It’s all about speed and costs:

The impact of digital technology on the insurance market structure

Martin Eling, Ruo Jia, Casey Rothschild[1]

(Work in progress; please do not cite or distribute; this version: April 9, 2018)

Abstract Digital technology is costly, but it allows insurers to more quickly learn about the risk type of their policyholders. We study the implications of this speed-versus-cost tradeoff for equilibrium pricing and coverage decisions in an insurance market featuring adverse selection. In particular, we develop a theoretical model of dynamic competitive equilibrium featuring individuals who differ in their privately known risk types, and a large number of two types of insurers: conventional insurers and “tech” insurers who employ digital technologies. We consider three distinct dynamic equilibrium concepts: a finite horizon structure with foresight, an infinite horizon “overlapping generations” structure, and an infinite horizon myopic structure. Equilibrium in each setting features a sorting of low-risk types into tech firms and high-risk types into conventional firms. Depending on the setting, however, the equilibrium tech-firm market share may negatively or positively correlate with the learning speed of conventional insurers.

Keywords Digitalization, Speed of learning, Asymmetric learning, Dynamic equilibrium, Adverse selection, Market structure

JEL Classification G22, D40, O33

[1] Martin Eling (martin.eling@unisg.ch) is affiliated with the Institute of Insurance Economics, School of Finance, University of St. Gallen; Ruo Jia (ruo.jia@pku.edu.cn) is affiliated with the Department of Risk Management and Insurance, School of Economics, Peking University; Casey G. Rothschild (crothsch@wellesley.edu) is affiliated with the Department of Economics, Wellesley College.

1 Introduction

Digital technology generates data, enabling insurance firms to more effectively and efficiently understand the risk they underwrite. A telematics device installed in a car, for example, provides an insurer with much more risk-relevant information than the conventional underwriting information. Moreover, it opens up opportunities for real-time monitoring and risk assessment. These digital technologies can therefore—at least in some contexts—be expected to mitigate information asymmetries, both of the adverse selection and moral hazard varieties. Qualitatively speaking, digital technologies such as telematics devices will operate in much the same way as "old fashioned" insurer learning based on claims experience and policyholder disclosure (Nilssen, 2000; de Garidel-Thoron, 2005; Cohen, 2012; Kofman and Nini, 2013)—learning processes which have been shown to be effective tools to mitigate adverse selection (Eling, Jia, and Yao, 2017) and moral hazard (Dionne, Michaud, and Dahchour, 2013). Quantitatively speaking, however, there are two important differences: speed and cost. As shown empirically in Eling, Jia, and Yao (2017), learning by conventional technologies is slow—and learning via digital technologies is fast or even instantaneous (Gemmo, Browne, and Gründl, 2017). Since telematics and other digital technologies involve installation and potentially data processing costs, they are—at least for the moment—more costly than conventional learning techniques.[2]

This paper studies the implications of this speed versus cost tradeoff for equilibrium pricing and coverage decisions in an insurance market featuring adverse selection. In particular, we develop a theoretical model of dynamic competitive equilibrium featuring a continuum of individuals who differ in their risk types and featuring a large number of two types of insurers: those using conventional underwriting techniques only (or "conv" insurers) and those "tech" insurers who employ digital learning technologies. Tech insurers have an advantage because they can rapidly learn the risk type of any given individual and thus can effectively price based on each risk; but installing (and/or maintaining) their technology is costly. Conventional insurers cannot immediately observe risk types, and hence cannot employ risk-based pricing and face adverse selection, but we do allow them to learn over time about the risk type of their customers.

Within this broad modeling framework, we consider three distinct dynamic equilibrium concepts. One is a "myopic" (behavioral) model where customers choose firms based on current prices. In this setting, conventional insurers understand their risk pool based on backward-looking experience and update prices with a lag as they "bleed" customers to tech insurers. This is best seen as a heuristic, but realistic model to describe the learning process of conventional insurers and the market evolution in response to the introduction of new technologies. The other two models consider firms and individuals

[2] The costs for implementing and maintaining new technology can be substantial. For example, costs for implementing a telematics device in a car are around 100 USD in 2018. The subsequent costs to maintain the technology can also be substantial and might be both of variable and fixed nature, depending on how the technology is operated (i.e. self-operated versus outsourced to a technology partner). Our analyses will also yield equilibrium implications, if the technology costs will go substantially down in the future.

to be fully rational and forward-looking. One is a finite horizon (FH) two-period model where conventional firms compete with each other for new customers only in the first period. The other is a stationary, infinite horizon, overlapping generations (OLG) model where conventional firms compete in each period for new customers. The key assumption in this model is that insurers cannot distinguish customers from the new generation from older customers who are fleeing other insurers who have learned that they are a high-risk type and raised the premiums accordingly.[3]

We show that, in all three settings, equilibria will feature a sorting of low-risk types into tech firms and high-risk types into conventional firms. We then consider comparative statics of equilibrium with respect to two key underlying parameters: the cost of digital technology and the speed at which conventional insurers learn the risk types of their buyers. Across all three settings, we establish the intuitive result that lowering the cost of the digital technology will raise the equilibrium market share of tech firms. Interestingly, however, the effect of the speed of learning on the tech market share depends on the dynamic model. In the FH model, faster learning leads to a higher market share for conventional firms. Intuitively, this is because incumbent firms can earn information rents once they learn about their clients, so the faster they learn the more they will "lowball" prices for new customers (Kunreuther and Pauly, 1985; Nilssen, 2000). In the OLG model, the same lowballing incentive is present, but it is counteracted by the fact that individuals who reveal themselves to be high risk will be dumped (via high prices) by their insurers. They will then return to the conventional market, "polluting" the risk pool for conventional insurers writing new business and hence raising the equilibrium price of conventional insurance. This "polluting" effect *always* dominates, so in the OLG model, faster learning by conventional firms actually *raises* the equilibrium share of tech insurers. In the myopic model, we show that the effect of faster learning of conventional insurers on the market share of tech insurers is ambiguous.

Our paper fits in a thick literature on asymmetric learning and dynamic pricing. For example, our model of conventional firm learning and consequent "lowballing" is consistent with the asymmetric learning model used in Kunreuther and Pauly (1985), Nilssen (2000) and de Garidel-Thoron (2005) in which the incumbent insurer knows the risk type of the policyholder, but competing insurers do not (and in contrast to the symmetric learning model used e.g. in Watt and Vazquez (1997) and Hendel (2016)) wherein all insurers have the same information in all periods). The key difference in our paper is the simultaneous presence of tech insurers who effectively know the risk type in real time and have no incumbency advantage.

Our work also contributes to the ongoing discussion on the impact of digitalization on the equilibrium of insurance market (Filipova-Neumann and Welzel, 2010; Gemmo et al., 2017). Gemmo et al. (2017) use a one-period menu contract framework to analyze the trade-off between the reduction of information asymmetry and the willingness to share private information (transparency aversion).

[3] Both cases can be motivated empirically, given that in many countries for certain types of insurance products information exchange platforms exists, while for other products such platforms do not exists.

While Gemmo et al. (2017) assume perfect observation of risk type at contract inception, Filipova-Neumann and Welzel (2010) analyze a similar setting where risk type is revealed only after an accident. Both studies find welfare increases greater or equal to zero in a separating equilibrium setting. We assume risks do not have the problem of transparency aversion and instead focus on the trade-off between the speed of learning and the digital technology cost. To the best of our knowledge, we are the first to analyze the impact of learning speed on the market structure at the competitive dynamic equilibrium.[4]

The rest of the paper is structured as follows. In Section 2, we introduce our theoretical framework by summarizing the common features and qualitative differences of the three models. In Section 3, we describe the three models, their (quasi-)equilibriums, and the corresponding propositions at the (quasi-) equilibriums. We conclude in Section 4.

