implement such partial liability in the form of a deductible that the individual has to pay out-of-pocket with the size of the deductible tied to the individual's car value. We leave the price of the full liability contract unchanged in order to hold the enrollment decisions fixed. (One can imagine the price of the full liability is quoted when the individual first contacts the company whereas details of other available coverages are discussed if the individual expresses interest in enrolling.³¹) Specifically, the price of a new contract is set equal to the price of a full liability contract minus an additive discount (held constant across risk classes).

7.1 Drivers' Choices under Counterfactual Menu of Contracts

We begin by investigating decisions of individual drivers who are faced with such an alternative menu of contracts. This analysis focuses on the subset of drivers with low car values since the choice between the full and partial liability coverage is most relevant for this group. The drivers with high car values tend to choose either full liability or comprehensive coverage.

Figure 5 illustrates selection and chosen levels of risk. We construct these graphs by exogenously associating individuals with risk classes as in Section 6.5 to abstract away from endogenous sorting associated with experience rating. We maintain, however, that agents expect their risk classes to be adjusted according to the rules of experience rating in the future. They may also reconsider their choice of contract each period. Hence, we have a non-selected population associated with each risk class and all individuals in a given class face the same consequences of having 'at fault' accident regardless of the contract they choose.

The lower panel of figure 5 depicts the average risk premium and marginal costs for the individuals who choose the full liability coverage and the liability coverage with the deductible set at 20% of individual's car value. The graphs thus reflect sorting into the contracts on risk aversion and marginal cost of effort. The contract with partial liability coverage attracts individuals with lower cost of effort uniformly across risk classes. The graph on the right in the top panel explains this regularity. Upon choosing the contract with partial liability coverage these individuals reduce

³¹This idea was first proposed in Cohen and Einav (2007).




Figure 5: Implications of Partial Liability Coverage. These figures illustrate sorting into contracts and the levels of risk chosen by drivers when faced with a menu of contracts which includes a partial liability contract. The top panel shows the share of low car value drivers who chose partial liability coverage (left) and the levels of risk chosen by different subsets of drivers (right). The lower panel depicts the average risk premium (left) and the average marginal costs for the groups of drivers who choose different coverage (right).

their risk below the position they would have taken had they chosen the full liability coverage. Indeed these individuals choose lower coverage because by adjusting their risk they are able to reduce their risk exposure to the level where the contract with the partial liability coverage (and lower premium) becomes more attractive than the full liability coverage. The marginal cost graph indicates that the attractiveness of the contract with liability increases with risk class (up until class 15). Recall that the premium schedule for the contract with the deductible is equal to the premium schedule of the contract with full liability coverage net of additive discount. The incentives for the risk provision increase with the risk class which makes it easier for an individual (holding his risk aversion fixed) to achieve the level of risk which makes partial coverage attractive. In other words the contract with partial coverage attracts higher marginal cost individuals on average in higher classes.

Further, the left-hand side graph in the lower panel indicates that the contract with partial

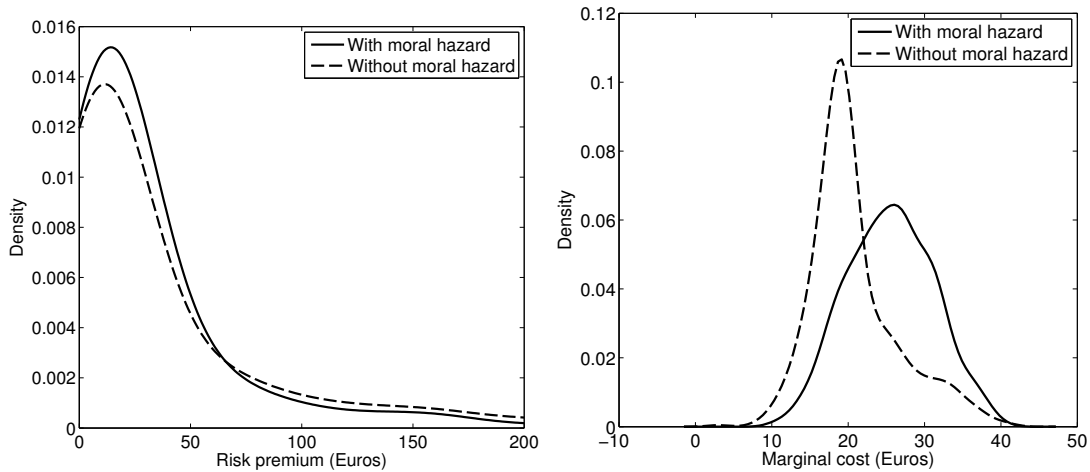


Figure 6: Comparing Adjustable and Fixed Risks. These figures illustrate sorting into partial liability contract under the baseline model (with adjustable risk) and the case when the risk of each driver is fixed at the level he would have chosen in the risk class one under exogenous allocation.

coverage attracts individuals with higher tolerance for risk. This is not surprising since these individuals require lower compensation for taking on extra risk. Interestingly, this effect works against the one described above since these individuals are also less prone to adjust their risk in response to the incentives. That is why, switching is associated only with moderate adjustment of risk levels.

