

Dynamic Selection and Reclassification Risk: Theory and Empirics*

Igal Hendel[†]

ESWC, Montreal

Abstract

This chapter surveys the theory and evidence on contracting under learning and imperfect commitment. We present a simple model of long-term insurance à la Harris and Holmstrom (1982) to show the relevance and insights of the theory. Different variations of the model encompass many situations that have been studied in diverse areas of Economics, including Labor, Finance, and Insurance.

The model is useful for understanding issues such as dynamic selection and reclassification risk. Imperfect commitment is shown to be the source of adverse selection and partial insurance in environments with learning, even when information is symmetric.

The empirical literature has looked at the testable implications regarding selection and optimal contracts. Recent work has focused on the welfare loss from lack of commitment, which has been found to be substantial. The theory offers policy prescriptions on how to contend with the market distortions associated with limited commitment.

1 Introduction

The provision of insurance is one of the main determinants of how societies organize and regulate economic activity. Since insurance may create perverse incentives, different economic systems – and forms of capitalism – represent distinct alternatives over such trade-offs. The provision of long-term insurance requires commitment to prevent the exclusion of those with the most unfortunate realizations and maintain participation of the most fortunate. Such interactions between insurance provision, incentives, and commitment have been central to the contract theory literature since the 1980s.

*Support from NSF grant SES-1259770 and comments from Alessandro Lizzeri, Soheil Ghilli, Victoria Marone, Ariel Pakes, Bernard Salanie, and Steve Tadelis are gratefully acknowledged.

[†]Department of Economics, Northwestern University; igal@northwestern.edu

Economists have studied the provision of long-term insurance in the context of labor contracts (Harris and Holmstrom (1982), Holmstrom (1983)), insurance markets (Pauly, Kunreuther and Hirth (1985)), consumption and savings (Hall (1978)), as well as development economics, where villagers who lack sophisticated financial instruments may rely on mutual insurance within the village (Townsend (1994)). These agency problems arise between firms and customers, firms and employees, or among firms.

This chapter surveys the empirical literature on dynamic contracting. We focus on long-lasting relations between parties, in which the dynamics are driven by information revelation. Evolving information generates gains from long-term contracts to cope with reclassification risk. Reclassification risk is a concern in many markets, such as health or life insurance. The literature is useful for understanding insurance provision and market design.

Empirical work on dynamic contracting is sparse.¹ The typical challenge to empirical work on contracting, even in static situations, is the nature of the agency problem. Agency conflicts arise when critical information on types or actions is not observed by one of the parties. Typically, researchers are only as informed as the least informed party, which renders empirical work difficult. Dealing with dynamic agency relations entails clearing additional hurdles. In dynamic agency problems, the challenge is compounded by the fragility of the theory. Theoretical predictions are quite sensitive to the specific institutions. Commitment, renegotiation, timing, and the extent of information revelation over time have substantial impact on predictions. The theory's sensitivity can help leverage the empirics. For example, the sharply distinct predictions under different forms of learning (more on this in the next paragraph) can guide empirical strategies. However, empirical researchers face the challenge of finding situations and data with clean institutional arrangements that can be properly mapped into the appropriate testable predictions.

The literature has considered situations with symmetric and asymmetric learning. Learning is symmetric when all parties remain equally informed as new information arrives. An example of symmetric learning is health insurance, when all potential insurers have equal access to prior diagnostics and treatment information. This is also typically the case in life insurance; all parties use the same information at underwriting (assuming the risk). Alternatively, learning is asymmetric if only the current insurer has access to such information. This is the case when automobile insurers do not share accident histories. The current insurer, who observes the insuree's record, is at an informational advantage relative to competitors.

¹While empirics lagged behind Contract Theory, there is now a large empirical literature on static contracting, mainly focusing on informational asymmetries. For example, Cardon and Hendel (2001), Chiappori and Salanie (2001), and Cohen (2008) test for asymmetric information in different insurance markets. Later work quantified the welfare costs of adverse selection (Einav et al. (2013) and Handel et al. (2015)) in health insurance markets. Handel et al. (2015) studied equilibrium and welfare under different contracting regulations, such as rules on pricing pre-existing conditions.

Evolving information may lead to reclassification risk, dynamic selection, and their respective inefficiencies. Long-term contracts play a critical role in market performance. Consider the example of health insurance coverage at, say, age 50. Spot contracts insure the event risk associated with uncertain health expenses. By the time the spot market opens, however, quite a bit of information might be known about the insuree's expected health expenses. Since premiums depend on expected costs, spot contracts fail to insure the risk associated with the information revealed before the market opens. Premium insurance could be transacted before the type is known – say at age 25 – only if long-term contracts are feasible. The role and viability of long-term contracts depends on the nature of commitment to the contract by firms and consumers. If information is ex ante symmetric, full commitment (by both parties) yields efficiency. Parties can design a contract that equates marginal utility from consumption over time and states.

For legal and practical reasons, in most markets commitment is at best one-sided (unilateral). Firms can be held accountable, but consumers can walk away from the relationship without penalties (more on this later). The study of optimal unilateral contracts under symmetric learning was pioneered by Harris and Holmstrom (1982), who find that lack of consumer commitment compromises insurance, and thus welfare. Contracts suffer from negative dynamic selection. Adverse selection is the result of consumers' imperfect ability to commit to remain in the pool, rather than being a consequence of asymmetric information. As shown by Harris and Holmstrom, optimal contracts involve delayed reward to customers, which enhances customer retention. The delayed reward is funded through front-loading premiums in the case of insurance, or low initial wages in the case of employment. These mechanisms help alleviate reclassification (premium or wage) risk, at the expense of consumption smoothing.

To highlight the main forces at play, we will present the simplest version of Harris and Holmstrom (1982). After considering the theory and its testable implications, we review the empirical literature.

We look at several pieces of evidence that attest to the theory's relevance. First, we consider the working assumptions. When is learning important? What is the nature of commitment? Second, we review the literature that explores the theory's testable implications. Is selection negative? What is the relation between contractual terms and the likelihood that contracts will lapse? Third, do observed contracts resemble optimal ones? Finally, we move on to more recent work that has looked at welfare. What is the magnitude of the welfare loss from lack of commitment to long-term contracts? What proportion of the welfare loss from lack of commitment is restored by unilateral contracts?

Supporting evidence is found in the life insurance industry by Hendel and Lizzeri (2003). They show that, as predicted, virtually all contracts offered in the US and Canada were

front-loaded in their sample. Front-loading was found to be negatively associated with the likelihood that contracts will be allowed to lapse, which in turn reduces reclassification risk. Further evidence is presented from other insurance markets, such as health (Browne and Hoffmann (2013) and Atal (2015)), long-term care (Finkelstein, McGarry and Sufi (2005)), labor markets (Chiappori, Salanie and Valentin (1999)), and the evolution of the Kibbutz (Abramitzky (2008)).

We then discuss the literature on asymmetric learning, where outside parties are at an informational disadvantage. As noted previously, competing automobile insurers may not observe accident histories; similarly, potential employers may not observe prior job performance. Informational asymmetries make employers suspicious of new applicants, believing that they must be low types, as high types, presumably, remain at their current jobs. This suspicion lowers outside offers, locking workers into their current jobs, enhancing their commitment to the long-term relationship (Greenwald (1986), Waldman (1984)). Since lack of commitment compromises welfare, the endogenous lock-in might become a blessing.

The key testable predictions of the model with asymmetric learning are opposite to those under symmetric learning. Instead of negative retention, under asymmetric learning it is the bad draws who drop from the relationship. In turn, good draws who become locked in are profitable to the insurer (or employer). Ex ante competition dissipates those future rents, leading to lowballing in equilibrium, rather than to the front-loaded insurance contracts observed under symmetric learning. Insurers charge less than actuarial premiums to invest in customers, a proportion of whom will turn out to be profitable later on.

These distinct predictions across informational environments enable testing. Cohen (2012) tests for whether learning is asymmetric in the automobile insurance market in Israel, where insurers do not share information. In contrast to the symmetric learning case, Cohen finds positive selection and contract lowballing.

Information sharing between insurers would turn asymmetric learning into symmetric. As shown by de Garidel-Thoron (2005), such a move is welfare decreasing because asymmetric learning locks the non committed side, insurees, into the contract. Enhanced commitment (as in Crocker and Moran (2003)) enables the transfer of resources from the best states – i.e., the state in which in the individual would otherwise lapse – to less fortunate ones. The transfer lowers consumption variance, and thereby increases welfare. This idea has been studied in the banking literature in the context of borrower-lender relationships (Sharpe (1990)).