2 Theoretical Framework

2.1 Common features of all three models

We consider three distinct models of competitive market structure. In each model, there are a large numbers of conventional and tech insurers that compete on price to sell an insurance product with homogeneous coverage to a large number of buyers who are differentiated by their loss probability. All models feature a set of periods $t = 0,1,2 \ldots, T$, where $T \geq 1$ and $T \leq \infty$. They all feature a set of risk types, indexed by the per-period risk of loss p. Losses are assumed to be independent across periods, and of constant size L out of a per-period income W. We abstract from saving, so that, absent insurance, individuals have a net consumption of $y_t^L = W - L$ or $y_t^N = W$ in the event of a loss or no loss, respectively.

Firms sell full insurance as one period contracts at a premium q, which, depending on that firm's ability to observe information, may depend on risk types of the buyer. Tech firms observe risk types and hence offer premiums $q^\tau(p)$ depending directly on the risk type. Conventional firms can distinguish between their *incumbent* insureds and other potential insureds, and hence can offer them different premiums; conventional firms may learn information about their incumbent insureds over time, in which case they can potentially offer different incumbent insureds different premiums.

An individual who buys an insurance contract at a price q in period t will have loss-independent consumption $y_t^L = y_t^N = W - q$. A conventional firm who sells such a contract to a risk type p will earn period profits $\pi^c = q - pL$. A tech firm who sells such a contract earns $\pi^\tau = q - pL - C$, where

[4] Our paper is also related to the debate on the usage of genetic information for risk calculation. A genetic test is also costly, but if shared with the insurance company decreases information asymmetry. Hoy and Ruse (2005) argue that the reduction of adverse selection and therefore the increasing efficiency is accompanied by the effect that people who are in poor health are punished twice (higher premium and health problems); some people will refuse to take a genetic test, because they are afraid that it will increase their insurance premium. Doherty and Posey (1998) find that for uninformed individuals a genetic test has a positive private value if prevention is sufficiently effective in lowering the premium, even though the information must be shared with the insurer.

$C \geq 0$ are monitoring costs, which we potentially allow to be history-dependent in order to allow for one-time installation costs as well as ongoing maintenance and monitoring costs.

Individuals are assumed to have standard risk averse von Neumann-Morgenstern preferences over consumptions within periods; we assume the utility of individuals is additively separable and time-discounted across periods, i.e., given by:

$$\sum_t \beta^t [pU(y_t^L) + (1-p)U(y_t^N)\ ].$$

We assume that U is twice differentiable, increasing, and strictly concave, and $\beta \leq 1$ is the time discount factor of individuals. Firms are risk neutral and they all discount their per-period profits over time with the same market discount rate β.

Finally, we assume that there is a large number of firms, so that competition drives expected profits down to zero on a contract-by-contract basis. As we describe below, the precise meaning of this will vary depending on whether firm beliefs are myopic or forward looking. In all cases, however, competition will drive tech firms to price at $q = pL + C$, so that they break even given the realized (and known) risk p, and this tech market will provide a static "outside option" to insurance buyers who are considering buying from a conventional firm. We assume, purely for expositional simplicity, that

$$pU(W-L) + (1-p)U(W) < U(W - pL - C)$$

for all $p \in \left[\underline{p}, \bar{p}\right]$. This ensures that all individuals prefer insuring with a (break-even pricing) tech firm to going without insurance (which is always true for sufficiently low C). We will therefore not discuss the case of opting out of the insurance market in the following.

2.2 Qualitative differences across models

We consider three families of models differ in their assumptions about *dynamics* and *rationality*. We describe these differences qualitatively here and then formalize them in the subsequent section.

The first "finite horizon" (FH) model is a two-period rational expectations model in the spirit of Nilssen (2000) and de Garidel-Thoron (2005). In this model, conventional firms compete in period 0, taking into account two things that firms in the myopic model ignore. First, they take into account the *endogenous distribution of risk types who actually buy from them*. In other words, they correctly anticipate that some buyers (in practice, the low risks) will buy from tech firms (to opt out of the market is excluded by definition) in the current period, and adjust their prices accordingly. Second, they take into account the future consequences of their current actions. Specifically, conventional insurers may asymmetrically learn between periods 0 and 1 about the risk type of the individual they insured. If so, they will earn some information quasi-rents in period 1—though these rents are limited by the competitive tech market that coexists with it. A central dynamic in this model is "lowballing" of conventional prices in period 1, in anticipation of these period 2 quasi-rents.

The second "overlapping generations" (OLG) model is an infinite horizon rational expectations model. As in the FH model, firms (and buyers) are forward-looking and correctly anticipate the equilibrium distribution of risk types who actually buy from them at each moment in time. As such, the same "lowballing" dynamic applies, wherein conventional firms make losses on initials sales in anticipation of the information quasi-rents they will earn when they learn the risk types of their incumbent insureds. A key difference from the FH model is that these quasi-rents are (also) limited by the ability of individuals to return to the conventional insurance market, where they will be indistinguishable from new entrants into the market. Indeed, a central dynamic in this model is that high risk types "cycle through" different conventional firms, sticking with their new firm until that firm learns their type and then returning to the broad pool for a new, uninformed conventional insurer.

The third "myopic" model makes stylized—but plausible—behavioral assumptions about pricing and purchasing behavior, in the spirit of Kunreuther and Pauly (1985). Individuals in this model are assumed to choose by (myopically) maximizing their current-period payoff—i.e., they choose the cheapest contract available to them today. Similarly, competition is assumed to drive *current period* profits down to zero, however, some conventional insurers may suffer from negative profits since they do not know the risk type of new customers and thus underprice them, given that the beliefs of firms are also myopic. Furthermore, the beliefs of conventional firms are assumed to be formed in a backward-looking way, based initially on the population distribution of risk types, and then, later, from the risk distribution realized at conventional firms in the preceding period. We consider a family of these myopic models, where conventional firms learn over time about the risk types of their incumbent insureds. A central dynamic in these models is an unraveling over time a la Akerlof (1970), as conventional firms realize worse and worse risks over time, raising prices and driving more individuals towards tech firms.

3. Models

3.1 Finite Horizon Rational Expectations Model

In this model, there are two periods: $t = 0,1$. A continuum of types with (cumulative) distribution $F(p)$ and continuous pdf $f(p)$ with support on $\left[\underline{p}, \bar{p}\right] \subset (0,1)$ contemplate buying insurance in each period. We make the technical assumption that $\frac{1-F(p)}{f(p)}$ is decreasing in p (the standard monotone hazard rate property which is obviously true, e.g., for a uniform distribution).

The cost C for tech firms to provide insurance is (potentially) time dependent, with cost C_0 for a first-time tech buyer, and $C_1 \leq C_0$ for an individual who was insured with a tech firm in period 0 and buys insurance from a tech firm again in period 1. We assume that *any* tech firm will have a cost of C_1 in period 1 for such a buyer; one interpretation is that there is a fixed cost of putting in equipment and then a variable cost of using it in each period—but all firms can use the same equipment once it is installed and only bear the monitoring cost.[5] This assumption ensures that tech firms are effectively competing *statically* in each of the two periods over the premium $q_t^\tau(p)$ for each potential buyer, so that, in equilibrium, $q_0^\tau(p) = q_1^{\tau,n}(p) = pL + C_0$ and $q_1^{\tau,i}(p) = pL + C_1$, where $q_1^{\tau,n}(p)$ denotes prices offered in period 1 to individuals who were not insured with tech firms in period 0 and $q_1^{\tau,i}(p)$ denotes prices for those who were.