Figure 6 illustrates the impact of risk adjustment on sorting into the contract with partial liability coverage. Specifically, it graphs the distribution of the risk premiums and marginal costs for two cases: (a) when agents are allowed to adjust risk as in our baseline model; (b) when an individual's risk is fixed at the level he would have chosen under a standard menu of contracts (which includes only full liability and comprehensive coverage). The graphs illustrate that selection is more pronounced in the later case since only individuals with sufficiently low (fixed) levels of risk would choose partial coverage. In contrast, with risk adjustment a larger set of individuals prefer partial coverage. It includes not only individuals who choose sufficiently low risk under full liability coverage but also those who can reduce their risk to the level which makes partial coverage attractive.

7.2 Welfare Analysis

In this section we investigate the welfare implications of offering an additional contract with partial liability coverage. Before we describe our findings a comment is in order. Mandatory liability insurance is motivated by the need to ensure that victims of the auto accidents get compensated

for their losses.³² However, there are other ways in which this goal can be achieved. Our objective in this analysis is not to evaluate whether the mandatory insurance should be abolished. Instead, we would like to point out that the possibility of moral hazard adds another dimension to the assessment of the costs associated with this legislature.

We consider several levels of deductible and copay associated with different price discounts. To maintain the constancy of additional risk exposure we assume that deductibles and copays are additive rather than multiplicative. The difference between copay and deductible is that within a given contract period an individual pays out-of-pocket until the deductible is met, and after that the insurance company compensates for the relevant driving-related losses; in contrast, a copay specifies the portion of each claim that an individual should cover out of pocket. We measure the change in consumer welfare in terms of compensating variation. Specifically, we are looking for the amount an individual would have to be paid each period in order to remain indifferent between the settings with and without the contract with partial coverage. Similarly, the change in value appropriated by the insurance company is computed as an amount the company would be willing to pay per period to stay in the environment with the additional contract rather than in the baseline world where only the full liability contract is offered. In this analysis we ignore possible changes in administrative costs since they are likely to be of a second order importance relative to the costs associated with risk related expenses. The results of our analysis are summarized in Table 16.

In all cases we consider, offering a contract with partial liability coverage reduces value to the company. Specifically, the company gives up more in price reduction than they are able to save due to the reduced coverage and because individuals reduce their risk upon switching. This appears to happen because the risk is already quite low in the system and the premiums are quite high.

In contrast, in all cases we consider the consumers gain from the introduction of a contract with partial coverage. As we explained above the welfare gains are mostly associated with the reduction in price which compensates for the cost of additional effort incurred in order to minimize risk exposure and for this additional exposure. We show that gains in consumer welfare sometimes exceed the loss of value to the company leading to higher overall welfare (e.g., with a deductible equal to 20% of individual's car value (€200 for an average insuree) and price reduction of €8;

³²Another motivation for enforceable mandatory coverage laws is that they help overcome the adverse selection associated with purchasing car insurance. This problem appears to be quite limited in Portugal (see footnote 32).

Additive Discount	Market share	Accidents per 800,000	Δ CS 1,000€	Δ PS 1,000€	Δ TS 1,000€	Δ TS (€) per accident	
Contracts with Deductible							
20%	€6	14%	-481	254.8	-231.0	23.8	49.4
20%	€8	22%	-901	507.7	-459.4	48.4	53.7
20%	€10	46%	-1,518	971.7	-928.2	43.5	28.6
20%	€12	85%	-1,853	2,207.7	-2,411.5	-203.8	-109.9
30%	€6	8%	-308	188.0	-182.6	5.4	17.4
30%	€12	25%	-1,464	843.7	-797.2	46.4	31.7
30%	€14	36%	-2,102	1,299.8	-1,246.1	53.8	25.6
30%	€16	76%	-2,551	2,353.7	-2,587.2	-233.5	-91.5
Contract with Copay							
20%	€6	12%	-442	249.8	-229.2	20.6	46.5
20%	€8	20%	-859	495.8	-468.0	27.7	32.3
20%	€10	45%	-1,435	1,002.5	-1,062.6	-60.1	-41.8