This survey starts with a simplified version of Harris and Holmstrom (1982) to explain the main forces at play and motivate the different evidence offered in the literature in the symmetric learning case. We discuss limitations and extensions of the basic framework. Among other we consider market features such as imperfect competition and switching costs, that

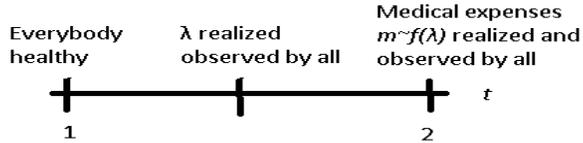


Figure 1

have received little attention in this literature. We then present the theoretical predictions in an asymmetric learning environment, and end by reviewing the asymmetric learning literature.

2 Symmetric Learning

2.1 A Simple Model

We consider a two-period model with a single risk to insure, suffered in period two. Consumers are healthy (i.e., face no loss) in period one. A signal, λ , that determines the distribution of period 2 medical expense m is observed by all parties prior to period 2 (symmetric learning). Expected medical expenses $E(m|\lambda)$ increase in λ . Consumers are risk averse, with preferences $u(c_t)$ and generate income y_t for $t = 1$ and 2. The insurance industry is competitive; namely, several firms offer an homogeneous product at each period and state.² Competition drives premiums for short and long-term contracts to expected actuarial costs. For simplicity, we assume no discounting and no borrowing.³ The timing is shown in Figure 1:

We assume a single period of uncertain expenses to highlight forces in the simplest way. We can interpret $t = 1$ as age 25, when no information has been revealed yet, and $t = 2$ as age 50. A more general model with yearly updates and health expenses is presented later.

Two benchmarks are useful: two-sided commitment and no commitment (spot contracts). The former delivers the first-best allocation, while spot contracts provide no reclassification-risk insurance. We will judge the effectiveness of dynamic unilateral contracts by comparison to these benchmarks.

²Product homogeneity fits a situation in which insurance is purely a financial arrangement, unbundled from health-care delivery. The delivery of care is differentiated by hospital and physician networks. See Section 5 for a discussion of imperfect competition.

³We later elaborate on this assumption. The key is imperfect capital markets that make borrowing costly (not necessarily absent).

2.2 Benchmark I: Full Commitment to Long-Term Contracts

Both parties commit to contracts offered at $t = 1$ (age 25), before the health type is known. Because contracting takes place prior to learning λ and the two sides are committed, both the risk associated with m given λ and the risk associated with λ are fully insured, equating marginal utilities, $u'(c)$, across states.

Competition drives the sum (present value) of first- and second- period premiums p_1 and p_2 to $p_1 + p_2 = E_\lambda(E(m|\lambda))$. Since both parties commit to the contract, p_1 and p_2 can be timed to smooth consumption, equating $u'(c)$ across periods as well. The allocation is first best.

2.3 Benchmark II: No Commitment

Absent commitment, long-term contracts are not feasible. The spot insurance market opens once λ is known. Under competition, full-event insurance is offered at actuarially fair premiums: $p(\lambda) = E(m|\lambda)$. Thus, uncertainty in m is fully insured in equilibrium, but uncertainty in λ (reclassification risk) is left uninsured.

Lack of commitment prevents eliminating the risk associated with λ . Individuals revealed to be bad risks in period 2 end up paying high premiums.

Reclassification risk represents one of the main motivations behind States' regulation of health insurance market, specifically the 1996 Health Insurance Portability and Accountability Act (HIPAA) and the Affordable Care Act (ACA). Both States' regulations and HIPAA impose guaranteed renewability of insurance, and most States forbid individualized premium hikes. The ACA goes further to forbid the pricing of pre-existing conditions, which in effects eliminates both renewability concerns and individualized premium hikes. Lack of commitment might be a reason for the poor performance of unregulated individual insurance markets: If parties committed to future contractual terms (and especially if firms did so) there would be no need for pricing regulations..

2.4 Long-Term Contracts: One-Sided Commitment

With the two benchmarks at hand, we can evaluate the performance of unilateral dynamic contracts. Firms offer long-term contracts $\{p_1, p_2\}$, which entail a first period-premium and commitment to a second-period premium.⁴ Consumers do not commit to the policy, and can let the contract lapse without penalty. We assume that the second-period premium is not contingent on the realized λ . We do so without loss of generality; as shown by Hendel

⁴In principle, a contract could also specify a co-insurance rate. Handel et al. (2015) show that equilibrium unilateral contracts offer full-event insurance, namely, zero co-insurance.

and Lizzeri (2003), the competitive allocation can be equally achieved with or without state contingent contracts.

Since information is symmetric, at $t = 2$ consumers can get coverage for $E(m|\lambda)$ on the spot market. They let the long-term contract lapse in states for which:

$$E(m|\lambda) < p_2.$$

The last inequality implies adverse selection in the second period: Better risks are those that lapse the long-term contract.

To find the equilibrium, we can use the fact that the competitive equilibrium contract maximizes consumer welfare subject to (i) the ex ante break-even constraint and (ii) the lapsation constraint, which accounts for the states λ in which buyers remain in the pool:

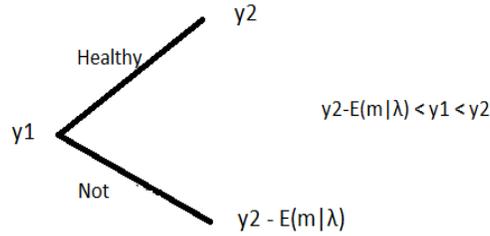
$$p_2 \leq E(m|\lambda). \tag{1}$$

In equilibrium there is full event insurance, namely, all medical expenses are insured (i.e., no out-of-pocket payments). For a proof, see Handel et al. (2016).

To understand how the optimal contract works, notice that the lower p_2 is, the fewer the states in which the insuree lapses. Fewer states in which the lapsation constraint binds means more states across which $u'(c_2)$ is equated. In other words, lower p_2 means more second-period premium insurance.

A low enough p_2 could fully eliminate reclassification risk. However, for the contract to break even ex ante, a low p_2 requires a high p_1 ; that is, front-loading is necessary. Front-loading is costly in terms of consumption smoothing, and therefore optimal contracts trade off reclassification-risk insurance and consumption smoothing.

The cost of front-loading depends on y_1 . Individuals with more limited initial resources, or more limited access to borrowing, prefer less front-loaded contracts, and in turn end up suffering more reclassification risk. More premium uncertainty (namely, fewer states being pooled due to higher lapsation) leads to more intense adverse dynamic selection, which translates into higher present value of premiums. This is formally proved by Hendel and Lizzeri (2003). Intuitively, more front-loading means lower premiums later in the contract. In this example, it means a lower p_2 . With a lower second-period premium, fewer types drop coverage. Since it is the good draws who lapse, lower lapsation means a healthier pool is retained. Under competition, firms break even, so that the present value of premiums proxies expected actuarial costs. Thus, the worse pool kept in the long run, when front-loading is low, translates into higher present value of premiums. To illustrate, if $p_2 = 0$, all types stay with the contract so that $p_1 + p_2 = E_\lambda(E(m|\lambda))$. Instead, absent front-loading, $p_1 = 0$, all but the worse type find a better price on the market and lapse; then $p_1 + p_2 = E(m|\bar{\lambda})$,



where $\bar{\lambda}$ represents the worse health state. Naturally, $E(m|\bar{\lambda}) > E_{\lambda}(E(m|\lambda))$ for all non degenerate distribution of λ , illustrating the negative link between front-loading and present value of premiums.⁵

It is worth emphasizing that the source of adverse selection is not asymmetric information, but rather lack of commitment. Information is symmetric, but consumers are unable to commit to remain in the contract in good states. Selection, or adverse retention, is driven by the inability to retain good risks.

Consumers would be better off ex ante if they could commit to stay with the policy in good states, so that resources in those states could be transferred to the less fortunate ones, up to the point where marginal utility is equated across states.

The following figure illustrates the problem in a two-state world, in which the healthy state entails no medical costs. Absent costs, insurees simply walk away in the good state. Transferring resources from the good state to less fortunate ones is not feasible. Up-front payments, however, are feasible. If the first-period income, y_1 , is higher than $y_2 - E(m|\bar{\lambda})$, the individual can transfer resources from the first period to the unhealthy state of the second period.

Notice that front-loading differs from saving. Savings transfer resources to all period-2 states. In contrast, the up-front payment goes exclusively to the bad state. As we will see later, with more states, the front-loaded amount goes to those states with higher marginal utility of consumption than the first period. The first-period marginal utility increases as more resources are transferred to the future. The optimal amount of front-loading, and the states pooled, are determined by equating marginal utilities across the first period and the bad second-period states.