Conventional firms cannot observe buyer type when setting prices in period 0 and thus compete over a single premium q_0^c. If they insure an individual in period 0, then, with probability α, they will learn the buyer's type before setting prices again in period 1, and will choose a price $q_1^{c,L}(p)$. With probability $1-\alpha$, they do not observe the type, and can offer a single premium $q_1^{c,U}$ to these unlearned types. We also assume, again purely for analytical simplicity, that in period 1 conventional firms cannot (or do not) offer contracts to non-incumbent insureds—so individuals who do not wish to stay with their original conventional insurer can either forgo insurance or switch to tech firms.[6]

With the goal of defining a competitive, rational expectations equilibrium for this economy, consider first the sequentially rational decisions of individuals and firms in period 1. Individuals who bought from a tech firm in period 0 choose rationally in period 1 between forgoing insurance and buying from a tech firm at their competitive price $q_1^{\tau,i}(p) = pL + C_1$. Because of our assumption on $[\underline{p}, \overline{p}]$ and C_1, they will all choose to purchase insurance.

Similarly, individuals who purchased from conventional firm in period 0 in principle face a choice among three options: foregoing insurance, switching to buy from a tech firm at the premium $q_1^{\tau,n}(p) = pL + C_0$, or remaining with their conventional firm at the offered price of $q_1^{c,L}(p)$ or $q_1^{c,U}$, depending on whether their type was learned or not. However, as we will show below they will never choose to forego insurance in equilibrium, and they effectively choose between switching to a tech firm and remaining with their incumbent firm. We assume that they remain with their conventional firm if they are indifferent. In other words, they choose to remain precisely when $q_1^{c,L}(p) \le q_1^{\tau,n}(p)$ if their type was learned and when $q_1^{c,U} \le q_1^{\tau,n}(p)$ if it was not.

In light of the optimizing behavior by individuals, an incumbent firm in period 1 who has learned the type p of a given insured maximizes profits by choosing $q_1^{c,L}(p) = pL + C_0$, i.e., the highest price at which they will retain their customer—which is a profitable price for the incumbent firm since $C_0 > 0$. The profit-maximizing price for their unlearned customers, namely $q_1^{c,U}$, is less obvious. At any price

[5] Allowing an incumbency advantage for tech firms, so that new firms have to pay $C_0 > C_1$ for one period does not change any of the qualitative results, though it is mildly more expositionally cumbersome.

[6] This assumption is intuitively plausible, as incumbent conventional firms have an informational advantage over other conventional firms for these individuals: even though they do observe their type, they *know* that they have not been revealed to be a high risk types. Non-incumbent firms would attract risks of both the unrevealed types and the high-risk revealed types who the incumbent firms would be happy to offload. For this reason, it is not hard to show that non-incumbent conventional firms will never make positive sales in equilibrium in which such sales are allowed. But the presence of these potential competitive firms makes the analysis of the equilibrium less transparent. Hence our technical assumption.

$q_1^{c,U}$, they will sell to those types for whom $q_1^{c,U} \leq q_1^{\tau,n}(p) = pL + C_0$, i.e., to any customers with $p \in \left[\frac{(q_1^{c,U} - C_0)}{L}, \overline{p}\right]$. Any incumbent unlearned types with $p < \frac{(q_1^{c,U} - C_0)}{L}$ will depart for tech firms. Lowering the premium thus involves a tradeoff: on the one hand, the firm retains additional, profitable lower risks. But, on the other hand they lower their per-unit profits on the higher-risk customers who were already planning to purchase at the higher premium. The assumption that $\frac{1-F(p)}{f(p)}$ is decreasing in p will ensure that there is a unique profit maximizing price $q_1^{c,U*}(\hat{p})$ whenever the set of period-0 conventional firm purchasers has the interval form $[\hat{p}, \overline{p}]$ (as it will in equilibrium), and that there will also be a cutoff individual $p^*(\hat{p}) \in [\hat{p}, \overline{p})$ for whom $q_1^{c,U*}(\hat{p}) \equiv p^*(\hat{p})L + C_0$ who is indifferent between staying with the conventional insurer and switching to a tech firm.

In period 0, we aim to describe a competitive equilibrium with perfect foresight. Intuitively, individuals will choose between tech and conventional firms in period 0, fully anticipating what pricing will be in period 1 (and hence consistent with the preceding). Conventional firms will compete over premiums and drive that premium to the level at which such firms earn zero lifetime profits. Since conventional firms earn positive profits in period 1, this will imply "lowballing".

To formalize this basic intuition, note first that individuals who choose a tech firm pay $q_0^{\tau}(p) = pL + C_0$ in period 0 and rationally anticipate paying a premium $q_1^{\tau,i}(p) = pL + C_1$ in period 1. Individuals who choose a conventional firm pay some premium q_0^C and then rationally anticipate their conventional firm's period 1 pricing. They thus anticipate paying $q_1^{c,L}(p) = pL + C_0$ when, with probability α, their type is learned. With probability $1 - \alpha$, their type is not learned. If $p < \frac{(q_1^{c,U*} - C_0)}{L}$, they will switch to a tech firm in period 1 and again pay $q_1^{\tau,n}(p) = pL + C_0$. Higher p types will instead anticipate staying with their conventional firm and paying some rationally anticipated $q_1^{c,U*}$.

It is intuitive that in period 0 purchase decisions will take a cutoff form: higher p individuals will buy from conventional firms and lower p individuals from tech firms.[7] The cutoff individual $\hat{p}$ who buys from a conventional firm will *always* buy insurance at a price $\hat{p}L + C_0$ in period 1. This follows from period 1 profit maximization, if their type is learned. If it is not learned, it follows from the fact that $q_1^{c,U} \geq \hat{p}L + C_0$ as, at $q_1^{c,U} = \hat{p}L + C_0$, the firm will retain *all* of their incumbent customers, and there is no "additional sales" benefit to lower prices. If interior, the cutoff type $\hat{p}$ is thus indifferent between buying from a conventional firm and paying q_0^c then $\hat{p}L + C_0$ or buying from a tech firm and paying $\hat{p}L + C_0$ and then $\hat{p}L + C_1$. This implies a cutoff type $\hat{p}(q_0^c; C_0, C_1)$ which is the maximum of $\underline{p}$ (if all buy from conventional firms) and the solution to:

[7] It follows formally from two simple observations given any fixed prices. First, the (non-stochastic) utility difference between buying from a tech firm in both periods and buying from a conventional firm in period 1 and then a tech firm in period 2 is $U(W - pL - C_1) - U(W - q_0^C)$, which is decreasing in p. Second, the period 1 option value of staying at a conventional firm is increasing in p.

$$u(W - q_0^c) + \beta u(W - \hat{p}(q_0^c)L - C_0) = u(W - \hat{p}(q_0^c)L - C_0) + \beta u(W - \hat{p}(q_0^c)L - C_1). \quad (\boldsymbol{Eqn.\hat{p}})$$

Notice that this solution is decreasing in both C_0 and C_1.

If the market price is q_0^c, then the lifetime profits of all conventional firms will be:

$$\pi^c(q_0^c) = \int_{\hat{p}(q_0^c)}^{\overline{p}} (q_0^c - Lp)f(p)dp + \beta\alpha C_0 \left(1 - F\big(\hat{p}(q_0^c)\big)\right) + \beta(1-\alpha)\int_{p^*\left(\hat{p}(q_0^c)\right)}^{\overline{p}} (q_1^{cU*}\big(\hat{p}(q_0^c)\big) - Lp)f(p)dp.$$

The first term is the period-0 profits. The second term is the profits in period 1 for individuals whose types have been learned (i.e., C_0, since the profit-maximizing price is $q^{c,L}(p) = pL + C_0 = q^{\tau n}(p)$). The final term is the profit from unlearned types. The equilibrium price q_0^c is determined by a zero lifetime profit condition, $\pi^c(q_0^c) = 0$. We say that a zero-profit equilibrium is *locally stable* if $\pi^c(q_0^c)$ is increasing in q_0^c (otherwise, a small decrease in premium by a single firm will be profitable).

We gather the preceding reasoning into the following definition of an FH-equilibrium.