Table 16: Welfare counterfactuals. In this table ‘deductible’ refers to the amount an individual has to cover out of pocket before he is compensated by the insurance company; ‘copay’ refers to the amount an individual has to pay out-of-pocket per claim. ‘Life-time value’ reflects the loss of value to the company from offering the contract with deductible/ copay in perpetuity for the current set of customers.

or alternatively for the contract with 20% copay and €8 reduction in contract price). More importantly though, an introduction of the contract with partial coverage leads to substantial reduction in the number of accidents. For example, the annual number of accident perpetrated by individuals associated with this insurance company is reduced by 1,518 for the deductible of 20% and the price discount of €10 . Even in the cases when total welfare is reduced by the introduction of the partial liability contract, the lost welfare per eliminated accident is quite small. In some cases it is close to €30 -€40. Thus we would expect that if the social gains associated with the reduced number of accidents were taken into account the overall welfare would be increased.³³

8 Conclusion

The reason why measuring private information about idiosyncratic risks has been the focus of a large empirical literature is clear. Such measures are essential for assessing the efficiency of insurance market and for interpreting the objectives pursued by the insurance companies through contract design, such as sorting market participants on the amount of private risk. Primarily due

³³Our analysis also does not take into account illegal (uninsured) drivers who will default under the current system but may consider enrolling if partial insurance were available. Note that the prevalence of uninsured drivers is small in Portugal relative to other European countries [http : //ec.europa.eu/internal_market/insurance/docs/2006-market-docs/2508-06-annex-uninsured-cars-en.pdf](http://ec.europa.eu/internal_market/insurance/docs/2006-market-docs/2508-06-annex-uninsured-cars-en.pdf)

to data availability, the early seminal empirical results in this literature were established in the context of a car insurance market but were thought to offer insights that are more generally valid. However, in the market for auto insurance, and likely many other markets, individuals might be able to modify their risk in response to incentives. Clearly, the probability of accidents depends on, e.g traffic conditions in which an individual drives and insurance companies are neither fully informed on an individual's flexibility in choosing how much and when to drive nor are able to condition the contract on these variables. It is thus not surprising that the econometric tests have indeed detected presence of significant moral hazard in this market. We argue that it is essential to take moral hazard into account both when measuring the extent of private information and when interpreting the contract design which may serve a dual purpose of sorting on risk and providing incentives for modifying risks. Moreover, even the content of relevant private information might shift from information about risks to information about the ability to modify risks with associated implications for the market performance and the potential regulatory policy response.

Thus, the first substantive contribution of the paper is to generalize the existing framework to simultaneously allow for individuals' heterogeneity in risk-related factors and their ex ante ability to adjust risk in response to incentives. We show how additional sources of identifying variation can be used to identify model primitives and explain the adverse consequences of following the traditional identification strategy that neglects moral hazard. As a result, we are able to obtain a more reliable measure of private variation in risk and risk-related factors.

As could be expected, our estimates indeed reveal that an average agent is able to importantly modify his risk. For example, the average individual chooses 18% less risk when placed in risk class 10 rather than class 1. Even more dramatically, the average individual would choose twice the amount of risk when placed in the comprehensive coverage rather than liability coverage. At the same time, the market is characterized by a significant variation in the degree of feasible risk adjustments due to individuals' heterogeneity in the cost of effort and risk aversion. Indeed, the standard deviation in risk levels chosen by individuals with the same observable characteristics within a given risk class is 30% of the average risk level for that class. Importantly, the model closely predicts the degrees of sorting and of risk adjustment associated with a large set of qualitatively different incentives (starting with differing degrees of coverage to movement across risk classes and different levels of price discounts). The fact that the model is capable of rationalizing the data along multiple dimensions lends credibility to the findings.

The second contribution of the paper is to use the estimated model to gain insight into possible

motivation underlying the structure of contracts offered by this industry. Our analysis suggests that the insurance industry appears to place an important emphasis on incentivizing effort provision rather than on sorting individuals on the basis of their risk. Indeed, experience rating provides strong incentives for risk reduction especially for young drivers. At the same time it induces only a limited sorting. The pricing of contracts with different coverage is also not targeted to sort drivers on private risk. This is evident from the fact that the purchase of extended coverage is largely determined by the observable price of the car. The observed pricing thus ensures that only sufficiently risk averse (and thus relatively safe) drivers are offered a choice among multiple contracts. In addition, observed pricing excludes young drivers from purchasing comprehensive coverage, because they are less flexible (high cost of effort) and therefore risky drivers.