⁵Naturally, perfect capital markets would enable paying all premiums up front, restoring consumer commitment. As we will see later, full insurance would require payments in the tens of thousands of dollars for health or life insurance when the consumer is in her late 20s, which does not seem affordable for most buyers. Pauly, Kunreuther and Hirth (2005) find the minimum premiums that guarantee no unraveling. Absent capital market imperfections, such an allocation would be first best. Cochrane (1995) discusses a market for premium insurance, which under perfect capital markets achieves full insurance.

In sum, lack of consumer commitment compromises reclassification-risk insurance and creates adverse retention. Optimal contracts resort to front-loading to partially restore consumer commitment. The next section shows the relevance of these simple insights through evidence from different areas of Economics.

3 Empirics

The relevance of the theory can be assessed in several ways. For instance, one can identify the markets in which the main working assumptions of the theory are not only relevant, but actually capture the main forces at play. After identifying those markets, the natural next step is to test the implications of the theory on selection, lapsation, and reclassification risk. We then look at contracts. Do observed contracts resemble those predicted by the model? Finally, we ask what is the welfare loss associated with limited commitment?

3.1 Main Assumptions

The theory's predictions are very sensitive to the main market characteristics, such as the type of commitment and the nature of learning. Sharply different predictions are in principle a blessing for testing. The challenge for the empirical researcher, given the predictions' sensitivity, is to find institutions that can be reasonably mapped into specific assumptions; Such a mapping is complicated when too many forces are at play. It is important, therefore, to identify simple-enough situations to isolate the key forces. We concentrate on markets where learning and imperfect commitment are the key determinants of contractual relations.

3.1.1 Symmetric Learning

Learning is prevalent in many economic situations; for example, learning about product quality, worker productivity, or the quality of a match.

Learning might or might not be symmetric. We refer to learning as symmetric when all parties, including competitors, receive the same signals. In many markets learning might not be symmetric. For instance, signals about workers' performance are revealed over time – yet, work is often done in teams, hindering worker-specific learning. In many situations, supervisors and close co-workers are able to monitor performance, but potential employers cannot. Such asymmetries will be discussed in the next section.

An example of a labor market with arguably symmetric learning is academics. Output, both teaching and research, is made public or done in public. Presentations and publications

generate observable signals. While coauthoring may interfere with symmetric learning, productivity is to a large extent attributable. As we see later, academics' compensation patterns seem to resemble those predicted by Harris and Holmstrom (1982).

Many lines of insurance involve learning. For example, health status evolves over time. Both existing conditions and prior diagnostics are observable, which make learning arguably symmetric. Life and health insurers can discover such signals prior to underwriting, namely, at the time the insurer assesses and assumes the risk.⁶ Regulation often prevents pricing pre-existing conditions, but the underlying information structure is well approximated by symmetric learning.

Accident history in automobile insurance can be public or not, depending on regulation or firms' decision to share information. The US has a national loss-underwriting database, to which insurers provide information about every home and auto insurance claim. The database can be accessed by all insurance companies at the time of underwriting. In other countries, information is not shared among insurers (Cohen (2013)). While there is learning in the automobile market, whether learning is symmetric or not depends on market-specific institutions. As another example, credit-scoring agencies arguably make learning symmetric in credit and mortgage markets.

3.1.2 Unilateral Commitment

In most markets, consumers and workers are free to withdraw from contractual arrangements. Legal considerations make termination fees hard to enforce, and are thus rarely used.⁷ Other reasons for partial commitment are discussed by Daily, Hendel and Lizzeri (2008), Fang and Kung (2012), and Bayot (2015). As these authors argue, uncertainty about future need for the product makes commitment costly. Individuals may not want to commit, or to pay health premiums up front, in the private market if they might later switch to an employer who offers more generous health-insurance coverage. Similarly, the possibility of a divorce may detract from the value of committing to life insurance coverage.

Regardless of the reason, we observe contracts with one-sided commitment in many industries: life insurance, health insurance, mortgages, and academics' compensation.

Long-term contracts arise in some markets without intervention, such as US and Canadian life insurance. In other markets – for instance, the health insurance markets in Germany and Chile – government regulation requires that firms offer renewable contracts in a specific

⁶Concealing information at underwriting leads to contestability by life insurers.

⁷The reason termination fees are hard to enforce is that damage due to lapsation is difficult to prove (as a liquidated damage). An example is the infamous early-termination fee for cell phones, which was declared illegal by state courts.

format (Atal (2015), Browne and Hoffman (2013)), which resembles the theory's prediction. Both regulated and unregulated markets provide evidence on the theory's relevance.

3.2 Testable Implications

Learning and imperfect commitment seem to capture the key forces that shape contracts in many markets. We now turn to empirical work on those markets. One can look at two types of evidence. First, regardless of whether contracts are optimal or not, learning and limited commitment are predicted to have implications on reclassification risk, selection, lapsation, and the link between lapsation and front-loading. Second, one can compare actual and predicted contracts. The distinction between testing the implications of the assumptions, as opposed to the prediction on contracts, is particularly relevant in regulated markets. Regulated contracts need not resemble those predicted by the theory and, if they do, they do not represent evidence in support of the theory. Even so, regulated markets such as the Chilean health-insurance market, still help in testing for learning and consumer commitment.

3.2.1 Long-Term-Care Insurance

Finkelstein, McGarry and Sufi (2005) study the long-term-care insurance market in the US. It is a thin market, where only a minority of the elderly population with potential long-term-care needs buys coverage. The authors present evidence on pricing, which shows that long-term-care insurance contracts are front-loaded. Despite the front-loading, a substantial share of insurees lapse.

Under symmetric learning and limited commitment, the theory predicts that the better types are those that lapse. Using data from the Health and Retirement Survey, the authors regress eventual utilization on lapsation to provide evidence of adverse retention. The reported negative relation suggests that, in accordance with the theory, good risks are more likely to lapse.

3.2.2 Health Insurance

Patel and Pauly (2002) discuss contract renewability in the individual health-insurance market in the US. They argue contracts are dynamic, with renewability guaranteed by the HIPAA. The law is incomplete, however, as it fails to ban individual premium changes within a rating class. Namely, while individual coverage must be renewed, premiums could depend on recently developed health conditions. Patel and Pauly (2002) survey state insurance regulators and find that all but three states forbid individualized premium hikes: The pool can face a premium hike, but not specific individuals within it. State law rather than

HIPAA makes contracts dynamic. Patel and Pauly conclude that guaranteed renewability works to eliminate premium risk.

Judging by the proportion of the population that was uninsured prior to the ACA, the US individual health-insurance market seemingly did not function properly. Some commentators, including Patel and Pauly (2002), attribute the market's unraveling to state regulation that prevented firms from charging individualized prices at the time of underwriting. Inability to price pre-existing conditions is expected to result in an Akerlof-type unraveling. In other words, while dynamic contracts are offered in the US, adverse selection generated by restricted underwriting appears to have caused an inefficient level of trade.

Dynamic contracts are also offered in Germany and Chile (Browne and Hoffman (2013), Atal (2015)). Government regulation in Germany permits premiums to depend on health conditions when coverage begins, but must remain fixed thereafter. Basically, risk rating is legal at underwriting, but later premiums cannot vary based on new conditions and must remain constant (level premiums). Since health expenses increase with age, level premiums mean that contracts are front-loaded.

Consumers are allowed to switch insurers later in life, and thus these are unilateral contracts. However, the new policy will reflect their increased age and any changes in health conditions. The regulation that forces level premiums, which can only be updated by switching, locks insurees into their policies.

Browne and Hoffman (2013) using data from a large insurer in the private German health insurance market report evidence consistent with the model: Front-loading lowers lapsation, and better risks are those that lapse.

Atal (2015) studies the costs of lock-in. Regulation in Chile forbids individualized premium increases, thus creating consumer lock-in. Atal shows that actuarial costs divided by premiums increases over time, which confirms front-loading and thus lock-in.

When insurers are differentiated, lock-in may be costly. For example, as their conditions change, insurees may prefer a different hospital network; Atal (2015) quantifies the inefficiency due to inability to change insurers (more details in Section 5).

Another possible hurdle to long-term contracting is uncertainty about future health costs, which is a non-diversifiable risk that firms may be reluctant to take. Basically, if they are unable to predict costs far into the future, they may not want – or be able – to commit to future premiums. German regulation deals with the problem by indexing premiums to the aggregate cost of health care. That way, firms insure the idiosyncratic health risk, but do not suffer aggregate shocks to health costs. Indexing to a health-care price might assist the market in developing elsewhere.

3.2.3 Life Insurance

Hendel and Lizzeri (2003) use data on contract dynamics of term life insurance. Term insurance, unlike whole life, is a simple and homogeneous product that provides coverage for a specified period (often up to age 70) as long as the policy is renewed. The authors use explicit contract data from Compulife, a pricing software used by insurance agents. Compulife quotes not only premiums at underwriting for different health statuses, but also future premiums guaranteed by the insurer. Premiums to be paid in the future are, for some contracts, contingent on health status. These health-contingent premiums embed premium risk: Namely, individuals will pay different premiums depending on how healthy they remain.