Definition: Equilibrium in the FH Market

An equilibrium is a set of prices $q_0^c, q_1^{c,U}, q_1^{c,L}(p), q_0^\tau(p), q_1^{\tau,n}(p), q_1^{\tau,i}(p)$, and a pair of cutoff types p_0 and $p_1 \geq p_0$ such that:

(i) Tech firm competition: $q_0^\tau(p) = q_1^{\tau,n}(p) = pL + C_0$. $q_1^{\tau,i}(p) = pL + C_1$.

(ii) Conventional firm period 1 profit maximization for learned types: $q_1^{c,L}(p) = pL + C_0$.

(iii) Conventional firm period 1 profit maximization for unlearned types:
$$q_1^{c,U} = \arg\max_{\mathrm{q}} \int_{\tilde{p}(q)}^{\bar{p}} (q - pL)f(p)dp\,, \quad \text{where} \quad \tilde{p}(q) = \max\left\{p_0, \frac{q - C_0}{L}\right\}.$$

(iv) Period-0 conventional firm competition and stability: $\pi^c(q_0^c) = 0$ and $\frac{d\pi^c(q_0^c)}{dq_0^c} > 0$.

(v) Individual optimization: $p_1 = \frac{\left(q_1^{c,U} - C_0\right)}{L}$ and $p_0 = \max\left\{\underline{p}, \hat{p}(q_0^c)\right\}$, where $\hat{p}(q_0^c)$ is the solution to equation ($\boldsymbol{Eqn.\hat{p}}$) above. Individuals with $p < p_t$ $(t = 0,1)$ purchase from tech firms while those with $p > p_t$ purchase from conventional firms.

Equilibrium depends on the exogenous cost parameters C_0, C_1 and the speed of learning α. The following proposition is our main result for the FH model. It establishes formally that the equilibrium cutoffs are decreasing in α, C_0, and C_1. In other words, when the speed of learning or the cost of digital technology increase, the market shares of conventional firms in both periods increase. The formal proof is in the appendix.

Proposition FH: In any locally smooth family of equilibria:

(1) The equilibrium cutoff $p_0(\alpha, C_0, C_1)$ is (weakly) decreasing in α, C_0 and C_1, strictly so if $p_0 \neq \underline{p}$.

(2) The equilibrium cutoff $p_1(\alpha, C_0, C_1)$ is (weakly) decreasing in α, C_0 and C_1. It is strictly decreasing in C_0 if $p_0 \neq p_1$.

3.2 Infinite Horizon OLG Model

In this model, there is an infinite number of periods. There is a constant mass of potential insurance buyers in each period, which we normalize to 1. These potential buyers (individuals) make insurance decisions over one-period contracts at the beginning of the period, then losses are realized and payouts are made for covered losses. Before the next period starts, each individual has an independent probability $(1-\eta) \in [0,1)$ of "exiting" the market (e.g., dying), and a mass $(1-\eta)$ of new individuals enter the market (are "born").

Otherwise, the structure is broadly similar to the FH market. We assume that new types have the distribution $F(p)$ (with continuous pdf $f(p)$ with support on $\left[\underline{p}, \bar{p}\right] \subset (0,1)$) . The distribution of types in the market is thus constant and described by $F(p)$, as in the FH market, but here there is a constant inflow and outflow of new and old agents.

We again assume that the cost C for tech firms to provide insurance is again (potentially) time dependent, with cost C_0 for a first-time tech buyer, and $C_1 \leq C_0$ for an individual who was previously insured with any tech insurer. We allow any discount factor $\beta \in (0,1)$.

As in the FH model, we assume that if a conventional firm does not know the type of an individual to whom it sells a contract, it will learn that type with a time-independent probability α.[8]

The key difference in modeling assumptions are closely related to the finite vs infinite time horizon. With a finite horizon, a conventional firms knows in period 1 whether they insured a given individual in the preceding period, and thus would know that *any* new potential customer must have previously insured with another provider—and thus is likely to be high risk. It is thus natural to assume that individuals in period 1 will find it impossible to purchase insurance at competing, less informed, conventional firms in period 1. In contrast, conventional firms in the infinite horizon model know that there are newly born potential customers in each period, so they have an incentive to make new sales. We assume—critically—that these firms are unable to distinguish these new customers from customers who were previously insured at another firm in the infinite horizon model.

The tech firms still compete statically each period, setting $q_t^{\tau,n}(p) = pL + C_0$ for customers who have new tech customers and $q_t^{\tau,i}(p) = pL + C_1$ for old ones, and we drop the t subscript henceforth since the environment is, by construction, static over time. The ability for individuals to "return to the market" and find a new conventional firm significantly changes—and in some ways simplifies—conventional firm pricing. First, conventional firms who have incumbent types whose type they have

[8] Allowing α to be increasing over time for a conventional firm that repeatedly insures the same individual and has not yet learned their type does not change any of the subsequent analysis or qualitative results. It is significantly more notationally cumbersome, however.

learned set some price $q^{c,L}(p)$ so as to maximize profits, but the outside option for these types is no longer the same: they can go to a tech firm, as in the FH model, but now they additionally have the option of leaving for a new conventional firm. Second, conventional firms who do *not* learn the type of their incumbent buyer have no incentive to retain those buyers. Intuitively: they have learned nothing about them (other than the fact that they are high enough risk to prefer conventional to tech insurers), and because these buyers can always go to a new firm, they cannot charge a higher price or extract any rents. In effect, we can treat unlearned types as "returning" to the conventional pool and buying from a random new firm. Together, this means that the only conventional firm premiums needed to describe equilibrium in the OLG model are $q^{c,L}(p)$ and q^{c}, with the latter denoting the price in the competitive market for "new" insureds. We will be looking for a steady-state equilibrium in which these premiums are stable over time.

Towards describing that steady-state equilibrium, note that an individual can be in one of only four mutually exclusive and exhaustive states: they can be with an incumbent conventional firm who knows their type (L), or not (U), and they can have purchased from a tech firm in the past (i), or not (n). We describe the four states by Li, Ln, Ui, and Un.

Consider first an individual in state Ui. She has the option of choosing a tech firm at price $q^{\tau,i}(p)$ or from a conventional firm at price q^{c}. Contrast this with an individual in state Un. She has the option of choosing a tech firm at price $q^{\tau,n}(p)$ or from a conventional firm at price q^{c}. Because $q^{\tau,n}(p) \geq q^{\tau,i}(p)$ (with strict inequality if $C_1 < C_0$), it is obvious that if the Un individual finds it optimal to choose the tech firm, so will the Ui individual. But this means that an individual who purchases from a tech firm in the period in which she is born (into the Un state) will *always* purchase from a tech firm—that is, there will never be a type in the Li state. Conversely, an individual who purchases from a conventional firm when they are born will *never* purchase from a tech firm: she has revealed that she prefers q^{c} at a conventional firm over buying from a tech firm, and she will always have the option of doing so. Together, this means that individuals *sort* permanently into tech and conventional firms in the year they are born.