The third contribution of the paper is to illustrate the quantitative importance of jointly considering the heterogeneity in risk-related factors and individual ability to modify risk for policy design. To this end, we assess the welfare cost of mandatory liability insurance. We find that the ability of insurance companies to screen on risk and to further incentivize risk reduction is limited by existing laws which mandate full liability coverage. Interestingly, the primary cost of such restrictions comes not from the reduced efficiency of the market due to incomplete sorting but from raising the number of accidents that could have been prevented in the absence of such laws. In light of these findings, we believe that the effects of moral hazard have to form an integral part in future substantive analyses of many insurance markets. Moreover, the recognition that the amount of risk might be endogenous to policy must be taken into account when selecting among regulatory designs that may perhaps achieve the same policy objectives.

References

- ABBRING, J., P.-A. CHIAPPORI, J. HECKMAN, AND J. PIQUET (2003): “Adverse Selection and Moral Hazard in Insurance: Can Dynamic Data Help to Distinguish?,” *Journal of the European Economic Association*, 1(2-3), 512–521.
- ABBRING, J., P.-A. CHIAPPORI, AND J. PIQUET (2003): “Moral Hazard and Dynamic Insurance Data,” *Journal of the European Economic Association*, 1(4), 767–820.
- ABBRING, J., P.-A. CHIAPPORI, AND T. ZAVADIL (2008): “Better Safe than Sorry? Ex-ante and Ex-Post Moral Hazard in Dynamic Insurance Data,” Working paper, Columbia University.

- ARELLANO, M., AND S. BOND (1991): “Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations,” *Review of Economic Studies*, 58(2), 277–97.
- BARROS, P. P. (1996): “Competition Effects of Price Liberalization in Insurance,” *The Journal of Industrial Economics*, 44(3), 267–287.
- CHIAPPORI, P.-A., AND B. SALANIE (2000): “Testing for Asymmetric Information on Insurance Markets,” *Journal of Political Economy*, 108(1), 56–78.
- COHEN, A. (2005): “Asymmetric Information and Learning: Evidence from the Automobile Insurance Market,” *Review of Economics and Statistics*, 87(2), 197–2007.
- COHEN, A., AND L. EINAV (2007): “Estimating Risk Preferences from Deductible Choice,” *American Economic Review*, 97(3), 745–788.
- DIONNE, G., P.-C. MICHAUD, AND M. DAHCHOUR (2010): “Separating Moral Hazard from Adverse Selection and Learning in Automobile Insurance: Longitudinal Evidence from France,” *Available at SSRN 1673381*.
- (2013): “Separating Moral Hazard from Adverse Selection and Learning in Automobile Insurance: Longitudinal Evidence from France,” *Journal of the European Economic Association*, 11(4), 897–917.
- DIONNE, G., J. PINQUET, M. MAURICE, AND C. VANASSE (2011): “Incentive Mechanisms for Safe Driving: a Comparative Analysis with Dynamic Data,” *The Review of Economics and Statistics*, 93(1), 218–227.
- EINAV, L., A. FINKELSTEIN, AND J. LEVIN (2010): “Beyond Testing: Empirical Models of Insurance Markets,” *Annu. Rev. Econ.*, 2, 311–336.
- EINAV, L., A. FINKELSTEIN, S. RYAN, P. SCHRIMPF, AND M. CULLEN (2013): “Selection on Moral Hazard in Health Insurance,” *American Economic Review*, 1003(1), 1–44.
- HOLTZ-EAKIN, D., W. NEWEY, AND H. S. ROSEN (1988): “Estimating Vector Autoregressions with Panel Data,” *Econometrica*, 56(6), 1371–1395.

- ISRAEL, M. (2004): “Do We Drive More Safely When Accidents are More Expensive? Identifying Moral Hazard from Experience Rating Schemes.” Working Paper 43, University Center for the Study of Industrial Organization.
- MIMRA, W., AND A. WAMBACH (2014): “New Developments in the Theory of Adverse Selection in Competitive Insurance,” *The Geneva Risk and Insurance Review*, 39, 136–152.
- NICKELL, S. (1981): “Biases in Dynamic Models with Fixed Effects,” *Econometrica: Journal of the Econometric Society*, pp. 1417–1426.
- PAKES, A., AND D. POLLARD (1989): “Simulation and the Asymptotics of Optimization Estimators,” *Econometrica: Journal of the Econometric Society*, pp. 1027–1057.
- PUELTZ, R., AND A. SNOW (1994): “Evidence on Adverse Selection: Equilibrium Signalling and Cross-Subsidization in Insurance Market,” *Journal of Political Economy*, 102(2), 236–257.
- ROTHSCHILD, M., AND J. STIGLITZ (1976): “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *Quarterly Journal of Economics*, 90(3), 629–649.
- VISCUSI, W. K., AND J. ALDY (2003): “The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World,” *The Journal of Risk and Uncertainty*, 27(1), 5–76.
- WILSON, C. (1977): “A Model of Insurance Markets with Incomplete Information,” *Journal of Economic Theory*, 16, 167–207.