Does the competitive model with symmetric learning and unilateral commitment (as in Harris and Holmstrom (1982)) fit the life-insurance industry? It is a pretty competitive industry, with hundreds of life insurers. Insurers commit to future policy terms, and there are no termination fees. Finally, health type evolves over time (more on health transitions later), while medical examinations and questionnaires at underwriting arguably make learning symmetric.

Learning and Reclassification Risk Table 1 illustrates a typical term contract, popularized in the 1970s, called Select and Ultimate Annual Renewable Term. The table should be read as follows. The first row shows premiums guaranteed at age 40 to an insuree who has just gone through underwriting (i.e., a medical examination to make sure he qualifies as a preferred risk) as he ages, from the first to the twentieth year of coverage. The insuree would start paying \$370 at age 40, and in year 11, at age 50, he is guaranteed to pay no more than \$2,555. The second row shows the premiums paid if he remains healthy after a year of coverage: At age 41, he would pay \$385 if he is able to produce a letter from his physician attesting to his good health. While he was guaranteed a \$475 premium, insurees who requalify as good risks get a discount to \$385. At the time of contracting (age 40), the insuree knows what future premiums will cost under different contingencies.

As we move down column 1, we find premiums for an insuree who remains healthy. Once he fails to requalify, horizontal moves depict the premiums that will apply, regardless of how unhealthy he is. For example, the premium to be paid by a 51-year-old who remains in good health until age 49 is \$1,080.

These contracts became popular in the 1970s in response to market competition. Until then, annual renewable term (ART) policies were not state contingent – that is, premiums depended on age and not health condition; Good draws would simply lapse. These contingent contracts were designed to match the better terms offered by competitors, thereby preventing the lapsation of good draws.

Table 1 highlights the role of learning: Premiums depend on the information revealed during the contracting period. The contract embeds premium risk. At age 59, the insuree could be paying as little as \$1,340 or as much as \$6,375.

Notice that learning is essentially symmetric, in the sense that the same medical examination that triggers a discount can be used by competitors to generate the outside offers.

Table 1: S&U ART premiums

Policy Year							
Age	1	2	3	10	11	19	20
40	370	475	640	1,485	2,555	5,680	6,375
41	385	490	660	1,565	2,815	6,375	7,040
42	400	530	690	1,705	3,105	7,040	7,790
49	630	890	1,080	2,725	6,375	13,675	14,785
50	690	945	1,155	2,895	7,040	14,785	15,765
51	735	1,050	1,295	3,230	7,790	15,765	17,230
58	1,245	1,750	2,295	6,420	14,785	33,165	35,445
59	1,340	1,785	2,480	6,945	15,765	35,445	38,715

Note: Contracts offered in 7/1997 to a preferred male non smoker for \$500K coverage.

S&U ART=annual contract that allows for reclassification by showing good health.

Front-Loading Theory predicts that a variety of contracts should be offered that cater to individuals with different income profiles, who find front-loading more or less costly. While indeed a variety of contracts are offered in the US (and Canadian) market, virtually all contracts were front-loaded at the time of the study.

One can use the slope of premiums over time as a measure of the extent of front-loading. Actuarially fair premiums, which are expected to arise in a competitive industry that offers spot contracts, should increase at the same rate the death probability increases with age. Premiums that increase slower than death probabilities reflect front-loading. Such contracts do not break even period by period, which suggests that contracts are dynamic.

The slope of premiums can be proxied by the ratio $Q(1st)/Q(11th)$, the premium in the first year of the contract divided by the premium 10 years later. Table 2 shows that on average $Q(11th)$ is twice $Q(1st)$. However, the range of premium slopes is quite wide. Premiums increase as much as ninefold for some contracts, while premiums remain flat in level-term policies. Naturally, level-term policies entail substantial front-loading.

It is interesting to note that even the steepest premiums entail front-loading. In other words, death rates increase faster than even the steepest premiums. Using US actuarial

tables, conditional on being in good health at age 40 (namely, qualifying as a preferred risk), the probability of death at age 59 is 17.2 times higher than at age 40; the steepest premiums are only 9 times larger at 59 than at 40. Basically, all contracts in the sample are dynamic – that is, they do not break even period by period.

Table 2				
Slope and Cost Dispersion across Contracts				
	Mean	Std Div	Min	Max
Q(1st)/Q(11th)	0.43	0.31	0.11	1.00
PV	16,055	5,245	6,871	28,754
Q(1st)/Q(11th)=ratio of first to 11th premium.				
PV=present value of 20 years of coverage at r=0.08.				

The last row of the table displays the present value of 20 years of coverage starting at age 40. Under competition (zero profits), the present value of premiums proxies the actuarial cost of covering the respective pools retained by the different policies. The range of present values of coverage varies fourfold, from about \$6,800 to \$28,700.

Because of dynamic adverse selection, we expect a negative relation between $Q(1st)/Q(11th)$ and the present value of coverage. This prediction is put to the test in Table 3.

Dynamic Selection The theory’s main prediction is that increased commitment through front-loading leads to better dynamic selection. Since good draws are those predicted to lapse, lower lapsation means a better pool, which in turn translates into lower costs of coverage. The prediction is tested by regressing the cost of coverage on the premium slope. Hedonic-type regressions, presented in Table 3, show not only that the correlation is negative, but the main determinant of the cost of coverage is premium front-loading. Premium slope accounts for 60% of the log variation, while numerous other policy characteristics (convertibility, renewability, etc.) have limited impact on the cost of coverage.

Table 3		
	log(PV)	
	(1)	(2)
Q(1st) / Q(11th)	-1.06	—
	(-16.79)	—
Other contract characteristics	—	—
	Y	Y
R ²	74.4	16.6
N	125	125

Dependent Variable log(PV) is the log of the present value (r=8%) of the cost to the consumer of 20 years of coverage starting at age 40.

Key Findings Available contracts are state contingent, which confirms the prevalence of learning and relevance of reclassification risk. All contracts in the US and Canada were front-loaded at the time of the study, which suggests that front-loading plays an important role in coping with reclassification risk. Lapsation is higher for less front-loaded contracts. More front-loading is associated with lower present value of premiums, which reflects a healthier pool. The increased commitment through front-loading, therefore seems to reduce dynamic negative selection. Quantitatively, premium slope accounts for the majority of premium variation.

3.2.4 Labor Market

While Finkelstein et al. (2008) observe long-term-care utilization, and Abramitzky (2008) (discussed below) observes workers' types, in many situations, especially in the labor market, types are not observed by the researcher. To overcome this lack of observability, Chiappori, Salanie and Valentin (1999) derive testable implications on promotions rather than productivity. Promotion dynamics reflect the interaction between performance and the optimal contract, and thus help uncover symmetric learning.

Specifically, the authors show that in a situation with symmetric learning and downward wage rigidities, as in Harris and Holmstrom, wages, and thus promotions, display what they term the "late beginner property."

In short, if we compare two individuals currently at the same rank (or wage levels), but one of them rose in the ranks earlier than the other (having both started at the same initial level), the late beginner is expected to do better in the future. Intuitively, the current wage of the early riser, who later slowed down, because of downward rigidities might conceal sufficiently negative information that merits a wage downgrade. In contrast, the late beginner's

pay reflects her current performance. In other words, early starters' rank eventually conceal interim bad performance and are therefore expected, other things being equal, to do poorer than late beginners. Chiappori et al. (1999) test their predictions using data on French public servants' rank and the timing of promotions.

3.2.5 Alternative Forms of Commitment

Looking at alternative forms of commitment, other than front-loading, can attest to the role of imperfect commitment in contracting. Crocker and Moran (2003) study the bundling of health insurance and employment. The idea is that attachment to a job enhances individuals' lock-in to their employer-sponsored health coverage. Commitment is expected to translate into more generous coverage.

Crocker and Moran (2003) use the National Medical Expenditure Survey and information on job attachment to link the nature of insurance offered to the type of job. They show that health insurance generosity (coverage and lifetime limits) is associated with worker immobility.

Their findings are consistent with the pervasiveness of employer-sponsored health insurance coverage in the US, which suggests that employer-sponsored health insurance may work better than prepayment. As argued by Atal (2015), prepayment can be problematic in health insurance, especially when insurance (i.e., the financial side of coverage) is bundled with health care.