Since tech firms can observe risk type better than conventional firms, standard adverse selection intuition suggests that this initial sorting will take a "cutoff" form. To show this formally, consider again the pricing problem for a conventional firm who has learned a buyer's type and thus can earn some information quasi-rents from her. Because that buyer "sorted" into conventional firms, we know that the relevant "outside option" for this buyer is a return to the conventional market where she will pay q^{c}. The conventional firm can and therefore will charge up to q^{c} to incumbent buyers they wish to retain. (They will charge *more* to those buyers with $q^{c} < pL$, because there is no price they can profitably charge while retaining them, but without loss of generality we can imagine them charging a price of q^{c} to all and only retaining types with $q^{c} \geq pL$.) This means that conventional buyers *face a time independent insurance premium*, independent of their type—and this is true regardless of whether their type is learned or not. In light of this the payoff to "sorting" into conventional firms when born is independent of type, while the payoff to "sorting" into tech firms is decreasing in p. There is therefore some cutoff type $\hat{p}$ who is indifferent between the tech and conventional firms, with lower types sorting into tech and higher types sorting into conventional firms. In fact, we can easily characterize the cutoff $\hat{p}(q^c)$ as a function of the price q^c via the indifference condition:

$$\frac{1}{1-\beta\eta}U(W-q^c)=U\left(W-q^{\tau,n}\big(\hat{p}(q^c)\big)\right)+\frac{\beta\eta}{1-\beta\eta}\,U\left(W-q^{\tau,i}\big(\hat{p}(q^c)\big)\right).$$

The final thing needed to characterize equilibrium is the price q^c offered to conventional buyers in the "non-incumbent" pool. Towards describing the distribution of types in this pool, let $\tilde{p}(q^c)=\frac{q^c}{L}$, which is the maximum risk type that incumbent firms will retain after learning their types. The pool of non-incumbent buyers consists of *all* living individuals with $p>\tilde{p}(q^c)$ (a mass $\left(1-F\big(\tilde{p}(q^c)\big)\right)$ of them) and a fraction $\frac{1-\eta}{1-\eta(1-\alpha)}$ of individuals with $p\in[\hat{p}(q^c),\tilde{p}(q^c)]$,[9] (a mass $\left(F\big(\tilde{p}(q^c)\big)-F\big(\hat{p}(q^c)\big)\right)$. A randomly selected individual in the pool will thus be drawn from $p\in[\hat{p}(q^c),\tilde{p}(q^c)]$ with probability

$$Q_M(q^c)\equiv\frac{\frac{1-\eta}{1-\eta(1-\alpha)}\left(F\big(\tilde{p}(q^c)\big)-F\big(\hat{p}(q^c)\big)\right)}{\frac{1-\eta}{1-\eta(1-\alpha)}\left(F\big(\tilde{p}(q^c)\big)-F\big(\hat{p}(q^c)\big)\right)+1-F\big(\tilde{p}(q^c)\big)}.$$

Lifetime expected profits from such a sale are thus:

$$\pi^c(q^c)=\left(Q_M(q^c)\mathbb{E}\big[q^c-Lp|p\in[\hat{p}(q^c),\tilde{p}(q^c)]\big]+\big(1-Q_M(q^c)\big)\mathbb{E}\big[q^c-Lp|p\in[\tilde{p}(q^c),\overline{p}]\big]\right)$$
$$+\alpha\eta\beta Q_M(q^c)\frac{1}{1-\eta\beta}\left(\mathbb{E}\big[q^c-Lp|p\in[\hat{p}(q^c),\tilde{p}(q^c)]\big]\right).$$

The first line is just the premium minus the expected losses in the current period. The second term is the future profits if the type is learned (probability α), does not "die" (η) and is worth retaining ($Q_M(q^c)$). In this case, the firm earns positive information rents $q^c-L\,\mathbb{E}\big[p|p\in[\hat{p}(q^c),\tilde{p}(q^c)]\big]$ for all future periods in which the individual remains alive (the present discounted number of which is $\frac{1}{1-\eta\beta}$).

Definition: Equilibrium in the OLG Market

An equilibrium is a set of prices $q^{\tau,n}(p)$, $q^{\tau,i}(p)$, $q^{c,L}(p)$, and q^c and a cutoffs $\tilde{p}^*$ such that:

(i) Tech firm competition: $q^{\tau,n}(p)=pL+C_0$, $q^{\tau,i}(p)=pL+C_1$.

(ii) Incumbent conventional firm profit maximization: $q^{c,L}(p)=q^c$

(iii) Conventional firm competition: $\pi^c(q^c)=0$, and $\frac{d\pi^c(q^c)}{dq^c}>0$.

[9] These individuals remain in the pool as long as they are (a) alive and (b) have not yet had their type learned. The number of such individuals is $(1-\eta)\big(F(\tilde{p})-F(\hat{p})\big)(1+\eta(1-\alpha)+\cdots+\eta^t(1-\alpha)^t+\ldots)$. The total number of such individuals in total is $\big(F(\tilde{p})-F(\hat{p})\big)$.

(iv) Individual optimization $\frac{1}{1-\beta\eta} U(W - q^c) = U\left(W - q^{\tau,n}(\hat{p}^*)\right) + \frac{\beta\eta}{1-\beta\eta}\, U\left(W - q^{\tau,i}(\hat{p}^*)\right)$.

Equilibrium will again depend on the exogenous cost parameters C_0, C_1 and the speed of learning α. The following proposition is our main result for the OLG model. It establishes formally that the equilibrium cutoffs are again decreasing C_0 and C_1. In contrast to the FH model, however, faster learning by conventional firms actually *lowers* the conventional firm market share (raises the cutoff). The formal proof is in the appendix.

Proposition OLG: In any locally smooth family of equilibria:

(1) The equilibrium cutoff $\hat{p}^*$ is (weakly) decreasing C_0 and C_1, and strictly so if $\hat{p}^* \neq \underline{p}$.

(2) The equilibrium cutoff is *increasing* in α, strictly so if $\hat{p}^* \in (0,1)$ and $\beta < 1$.

3.3 Infinite Horizon Myopic Model

In this section, we present a model that describes the heuristic (myopic) decision making process that might reflect the real insurance market today. We again assume an infinite number of periods and a continuum of types with (cumulative) distribution $F(p)$ and continuous pdf $f(p) > 0$ on $\left[\underline{p}, \bar{p}\right] \subset (0,1)$, buying insurance in each period. The cost C for tech firms to provide insurance is (potentially) time dependent, with cost C_0 for a first-time tech buyer, and $C_1 \leq C_0$ for an individual who was previously insured with any tech insurer for k periods. (As we shall see, the C_1 will not end up being relevant, and neither would any C_k if the cost was decreasing further for longer incumbency periods.) In contrast to the FH and OLG model, customers choose firms based on the current-period prices without considering the consequences of current-period choices for prices in future periods. Also, conventional insurers use backward-looking expectations about their risk pool and update prices with a lag of one period. The tech firms compete in each period, offering $q_t^{\tau,n}(p) = pL + C_0$ for new customers and $q_t^{\tau,i}(p) = pL + C_\mathrm{k}$ for old ones.

In period 0, conventional firms cannot observe buyers' types when determining prices, and offer average premium $q_0^{c,U} \equiv \tilde{p}_0^* L$ equal to the expected loss of all risks in the market to each customer (where $\tilde{p}_0^* \equiv \int_{\underline{p}}^{\bar{p}} p f(p) dp$). In the subsequent periods ($t = 1, 2, \ldots$), if a conventional firm insured an individual in the period $t-1$, it will learn that policyholder's risk type with probability α.

In subsequent periods, conventional firms offer new customers whose type they don't know (either unlearned incumbents or new customers) $q_t^{c,U} = E[pL \middle| conv_{t-1}^U] \equiv \tilde{p}_t L$, the premium that would be fair given the losses they experienced within their unknown insured population in the previous period—whose mean risk is defined to be $\tilde{p}_t$. We assume symmetry across all conventional firms, so the firm-specific distribution $conv_{t-1}^U$ coincides with the market-wide distribution of customers who buy in the "unknown" market. As such, we can without loss of generality assume that all customers whose type was unlearned in period $t-1$ and who choose to buy from a conventional firm, re-enter a common pool and choose randomly among the conventional insurers.

We consider two different assumptions about $q_t^{c,L}(p)$, the premiums that the incumbent insurers offer to incumbents whom they have learned type. The first assumption is that competitive firms are savvy enough to recognize that knowing an individual's risk type potentially allows them to extract rents and that it knows (by introspection) the price that other conventional firms will offer to new customers. In this case, firms will set $q_t^{c,L}(p) = \min\{pL + C_0, q_t^{c,U}\}$ for any customer to whom it will be profitable to sell at this price (and otherwise will set a very high price to encourage the customer to go elsewhere). This pricing formula extracts the maximal rents possible given the effective competition. A second, simpler pricing scheme is to have $q_t^{c,L}(p) = pL + C_0$. This pricing rule is more naïve, in the sense that it effectively regards tech firms as the competitive fringe even when the other uninformed conventional firms may actually offer a better price.