On the other hand, employer-sponsored insurance is not ideal either, as it creates job lock-in (Currie and Madrian (1999)) potentially leading poor worker-employer matches. However, lock-in may also offer benefits. Lock-in to an insurer or an employer may increase incentives to invest in health. Health, as a form of general human capital, delivers immediate as well as future benefits. The lower the turnover, the more the current employer or insurer benefits from future savings. Fang and Gavazza (2010) present evidence of underinvestment during working years in jobs with high turnover, which translates into higher medical expenses after retirement.

Crocker and Moran (2003) interpret their findings as evidence that the key hurdle to the functioning of private health-insurance markets, especially the individual market, is commitment rather than asymmetric information. When commitment is restored by job lock-in, insurance provision improves.

3.2.6 Further Afield: Economic Institutions

Abramitzky (2008) presents a concrete example of how social institutions are designed to provide insurance, by examining economic equality in the Israeli kibbutzim (plural of kibbutz). A Kibbutz is a collective community – originally primarily an agricultural community – that aims to achieve full equality among its members. Equality is understood from an ex ante perspective, before the talent, human capital, and market opportunities of its members are known, as a form of insurance.

Kibbutz formation as an institution started in the early 1900s and reached about 130,000 members at its peak. Participation by kibbutz-born individuals is voluntary (people from the outside must to be admitted by the kibbutz); namely, commitment to the institution is unilateral.

For most of their history, kibbutzim offered full equality among its members. However, recent negative financial shocks, such as the loss of government subsidies, decline in world agricultural prices, and bad investments, led them to shift away from full equality toward different, more limited degrees of equality.

Members' earnings go to the kibbutz, which then budgets according to different compensation schemes. Some kibbutzim allocate all members an equal budget; others allocate based on individual earnings.

The theory predicts that higher wealth is associated with a more egalitarian distribution, which, in turn, translates into lower exit levels. Ex ante, members want equality (insurance), but at some point each individual learns their type (including human capital), which determines their market opportunities. At that stage, they may pursue alternative employment, which is mainly in the city. Commitment, in the form of initial resources, achieves partial reclassification-risk insurance.

Abramitzky combines data on about 180 kibbutzim with individual-level census data to test whether the wealth of the kibbutz worked as a lock-in mechanism. The key information on the kibbutz is their wealth, how egalitarian their budgeting is, as well as their ideology. Census data are used to track migration to and from the kibbutz, by profession.

The main findings are that wealthier kibbutzim retained higher equality. Entry and exit are associated with negative selection, namely, those that leave for the city are the individuals with higher talent, measured by their wages.

More specifically, evidence on the talent of those leaving the kibbutz is supported by the finding that less educated former kibbutz members earn more than similar individuals in the city; that is, the good draws (in terms of talent) are those that depart. In addition, more educated former kibbutz members earn less than similar individuals: for instance, most MBAs leave, not just the good ones. Entering members have lower wages than similarly educated individuals had before entering.

4 The Welfare Implications of Long-Term Contracts

The papers just reviewed provide evidence on the relevance of the theory. A natural next step is to assess the impact of the identified distortions on welfare. In short, how painful is reclassification risk in the absence of dynamic contracts? How effective are long-term contracts for eliminating reclassification risk?

4.1 Regulation of Health Insurance Premiums

Understanding the performance of different contractual arrangements in contending with reclassification risk is important for policy design. Reclassification risk is one of the key motivations behind health insurance market regulation. The ACA bans pricing health conditions. One of the goals of HIPPA, the ACA and numerous state regulations was to eliminate reclassification risk and coverage denials.

While effective in eliminating premium risk, banning the pricing of health conditions comes at a cost. By imposing a uniform price on a heterogeneous population, the ban results in adverse selection, and possibly the full collapse of the market. This concern is supported by the findings of Handel et al. (2015). They simulate the functioning of health exchanges using data from a large employer. They compute the actuarial risk of every person in the population using Adjusted Clinical Group software (described in the next section). Preferences towards risk are recovered through a choice model that fits observed coverage choices. Using preference estimates and the distribution of risk types, they simulate the equilibrium of a market in which high and low coverage policies are offered by multiple insurers. The Nash equilibrium is then computed for numerous populations and different pricing rules. Unraveling is pervasive when age is priced. The equilibrium involves full unravelling to the minimum coverage, 60% actuarial value.

As an alternative to the ban on pricing health conditions, long-term contracts may permit addressing reclassification risk without inducing adverse selection. For instance, the contracts described in Browne and Hoffmann (2013) avoid adverse selection by allowing for the pricing of pre-existing conditions at underwriting, while at the same time eliminating premium risk by forcing constant premiums over time. The welfare gains delivered by long-term contracts, in particular vis-a-vis ACA regulation, is studied by Handel et al. (2016).

The first step in assessing welfare is to characterize contracts under different regimes. The model described in the next section is used by Handel et al. (2016) jointly with an individual-level panel of workers at a large employer (25,000 covered lives) to simulate equilibria and compare welfare in different situations.

The key component of the data is the detailed information on health realizations: all medical claims, including the ICD-9 diagnostic codes. The diagnostic codes are used in

conjunction with Adjusted Clinical Group (ACG) software (which was developed at the Johns Hopkins Medical School to assess the actuarial risk of individuals) to generate a risk score for each individual in the population. In turn, the score is used to compute health state transition matrices. All those parameters are fed into the following model, to predict the shape of optimal dynamic contracts.

4.1.1 Model

The model in Handel et al. (2016) is a T -period version of the one described in Section 2.1. The T -period model is meant to capture a population aged 25 to 65, from college to Medicare. The model predicts optimal history-contingent premiums for a given income profile $\{y_t\}$, parametrized by risk preferences, the distribution of health costs, and transitions across states.

Consumer preferences are: $U = E [\sum_t \delta^t u(c_t)]$. Health state λ_t (ACG) determines expected health costs, $E[m_t|\lambda_t]$. Information is symmetrically revealed: m_t and λ_t are commonly observed by consumers and firms. The insurance industry is competitive, firms are risk neutral, and the discount factor is δ . Capital markets are imperfect: savings or borrowing.

As in the simple model, full commitment delivers the first best allocation. No commitment leads to full event insurance at actuarial premiums $E[m_t|\lambda_t]$, so that insurees are fully exposed to the information revealed by λ_t . The ex ante welfare of a representative consumer under these benchmarks is immediate to compute.

4.1.2 Equilibrium under Unilateral Commitment

Handel et al. (2016) show that equilibrium contracts involve full insurance against medical risk, are front-loaded, and premiums are such that consumption is downward rigid.

Premiums, at the start of a contract, are designed to guarantee a minimum consumption level. No matter how bad the health state turns out to be later in life, consumption is guaranteed never to decline. This guarantee is funded by the up-front premiums.

The consumption guarantee is bumped up to match the outside offers that arise every time the consumer receives good news – that is, when her type improves. The newly guaranteed consumption level is the first-period consumption of an optimal contract that would start at that date and improved state. Basically, the consumption level of the outside offer is matched.

The consumption guarantee parallels the downward rigid wages in Harris and Holmstrom (1982). Optimal wages increase when the worker's productivity is upgraded, but do not decline after bad news. These bad states are subsidized with low initial wages. Like front-loading, the delayed reward generates worker lock-in.

While first-best involves equating marginal utilities across all states and periods, under unilateral commitment only states and periods without tempting outside offers can be pooled. Why is reclassification risk left uninsured? Insurees would like to transfer resources from the healthier states to the less fortunate ones, but lack of commitment implies that they will exit the contract when a more attractive offer comes. Since resources can be sent costlessly forward, at actuarial value, consumption is not expected to decrease even if health deteriorates.

4.1.3 Optimal Contract for Flat Net Income

Optimal contracts are recursively computed to find consumption guarantees. Table 5 presents first period consumption, premiums, and front-loading for a consumer with flat net income, namely, constant $y_t - E[m_t]$. The main interest in this income profile is that gains from long-term contracting are not driven by life-cycle considerations (saving and borrowing), since individuals with such an income profile do not want to borrow or save.

State	1	2	3	4	5	6	7
Consumption	52,548	51,143	49,642	49,168	46,414	43,408	37,294
Premium	2,750	4,155	6,008	6,130	8,885	11,890	18,554
Front-loading	1,619	1,864	2,228	2,155	3,035	1,235	–

Except for state 7, in which the individual is as unhealthy as she will ever be, premiums are well above actuarial value. Front-loading is substantial in all but the worse state. Those savings are used to guarantee future consumption in those states in which bad news arrives. Subsequent premiums (not shown here) are contingent on more complex histories. Should good news arrive, consumption is upgraded. Otherwise, initial promises are kept in all future periods.

4.1.4 Welfare

Table 6 presents certainty equivalents CE_X . CE_X is the constant monetary amount that makes the individual indifferent to scenario X . The following contracting scenarios are considered: D denotes dynamic contracts with one-sided commitment, S denotes spot contracts, and TS denotes two-sided commitment.

	TS	D	S	ACA	
CE_x	53.67	52.77	46.27	51.30	
Estimated median CARA parameter $r=0.0004$					

CE_{TS} reflects welfare under the optimal allocation. The gap between CE_{TS} and CE_S represents the welfare loss from reclassification risk due to lack of commitment. The loss is about 14% of welfare. As expected, CE_D is between CE_S and CE_{TS} . One-sided commitment recaptures almost 90% of the gains from full commitment relative to spot contracts. Handel et al. (2016) study how these gains depend on risk aversion, switching costs and income profiles.

Handel et al. (2016) also compute welfare under the ACA, to assess the effectiveness of dynamic contracts relative to other ways of contending with reclassification risk.

5 Extensions and Limitations

Several additional market features deserve attention, especially for applied work. Switching costs have been documented to be substantial in insurance markets (Handel (2013)). We also consider non-health-related determinants of lapsation. For example, employment-related insurance may render participation in the individual market unnecessary. We look at imperfect competition, which is quite relevant in the exchanges; in many exchanges only a handful of insurers compete.

5.1 Switching Costs

Handel (2013) shows the importance of switching costs in health insurance demand, using the data described in the previous section. He does so in the context of static choices, studying the link between inertia and adverse selection. When switching costs are high, individuals are less likely to act on the information they learn, thus attenuating adverse selection.

Switching costs are typically associated with back-loading. Consumers are attracted with low prices, which are then raised to exploit lack of price responsiveness. It is thus important to consider how switching costs are expected to affect dynamic contracts, especially front-loaded ones.

In the context of dynamic contracts, switching costs induce commitment, reducing lapsation. The model described above predicts that good draws will lapse. Bad risks stay with the contract, as they find no better deals on the market. Thus, adding inertia to the model affects mostly the behavior of good risks: It becomes easier to retain good risks in the pool. While we have not seen any formal results with switching costs, the enhanced commitment due to inertia is expected to be welfare increasing, as in de Garidel-Thoron (2005). In de Garidel-Thoron lock-in is endogenous, due to asymmetric learning. Keeping good risk in the pool enables the transfer of resources from states with low to states with high marginal

utility from consumption, which lowers consumption variance. We present more details in Section 6.1.

5.2 Non-Health-Related Lapsation

Insurees lapse for many reasons besides being in good health. For instance, a divorce may render a life-insurance policy unnecessary or a new job may offer health-insurance. Those events might generate the lapsation of types that are expected to remain in the pool, namely, the bad risks; good risks are expected to lapse regardless.

Lapsation has conflicting effects on the willingness to commit to long-term contracts through front loading. On the one hand, uncertain about their future needs, buyers may be less willing to front-load future payments. On the other hand, the lapsation of bad risks reduces the front-loaded amount necessary for the policy to break even ex ante.⁸ The main features of the predicted contracts remains unchanged.

5.3 Imperfect Competition

Most of the literature assumes that insurance is a purely financial product. As such, one would expect little product differentiation – which jointly with the assumption of many sellers, justifies looking at the industry as competitive as a first approximation. One can argue that competition is a reasonable assumption for the life insurance industry, in which the product itself is a monetary transfer and hundreds of insurers compete.

This assumption is less realistic in other contexts. Most health insurance products available in the exchanges are bundled with health-care provision (Dafny et al. (2015)), and plans are differentiated by the hospital and physician networks they offer. In addition, only a few insurers compete in most states' exchanges. For health insurance markets, and especially in the exchanges, pure competition is not a good approximation.

Imperfect competition is harder to model. Characterizing the equilibrium dynamic contracts under competition entails maximizing consumer utility subject to the break-even and lapsation constraints. Instead, finding the equilibrium under imperfect competition requires finding the dynamic contract that maximizes each seller's profits subject to the lapsation constraints, given the contracts offered by other sellers. Namely, one has to find the fixed point of the contracts that best respond to each other, which is much more complicated than finding the competitive outcome.

⁸Interestingly, in the life insurance industry the term "lapsation based products" is quite common. It refers to firms' expectation that they will not fully compensate bad risks, under the assumption they are likely to lapse: if they did not, the product would not be profitable (Daily (1989)).

How would optimal contracts look in an imperfectly competitive market? The easiest set-up to consider involves two competitors, two periods, and the following preferences:

$$U_t(c_t, j) = u(c_t) + \varepsilon_{jt}, \quad j = a, b.$$

Product differentiation enters through a brand-specific shock that is additive to the utility from the financial terms of the policy. Consider stable preferences by assuming $\varepsilon_{j1} = \varepsilon_{j2}$.

As the second period spot market opens, for consumers in every state s , firms can poach customers currently covered by their competitor's long-term contract. Due to product differentiation, firms face downward-sloping demand for second-period coverage. Unlike the competitive case, firms charge a second-period spot premium above actuarial costs. Thus, the lapsation constraints of the first-period contracts are relaxed relative to the competitive solution, in which outside offers are at actuarial costs. In other words, spot market mark-ups reduce incentives to lapse. Contracts still display front-loading to enhance long-term insurance, and the good risks are those expected to lapse, but it appears that by relaxing the lapsation constraints, market power enhances commitment to the long-term contract.

5.4 Lock-In and Changing Needs

In many situations – and health insurance in particular – needs may change, so that $\varepsilon_{j1} \neq \varepsilon_{j2}$. Thus, committing to a policy may be far from optimal. That is the question studied by Atal (2016). Atal uses choice and utilization data from Chile to estimate preferences for different health insurance plans, taking into account preferences for specific hospital networks and how those preferences change as the health status of the insuree changes. Using the estimated preferences and health-state transitions, he assesses the cost of being stuck with a given hospital network. He finds that consumer would be willing to pay an additional 13% in premiums to avoid the lock-in.

How are optimal contracts affected by changing needs? When tastes change, $\varepsilon_{j1} \neq \varepsilon_{j2}$, some consumer regret their commitment to a policy and they may prefer to lapse, rendering the benefits of front-loading more limited. On the other hand, as in Section 5.2, the front-loading necessary for the contract to break even ex ante is more limited, since some bad risks will lapse as well. It is not clear that the gains from dynamic contracting are compromised by changing needs.

Interestingly, Atal (2016) reports that most of the lock-in in the Chilean case comes from fear of future denials, rather than from front-loading. Namely, consumers in good health remain insured – even, perhaps temporarily over insured – because lapsing entails the risk of not being able to return to the private market. Good draws could purchase cheaper

insurance, or switch to the public system (which is cheaper and lower quality), but they may be later denied coverage once they develop a health condition.

5.5 Aggregate Risks

So far we have considered insurance against idiosyncratic risks. In some markets, consumers suffer aggregate risks. For example, borrowers suffer a common risk associated with the evolution of the interest rate. Risk-neutral lenders can offer fixed rates to insure the risk of rising interest rates. Fixed rates, or at least temporary locks, are common in mortgages.

Learning is symmetric in mortgage markets: everybody observes the evolution of the interest rate and commitment is unilateral, since borrowers can pre-pay their balances, typically without penalty. Lenders suffer from dynamic selection. When the interest rate declines, borrowers refinance. If interest rates increase lenders are stuck, unable to shift their capital to higher-yield opportunities. Front-loading in the form of points attenuates dynamic selection.

The problem with aggregate risk and symmetric learning has been considered in the context of employment by Beaudry and DiNardo (1991). The state of the economy is likely to affect workers' productivity, making employers willing to pay more for workers when the economy is in good shape. Absent long-term contracts, workers would suffer uncertainty in wages due to the business cycle. With long-term contracts, risk-neutral employers are expected to insure workers, at least partially, against productivity shocks.

Using individual data from the Current Population Survey and the Panel Study of Income Dynamics, Beaudry and DiNardo (1991) study the link between wages and the state of the economy, specifically unemployment. Several contractual arrangements are considered. Under spot contracts, wages should depend on contemporaneous unemployment. Under two-sided commitment, wages are expected to be fixed at a level determined by unemployment at the time the worker was hired. Under one-sided commitment, wages should depend on the lowest unemployment rate realized since the worker was hired. Namely, wages are downward rigid, and get upgraded every time the worker finds better opportunities on the market.

In line with the one-sided commitment model's predictions, wages are found to be negatively related to the lowest unemployment rate realized since the worker started his or her present job. Interestingly, controlling for labor market conditions since the employment began, the contemporaneous unemployment rate no longer significantly affects wages.

6 Asymmetric Learning

In many markets, outsiders have inferior information; learning is not symmetric. Current employers arguably observe workers' performance better than competitors. Similarly, insurers observe their customers' accident histories while competitors may not.

Interestingly, the information structure has stark implications on equilibrium contracts and market performance. These distinct implications permit us to identify the kind of learning present in the market.