3.3.1 Naive Pricing Environment

We first consider the naive pricing case with $q_t^{c,L}(p) = pL + C_0$. Individuals who purchased from a conventional firm in period $t-1$ whose type was learned by their incumbent firm choose in period t among three options: (1) switching to a tech firm offering a premium of $q_t^{\tau,n}(p) = pL + C_0$, (2) stay with the incumbent conventional firm offering a premium of $q_t^{c,L}(p) = pL + C_0$, and (3) entering the "unknown" market paying a premium of $q_t^{c,U}$. Individuals who purchased from a conventional firm in period $t-1$ and did not have their type learned have only choices (1) and (3). Since customers are myopic, they choose the lowest-premium option. It follows that, among individuals who purchased in the "unknown" conventional market in period $t-1$:

- The fraction $(1-\alpha)$ whose insurers did not learn their type will leave for a tech firm if $q_t^{\tau,n}(p) < q_t^{c,U}$, i.e., $p < \tilde{p}_t - \frac{C_0}{L}$. Otherwise, they will continue to participate in the "unknown" conventional market.[10]
- The fraction α whose insurers learned their type will leave the "unknown" market and stay with their incumbent firm, paying the premium $pL + C_0$, if $p < \tilde{p}_t - \frac{C_0}{L}$, and otherwise will continue to participate in the "unknown" conventional market (by randomly choosing a new firm that does not know their type).[11]

Thus, all types below p_t^*=max$\left\{\tilde{p}_t - \frac{C_0}{L}, \underline{p}\right\}$ exit the unknown conventional market, while all types above p_t^* remain in the unknown conventional market.

[10] We assume that the indifferent types with $p = \tilde{p}_t + \frac{C_0}{L}$ remain in the unknown conventional market. Because there is a measure zero of such types, this assumption is unimportant.

[11] Again, the measure zero of indifferent types are (safely) assumed to remain in the unknown conventional market.

We prove these cutoffs, $\tilde{p}_t$ and p_t^*, are non-decreasing in t inductively. First, $\tilde{p}_0^* \leq \tilde{p}_1$ (with equality if and only if $\tilde{p}_0^* > \underline{p}$) since the lowest risk types (potentially) leave the unknown conventional market for tech firms. If $\tilde{p}_{t-1} \leq \tilde{p}_t$, then no individuals will re-enter the unknown conventional market (but will instead stay either in the tech market or with their incumbent conventional firm who knows their type). Per the preceding bullets, all *exit* from the unknown conventional market will occur from individuals with $p < \tilde{p}_t$—i.e., from individuals with below-average risk. Hence $\tilde{p}_{t+1} \geq \tilde{p}_t$ and thus $p_{t+1}^* > p_t^*$.

For any $\alpha > 0$ and under the condition of $p_t^* < \tilde{p}_t \leq \bar{p}$, there exist $p^* \in [\underline{p}, \bar{p}]$ and $\tilde{p} \in [\underline{p}, \bar{p}]$ such that $\lim_{t\to\infty} p_t^* = p^*$ and $\lim_{t\to\infty} \tilde{p}_t = \tilde{p}$. p^* and $\tilde{p}$ are determined by the formulas: $\tilde{p} = \int_{p^*}^{\bar{p}} \frac{pf(p)}{1-F(p^*)} dp$ and $p^* = \tilde{p} - \frac{C_0}{L}$. Thus, $\tilde{p}$ does not necessarily equal to $\bar{p}$. Intuitively, the risk type p exits the unknown market when $p < p_t^* = \tilde{p} - \frac{C_0}{L}$, thus there will be some highest risk types $p \in [p^*, \bar{p}]$ fish around over time in the unknown market, looking for new conventional firms to insure with.

Definition: Quasi-equilibrium in the Infinite Horizon Myopic Market if $\boldsymbol{q_t^{c,L}(p) = pL + C_0}$

In each period t (t= 1, 2, …), the quasi-equilibrium is a set of prices $q^{\tau,n}(p)$, $q^{\tau,i}(p)$, $q_t^{c,L}(p)$, $q_t^{c,U}$, and a cutoff $p_t^* = \max\left\{\tilde{p}_t - \frac{C_0}{L}, \underline{p}\right\}$ such that:

(i) Tech firm competition: $q^{\tau,n}(p) = pL + C_0$, $q^{\tau,i}(p) = pL + C_1$.

(ii) Conventional firm quasi-competition: $q^{c,L}(p) = pL + C_0$ and $q_t^{c,U} = \tilde{p}_t L$, with $\tilde{p}_t = E[p \mid conv_{t-1}^U]$.

(iii) Individual quasi-optimization:

a. In period 0, individuals with $p < \tilde{p}_0^*$ purchase from tech firms $p \geq \tilde{p}_0^*$ purchase from conventional firms.

b. In period t, individuals who purchased from a tech firm in $t-1$ continue to purchase from tech firms; individuals who purchased from a conventional firm who knew their type continue to purchase from that firm.

For individuals with $p_{t-1}^* < p < p_t^*$ who purchased in the "unknown type" market in $t-1$, a fraction α individuals who had their type learned continue to purchase from their current firm, and the fraction $1-\alpha$ of such individuals who did not have their type learned purchase from a tech firm. For individuals with $\geq p_t^*$, they stay in the "unknown type" market.

c. When $t \to \infty$, $\tilde{p} \to \int_{p^*}^{\bar{p}} \frac{pf(p)}{1-F(p^*)} dp$, and $p^* \to \tilde{p} - \frac{C_0}{L}$. Individuals with $p < p^*$ were either learned and continue to insure with the incumbent conventional firms or insure with the tech firms. Some highest risk types with $p \in [p^*, \bar{p}]$ fish around in the unknown market, looking for new conventional firms to insure with.

The quasi-equilibrium depends on the exogenous cost parameters C_0 and the speed of conventional insurer learning α. Note that p_t^* and $\tilde{p}_t$ do not depend on α: $\tilde{p}_0^* \equiv \int_{\underline{p}}^{\bar{p}} pf(p)dp$ is independent of α, and, inductively, if p_{t-1}^* and $\tilde{p}_{t-1}$ do not depend on α, then $\tilde{p}_t = E(p|p \in [\tilde{p}_{t-1}^*, \bar{p}])$ and $p_t^* = \max\left\{\tilde{p}_t - \frac{C_0}{L}, \underline{p}\right\}$ are also independent of α.

We also see that p_t^* is decreasing with C_0 for $t > 0$. Inductively: $\tilde{p}_0^* \equiv \int_{\underline{p}}^{\bar{p}} pf(p)dp$ is independent of C_0; thus, we have $\tilde{p}_1 = E(p|p \in [\tilde{p}_0^*, \bar{p}])$, and thus $p_1^* = \max\left\{E(p|p \in [\tilde{p}_0^*, \bar{p}]) - \frac{C_0}{L}, \underline{p}\right\}$ is decreasing with C_0; when p_{t-1}^* is decreasing with C_0, with $\tilde{p}_t = E(p|p \in [\tilde{p}_{t-1}^*, \bar{p}])$, $p_t^* = \max\left\{E(p|p \in\right.$ *[pt−1∗, p])−C0L,p* is decreasing with *C0*.