In some markets, the information structure is a matter of choice – for example, by forcing insurers to share (or forbidding them from sharing) customers' claims histories. The European Commission debated the legality of sharing accident information among automobile insurers in France and Belgium, which would have switched the market from one with asymmetric learning into one with symmetric learning. De Garidel-Thoron (2005) studies the welfare implications of making accident information public.

To ease exposition, let's call the insurers in the spot market "competitors," and the incumbent insurer who sold coverage in the first period "the insurer."

6.1 Theory

To study the role of dynamic contracts under asymmetric learning, we modify the model in de Garidel-Thoron (2005) to align the analysis with that of Section 2.1, so that the only difference between this section and Section 2.1 is the nature of learning.

De Garidel-Thoron (2005) presents a two-period model with symmetric priors, asymmetric learning, and one sided commitment. Ex ante identical risk-averse consumers with income y in both periods can suffer a financial loss m .

Unlike de Garidel-Thoron, for simplicity and consistency with Section 2.1, we assume there is no loss in the first period. The probabilities of a second-period loss are p_A and p_N , depending on the signal received between periods. The signal $\lambda = \{A, N\}$ can be interpreted as having a first-period accident or not (a costless accident, since we assume no first-period loss, for simplicity), so that $p_A > p_N$.⁹ Unlike Section 2.1, we capture asymmetric learning by assuming the signal is exclusively observed by the insuree and the insurer. We allow for changing income, $y_1 \neq y_2$.

Let's consider the two benchmarks presented in Section 2.1, but now with asymmetric learning in which competitors do not observe λ .

⁹It seems strange we assume there is no first-period risk, despite signals being observed. We do so for simplicity and to mimic the analysis in Section 2.1.

6.1.1 Benchmark I: Two-Sided Commitment

Because of commitment, insurance is sold before the signal is observed. Competitive premiums break even ex ante, so that $p_1 + p_2 = E_\lambda(E(m|\lambda))$. Premiums are independent of λ : Both event risk and premium risk are fully insured, equating $u'(c)$ across period-2 states.

Premiums p_1 and p_2 are timed to smooth consumption, equating $u'(c)$ across periods as well. The allocation is first best, as it is under two-sided commitment with symmetric learning. The asymmetry of information is immaterial for the allocation because parties commit before the information is revealed.

6.1.2 Benchmark II: No Commitment

Absent commitment, the spot market plays an active role. When the spot market opens, competitors are at an informational disadvantage relative to the insurer who observed λ . This asymmetry is the key shaping contracts under no commitment as well as under unilateral commitment (discussed next).

Because of the informational disadvantage competitors must offer Rothschild-Stiglitz type contracts to separate buyers. Notice separation is necessary in equilibrium. If a pooling contract was offered, the insurer would retain N and let A go, making the pooling contract unprofitable.

The period-2 separating spot contracts screen A by offering partial coverage to N . A gets full insurance at actuarially fair premiums, while N 's premium is set so that A is indifferent between the two contracts.

The insurer, having observed λ can offer full coverage to both types of customers. Having full information means it does not need to screen buyers.

The insurer enjoys an informational rent from N , whose outside option involves partial coverage. Although the market is competitive, unable to show her type in the spot market, N is locked-in to the insurer, paying more than actuarially fair premiums.

In the first period, insurers compete to attract customers, charging premiums below cost in anticipation of future rents. Insurers invest in customers to enjoy the second-period rent from those who end up becoming good drivers, N . First-period competition dissipates the rent ex ante.

Interestingly, despite of consumers' lack of commitment, N is stuck with the insurer. The lock-in enables the transfer of resources away from the most fortunate state – the state with the lowest marginal utility of consumption – to the first period. Asymmetric learning creates consumer commitment in the state in which it is valuable to commit to remain with the insurer, namely, in the good state.

By taking resources away from state N , the lock-in lowers consumption variance. Thus, as shown by de Garidel-Thoron (2005), welfare is higher under asymmetric learning than under symmetric learning when parties cannot commit.

6.1.3 One-Sided Commitment

Similar forces are at play under unilateral contracts. N finds herself locked-in with the insurer, who enjoys a second-period rent. The main distinction is that the insurer is now able to commit to the future terms of the contract, and promise to transfer resources to state A . The combination of lock-in to extract resources away from N and insurer commitment to send resources to state A improves welfare, relative to no commitment, as well as relative to one sided commitment with symmetric learning.

More specifically, second-period competitors offer separating contracts with full insurance for A and partial coverage for N . The terms offered differ from those under no commitment, since the inside option of each type depends on the long-term contract offered by the insurer in period 1.

The long-term contract offers full insurance to both A and N . As in the non commitment case, lock-in generates an informational rent from N . Competition in the first period dissipates the rent in the form of a transfer to period 1 and to state A in the second period, equating marginal utility across period 1 and state A in period 2. Absent commitment, such a promise is not feasible.

Relative to symmetric learning, N finds herself locked-in; committed to stay in the state she is tempted to lapse, enabling the transfer of resources to period 1 and state A , reducing consumption variance.

6.2 Policy Implications

What are the welfare consequences of disclosing accident information? As noted previously, the question was debated by the European Commission in the context of automobile insurance. On the plus side, symmetric information eases switching, perhaps increasing second-period competition. On the other hand, locking-in good types enhances commitment (as in Crocker-Moran (2003)).

De Garidel-Thoron (2005) shows that the extra commitment induced by asymmetric learning locks-in type N drivers, which reduces consumption variance, and is thus welfare improving. He does so by comparing market performance in the situation described in Section 2.4 under symmetric learning with the one with asymmetric learning in Section 6.1.3. In this stylized two-period world, banning information sharing is preferable.

6.3 Testable Implications

The two environments, symmetric and asymmetric learning, deliver starkly opposite predictions on premium profiles, lapsation, and selection.

Under asymmetric learning, good draws are unable to show their type and do not generate attractive offers. The insurer enjoys an informational rent, which is competed away through first-period premiums below actuarial costs, known in the literature as lowballing. Moreover, since good types are locked-in, we expect advantageous selection. Under symmetric learning, in contrast, we expect the good draws to get the best outsider offers and lapse from the contract, leading to adverse selection. These predictions were tested by D’Arcy and Doherty (1990) and Cohen (2012).

There is one important caveat to these theoretical results.”. While the symmetric learning predictions are derived from multi-period models (Harris and Holmstrom (1982), Handel et al. (2016)), we are not aware of any work showing the robustness of de Garidel-Thoron’s (2005) predictions to horizons longer than two periods. While it is not clear if predictions would change, more work is needed to understand asymmetric learning in multiple periods.

6.4 Labor Markets

Employment relations are natural situations in which asymmetric learning may arise. Performance might be monitored within the confines of the workplace, but is not easy to follow by outsiders. Not surprisingly, the earliest interest in asymmetric learning came from the work of Labor Economists.

Greenwald (1986) studied the frictions created by the asymmetry of information between current and potential employers. Similar to the analysis above, outsiders expect movers to be of lower quality and thus offer low wages, so that good draws remain in their current jobs, and the initial employer enjoys an informational rent.

Waldman (1984) and Ricart i Costa (1988) address a similar situation to Greenwald’s, but assume that while worker performance is not observed by competitors, rank and promotions are. In equilibrium, promotions are delayed to prevent dilution of the informational rent enjoyed by current employers.

6.5 Insurance

Work on insurance includes Kunreuther and Pauly (1985) and D’Arcy and Doherty (1990). These papers, unlike de Garidel-Thoron (2005), assume initial asymmetry in information between insurers and drivers. Nevertheless, predictions are similar to those derived above. The eventual asymmetry vis-a-vis competitors creates customer lock-in: Good draws cannot

prove their type in order to get better quotes from competitors. Locked-in drivers become profitable customers for their current insurers. These models predict lowballing: competition drives initial premiums below actuarial costs, due to later profits once good drivers become locked-in.

6.6 Credit Markets

Starting with Sharpe (1990), the banking literature has looked at the advantage banks have in ongoing relationships with repeat borrowers. Observing their performance over time, they enjoy an informational rent over new lenders. Rajan (1992) studies the impact banks' bargaining power, due to lock-in, has on firms' portfolio choice of borrowing sources. He shows that differential treatment due to lock-in can lead to inefficient capital allocation. Dell'Ariccia, Friedman, and Marquez (1999) study the impact of superior information by the incumbent lender on entry, which blockades entrants.

6.7 Evidence

Katz and Gibbons' (1991) is the first empirical paper to use labor data to assess the presence of asymmetric information between current and potential employers. They compare wages and unemployment duration under two different scenarios: when individuals are individually laid off versus layoffs due to plant closings. Plant closing convey no information about the individual worker's type, while an individual lay-off might be attributed to low performance. The authors find shorter unemployment spells and higher subsequent wages for workers displaced by a plant closing than those individually laid off. While the findings do not directly speak to issues of dynamic contracting, they do attest to the association of switchers with poor performance and worker lock-in.

D'Arcy and Doherty (1990) present evidence of lowballing in the US automobile insurance market. Loss ratios (losses over premiums) in their sample of seven insurers decline in the age of the policy, consistent with lowballing and lock-in of good types.

Cohen (2012) presents evidence from an automobile insurer in Israel. There is no information sharing among insurers in Israel, so learning is asymmetric. Cohen finds that drivers with a good record (good types) are less likely to lapse, and over time good types become more profitable. More precisely, drivers get a discount for their good driving record, but the discount does not fully account for the loss differential. The cost reduction is not fully passed to the insuree, which is consistent with the insurer having an informational advantage. Namely, good drivers are unable to fully convey their type to competitors.

In D'Arcy and Doherty (1990) and Cohen (2012), as well as in the model presented by Kunreuther and Pauly (1985), there is initial asymmetric information between the driver

and the insurer. While the initial asymmetry departs from the basic set-up presented above as well as from that of de Garidel-Thoron, predictions are to a large extent similar. We can think of the initial informational asymmetry as arising from a prior driving history under a different insurer – that is, as if the initial period in those models is the second stage of de Garidel-Thoron. A multi-period version of the latter might provide additional insights into the function of markets with asymmetric learning.

7 Conclusions

This chapter surveys the theory and evidence on contracting under learning à la Harris and Holmstrom (1982) and its variations. Models of learning and imperfect commitment are useful for understanding dynamic selection and reclassification risk. Imperfect commitment can be the source of adverse selection, even when information is symmetric. We draw from diverse areas of economics to show the relevance of the theory.

The empirical literature has looked at testable implications on selection and on optimal contracts, and more recently has estimated the welfare loss from lack of commitment to be substantial. The theory offers policy prescriptions on how to overcome the market distortions associated with limited commitment, mainly to prevent the lapsation of good draws.

References

- [1] Abramitzky, R.(2008) “The Limits of Equality: Evidence from the Israeli Kibbutz,” *Quarterly Journal of Economics* 123(3): 1111-1159.
- [2] Akerlof, G. (1970) “The Market for “Lemons”: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics* 84(3): 488-500.
- [3] Atal, J. P. "Lock-In in Dynamic Health Insurance Contracts: Evidence from Chile."
- [4] Beaudry, Paul and John DiNardo (1991) "The Effect of Implicit Contracts on the Movement of Wages Over the Business Cycle: Evidence from Micro Data *Journal of Political Economy*, Vol. 99, No. 4, pp. 665-688.
- [5] Browne, M. and A. Hoffman. (2013) "One-Sided Commitment in Dynamic Insurance Contracts: Evidence from Private Health Insurance in Germany" *Journal of Risk and Uncertainty*, February 2013, Volume 46, Issue 1, pp 81-112.

- [6] Cardon, J. and I. Hendel (2001), "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey," *RAND Journal of Economics* 32: 408-27.
- [7] Cochrane, J. (1995), "Time-Consistent Health Insurance," *Journal of Political Economy* 103: 445-73.
- [8] Cohen, A. (2005) "Asymmetric Information and Learning in the Automobile Insurance Market," *Review of Economics and Statistics*, 87: 197-207.
- [9] Cohen, A. (2012) "Asymmetric Learning in Repeated Contracting: An Empirical Study," *Review of Economics and Statistics*, Vol. 94, No. 2, Pages 419-432.
- [10] Crocker, K. and J. Moran (2003). "Contracting with limited commitment: evidence from employment-based health insurance contracts." *RAND Journal of Economics*, 694-718.
- [11] Currie, Janet and Brigitte Madrian, "Health, Health Insurance and the Labor Market" *Handbook of Health Economics*, Volume 3, Part C, 1999, Pages 3309–3416.
- [12] Dafny, Leemore, Igal Hendel and Nathan Wilson (2015) "Narrow Networks on the Health Insurance Exchanges: What Do They Look Like and How Do They Affect Pricing? A Case Study of Texas," *American Economic Review P&P* Vol 105, No 5, 110-114.
- [13] Daily, Glenn, *The Individuals Investor's Guide to Low-Load Insurance Products* (Chicago, IL: International Publishing Corporation, 1989).
- [14] Daily, G., I. Hendel, A. Lizzeri (2008). "Does the Secondary Life Insurance Market Threaten Dynamic Insurance?" *American Economic Review*, Vol. 98, No. 2, 151-156.
- [15] D'Arcy, Stephen P., and Doherty, Neil A. (1990). "Adverse Selection, Private Information, and Lowballing in Insurance Markets." *Journal of Business*, 63, April, 145-64.
- [16] DeGaridel-Thoron, T. (2005). "Welfare-Improving Asymmetric Information in Dynamic Insurance Markets," *Journal of Political Economy*, vol. 113, issue 1, pages 121-150.
- [17] Dell'Araccia, C., E. Friedman, and R. Maquez, (1999) "Adverse Selection as a Barrier to Entry in the Banking Industry, *RAND Journal of Economics* 30 (1999), 515-534.
- [18] Finkelstein, A, K. McGarry and A. Sufi. (2005) "Dynamic Inefficiencies in Insurance Markets: Evidence from Long-Term Care Insurance" *The American Economic Review, Papers and Proceedings*, 95 (2): 224-228.

- [19] Gavazza, Alessandro and Hanming Fang, (2010). "Dynamic Inefficiencies in an Employment-Based Health-Insurance System: Theory and Evidence," *American Economic Review*, Vol. 101, No. 7, pp. 3047-3077.
- [20] Gibbons, R. and L.F. Katz. (1991) "Layoffs and Lemons" *Journal of Labor Economics*, Vol. 9, No. 4, pp. 351-380.
- [21] Greenwald, Bruce, (1986) "Adverse Selection in the Labor Market," *Review of Economic Studies* 53, 325-347.
- [22] Hall, R. (1978) "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*. 86(6) 971-987.
- [23] Handel, B. (2013), "Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 103 (7), 2643-2682.
- [24] Handel, B., I. Hendel and M. Whinston (2015), "Equilibria in Health Exchanges: Adverse Selection vs. Reclassification Risk" *Econometrica*, Volume 83, Issue 4, 1261–1313.
- [25] Handel, B., I. Hendel and M. Whinston (2016), "The Welfare Gains from long-term Contracts: Health Insurance"
- [26] Harris, M. and B. Holmstrom, "A Theory of Wage Dynamics," *Review of Economic Studies*, XL (1982), 315-333.
- [27] Hendel, I. and A. Lizzeri (2003), "The Role of Commitment in Dynamic Contracts: Evidence from Life Insurance," *Quarterly Journal of Economics* 118: 299-327.
- [28] Herring, B. and M. Pauly (2006), "Incentive-Compatible Guaranteed Renewable Health Insurance Premiums," *Journal of Health Economics* 25: 395-417.
- [29] Holmstrom B. (1983) "Equilibrium Long-Term Labor Contracts." *Quarterly Journal of Economics*, Vol. 98, Supplement (1983), pp. 23-54
- [30] Koch, T. (2011), "One Pool to Insure Them All? Age, Risk, and the Price(s) of Medical Insurance," UC-Santa Barbara Working Paper.
- [31] Kunreuther, Howard, and Mark Pauly (1985). "Market Equilibrium with Private Knowledge: An Insurance Example." *Journal Public Economics* 26 April, 269-88.
- [32] Ligon, E., J. Thomas, and T. Worrall, (2002). "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies" *Review of Economic Studies* 69, 209-244.

- [33] Patel, V. and M. Pauly (2002). "Guaranteed Renewability And The Problem Of Risk Variation In Individual Health Insurance Markets," *Health Affairs*.
- [34] Pauly, M., H. Kunreuther and R. Hirth (1995), "Guaranteed Renewability in Insurance." *Journal of Risk and Uncertainty*, 10: 143-156.
- [35] Pauly, M., A. Percy, and B. Herring (1999), "Individual vs. Job-Based Health Insurance: Weighing the Pros and Cons," *Health Affairs* 18(6): 28-44.
- [36] Rajan, Raghuram, (1992) "Insiders and Outsiders: The Choice between Informed and Arm's Length Debt," *Journal of Finance* 47, 1367-1400.
- [37] Ricart I Costa, Joan, (1988). "Managerial Task Assignment and Promotions," *Econometrica* 56 (2), 449-66.
- [38] Sharpe, Steven, "Asymmetric Information, Bank Lending and Implicit Contracts: A Stylized Model of Customer Relationships," *Journal of Finance* 45 (1990), 1069-1087.