For any period *t* (*t*= 1, 2,…), the share of tech firms equals to $\mathrm{F}\left(\tilde{p}_0^* - \frac{C_0}{L}\right) + (1-\alpha)\left(\mathrm{F}\left(\tilde{p}_t - \frac{C_0}{L}\right) - \mathrm{F}\left(\tilde{p}_0^* - \frac{C_0}{L}\right)\right) = \alpha\mathrm{F}\left(\tilde{p}_0^* - \frac{C_0}{L}\right) + (1-\alpha)\mathrm{F}\left(\tilde{p}_t - \frac{C_0}{L}\right)$. It is easy to see the share of tech firms decreases with α and C_0. When $t \to \infty$, the market share of tech market is $\mathrm{F}\left(\tilde{p}_0^* - \frac{C_0}{L}\right) + (1-\alpha)\left(\mathrm{F}\left(\tilde{p} - \frac{C_0}{L}\right) - \mathrm{F}\left(\tilde{p}_0^* - \frac{C_0}{L}\right)\right) = \alpha\mathrm{F}\left(\tilde{p}_0^* - \frac{C_0}{L}\right) + (1-\alpha)\mathrm{F}\left(\tilde{p} - \frac{C_0}{L}\right)$. It is easy to see the market share of tech market is decreasing with α and C_0.

Proposition Myopic assuming naïve pricing i.e. $\boldsymbol{q_t^{c,L}(p) = pL + C_0}$:
In each period t, tech market share is decreasing in α and C_0 (and independent of C_1). The results are thus consistent with those from the FH model.

3.3.2 Savvy Pricing Environment

As a second case, we consider a savvy pricing rule: conventional firms who have learned a customer's type set $q_t^{c,L}(p) = \min\{pL + C_0, q_t^{c,u}\}$ for all customers with $p < \tilde{p}_t$, and any price at or above $q_t^{c,u}$ otherwise (so that they offload the higher risks who cannot be attracted at any profitable price).[12] Individuals face different pricing but the same basic choices as in the Naïve pricing environment and again myopically choose the best price. It is useful to define the terms $\tilde{p}_t$ and p_t^* exactly as in the Naïve model. Here:

- The fraction $(1-\alpha)$ whose insurers did not learn their type will leave for a tech firm if $q_t^{\tau,n}(p) < q_t^{c,U}$, i.e., $p < \tilde{p}_t - \frac{C_0}{L}$. Otherwise, they will continue to participate in the "unknown" conventional market.

[12] It is easy to show that all of the key results in this section will be the same if instead the incumbent conventional firms offer the acturially fair premium to the learned types, i.e. $q_t^{c,L}(p) = pL$. The proof is available from the authors upon request.

- The fraction α whose insurers learned their type will leave the "unknown" market and stay with their incumbent firm if $p < \tilde{p}_t$. Types with $p \geq \tilde{p}_t$ will continue to participate in the "unknown" conventional market (by randomly choosing a new firm that does not know their type).

Both $\tilde{p}_t$ and $p_t^* = \max\left\{\tilde{p}_t - \frac{C_0}{L}, \underline{p}\right\}$ play important cutoff roles here. First, the fraction α of all learned types below $\tilde{p}_t$ leave the unknown conventional market for a "learned" contract. Second, all unlearned types below p_t^* exit the unknown conventional market for a tech firm. . Following the same logic as in Section 3.3.1, it is straightforward to see that $\{p_t^*\}$ and $\{\tilde{p}_t\}$ are non-decreasing in t and $\lim_{t\to\infty} \tilde{p}_t = \bar{p}$ for any $\alpha > 0$.

Definition: Quasi-equilibrium in the Infinite Horizon Myopic Market with Savvy Firms

In each period t (t= 1, 2, …), the quasi-equilibrium is a set of prices $q^{\tau,n}(p)$, $q^{\tau,i}(p)$, $q_t^{c,L}(p)$, $q_t^{c,U}$, and a set of cutoffs $\tilde{p}_t$ and $p_t^* = \max\left\{\tilde{p}_t - \frac{C_0}{L}, \underline{p}\right\}$ such that:

(i) Tech firm competition: $q^{\tau,n}(p) = pL + C_0$, $q^{\tau,i}(p) = pL + C_1$.

(ii) Conventional firm quasi-competition: $q^{c,L}(p) = pL$ and $q_t^{c,U} = \tilde{p}_t L$, with $\tilde{p}_t = E[p|conv_{t-1}^U]$.

(iii) Individual quasi-optimization:

a. In period 0, individuals with $p < \tilde{p}_0^*$ purchase from tech firms and with $p \geq \tilde{p}_0^*$ purchase from conventional firms.

b. In period t, individuals who purchased from a tech firm in $t-1$ continue to purchase from tech firms; individuals who purchased from a conventional firm who knew their type continue to purchase from that conventional firm.

For individuals who purchased in the "unknown type" market in $t-1$, a fraction α individuals who had their type learned continue to purchase from their current firm if $p < \tilde{p}_t$ and otherwise buy in the "unknown type" market (from another random conventional firm), and the fraction $1-\alpha$ of such individuals who did not have their type learned purchase from a tech firm if $p < p_t^*$ and otherwise stay in the "unknown type" market.

c. When $t \to \infty$, $\tilde{p}_t \to \bar{p}$, and $p_t^* \to \bar{p} - \frac{C_0}{L}$. All risks were either learned and continue to insure with the incumbent conventional firms or insure with the tech firms. The unknown type market diminishes.

In this setup, the comparative statics of the share of individuals insuring with a tech firm may not be monotone in α. On the one hand, a higher α increases learning and means that more low risk types have their types learned early and thus stay with conventional firms forever. Thus, larger α would seem to imply lower tech shares. But there is an offsetting effect: it exacerbates adverse selection within the conventional unlearned market, as a higher α means that additional, low-risk individuals who would have stayed in the "unlearned type" market now exit for the learned type conventional market. Their exit raises the average riskiness of the "unlearned type" pool, implying a higher p_t^* and hence more types who leave for tech firms in the subsequent period. To put it another way: all else equal, a higher α implies more "unraveling" towards the *learned* conventional market instead of the tech market; but it also implies faster unraveling.

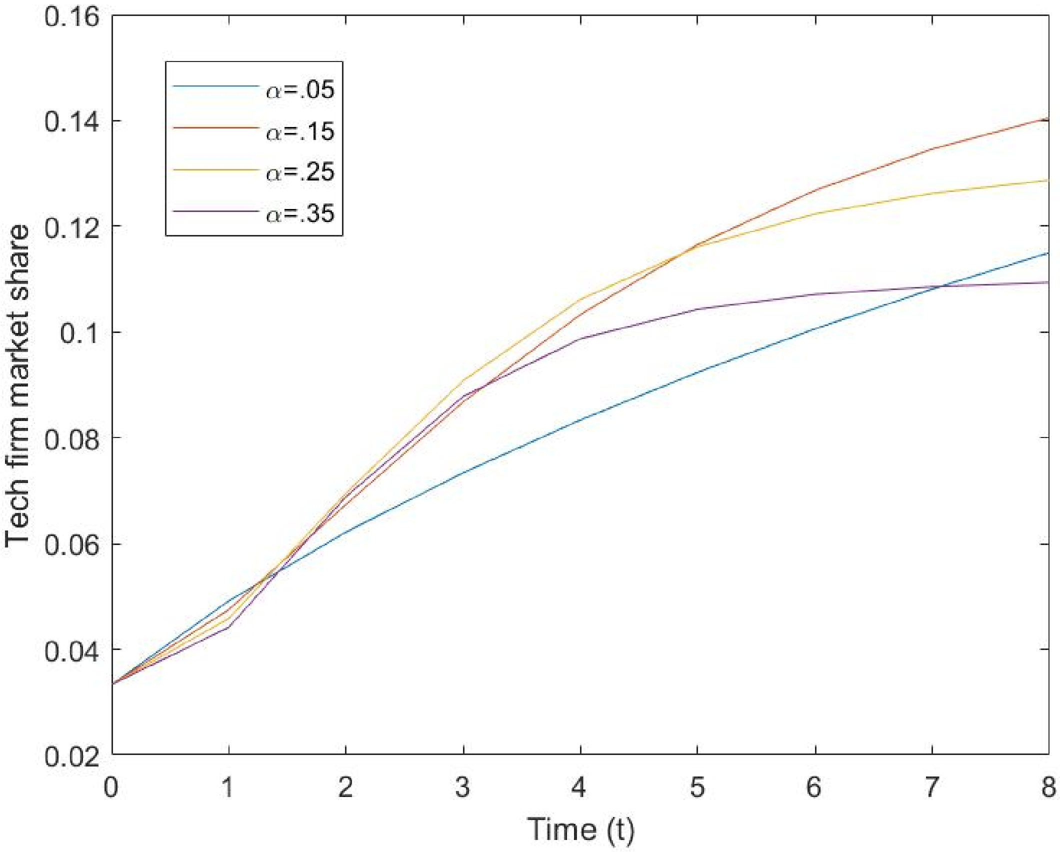

Figure 1: Temporal evolution of the tech firm market share. Assumes uniform distribution of p on [0.3, 0.9], $C_0 = 0.28L$.

Figure 1 illustrates the potential for non-monotone comparative statics with respect to α. The tech shares for four different levels of α "cross" each other—and hence switch order—over time. In period $t = 8$, for example, the tech share increases from $\alpha = 0.05$ to $\alpha = 0.15$ but then decreases when α is further increased to 0.25 and 0.35.

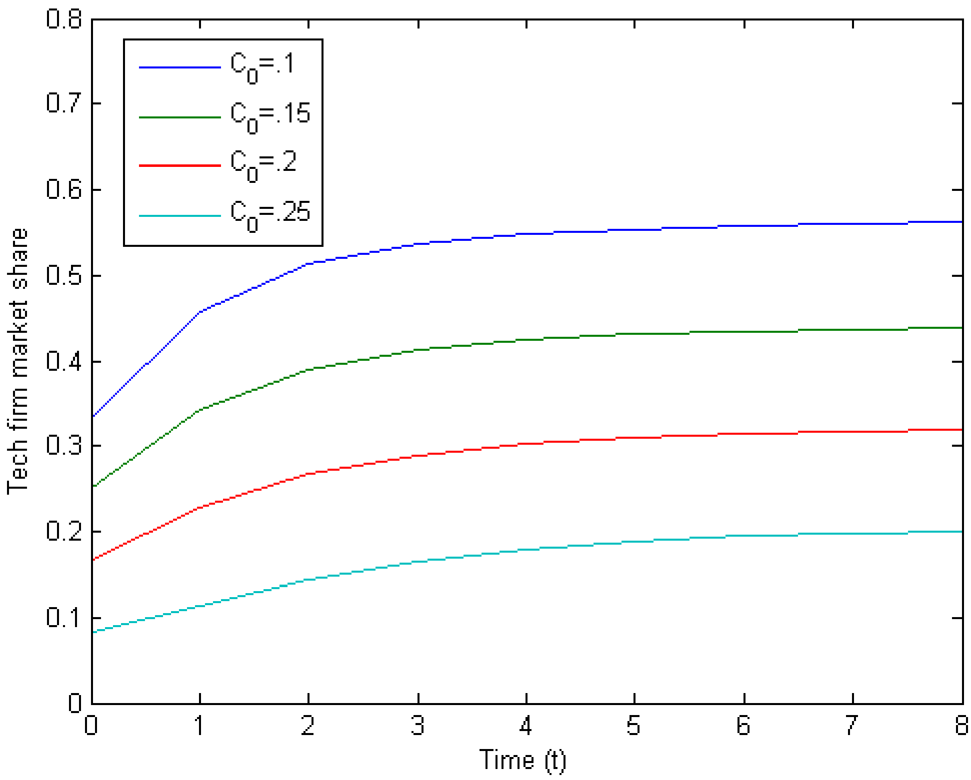

Figure 2: Temporal evolution of the tech firm market share. Assumes uniform distribution of p on [0.3, 0.9], $\alpha = 0.25$.

A higher C_0 implies, for any given distribution of types in the "unlearned type" market, a smaller share of individuals leaving for tech firms in the current period. On the other hand, it also implies slower unraveling of the market. Intuitively, this slower unraveling should give individuals in the "unlearned types" market more opportunities to be learned before the price is driven up to the point where they will exit to a tech firm—and hence should also lead to a smaller tech share. profit is straightforward to confirm the intuition that the tech market share is decreasing in C_0 for periods 0 and 1, and our simulations to date have all indicated this same pattern for all periods (see, e.g., the example in Figure 2). But we do not yet have a formal proof and, indeed, suspect there may be pathological examples with highly non-uniform distributions of types for which we can generate non-monotonicities.

"Proposition" Myopic if $q_t^{c,L}(p) = \min\{pL + C_0, q_t^{c,U}\}$:

The period t tech market share may be increasing *or* decreasing in α. The tech market share can be decreasing in C_0. Whether the tech market share can be increasing in C_0 is unknown. The tech market share is independent of C_1.

4. Conclusions

$Table$ 1 summarizes our main results, $i.e.$ the (quasi)-equilibrium predictions of the three models.

Table 1 Main Results

	Model structure / Questions	Finite horizon model with foresight	Infinite horizon OLG model	Infinite horizon myopic model
1	*How do risk types distribute between tech and conventional firms?*	There is a period 1 cutoff type p_1^*. Types $p < p_1^*$ buy from tech insurers. There is a a period 2 cutoff type $p_1^* \geq p_0^*$. Individuals with $p < p_1^*$ switch to tech insurers.	There is a cutoff type $\hat{p}^*$. All individuals with $p < \hat{p}^*$ purchase from tech insurers.	In each period, there is a cutoff p_t^* for current buyers in the conventional market with unlearned types. Types with $p < p_t^*$ buy from tech firms unless their type is learned by their insurer.
2	*How does the market share of tech insurer correlate with the learning speed of conventional insurers?*	p_1^* is decreasing and p_2^* is non-increasing with the speed of learning. Hence, the tech market share is weakly decreasing with the speed of learning.	$\hat{p}^*$(and hence the tech market share) is increasing with the speed of learning.	Tech market share is decreasing with the speed of learning under a naïve pricing environment. It is, however, ambiguous under a savvy pricing environment.
3	*How does the market share of tech insurers correlate with the additional cost of digital technology?*	p_1^* and p_2^* (and hence the tech market share) are decreasing in the cost of technology.	$\hat{p}^*$(and hence the tech market share) is decreasing with the cost of technology.	Tech market share is decreasing in the cost of technology.

Notes: For details see Propostions FH, OLG, and Myopic

All three models feature a "cutoff" structure, where low risk types buy from tech firms while higher risks buy from conventional firm. This cutoff structure reflects a basic tradeoff at the heart of all models: on the one hand, tech firms have an informational advantage; on the other hand, they bear a higher cost of operation, which they pass on to customers. In general, this tradeoff allows both types of firm to coexist.

In light of this tradeoff, one might be inclined to conjecture that anything which tilts this tradeoff more in favor of conventional firms or against tech firms will lead to a smaller tech market share. This is indeed the case for the costs: across all our market models, a higher cost disadvantage leads to a lower tech market share (with the possible exception of the myopic model with savvy firms, where the question remains open).

Interestingly, however, the effect of shrinking the *informational* disadvantage between tech and conventional firms—i.e., increasing the speed α at which conventional firms learn—does not always lead to falling tech market shares. Whether it does or does not depends on the setting. In the FH model, a higher α does, as expected, always lead to a *smaller* tech share. The same is true in the myopic model with naïve pricing. In the OLG setting, however, shrinking the informational disadvantage leads to a *larger* tech share. And, in the myopic model with savvy investors, the effect is ambiguous, and examples can be shown where the effects go in either direction.

Though perhaps ex-ante unexpected, there is some clear economic intuition for the differences across settings. In general, there are economic forces pushing in both directions. On the one hand, faster learning by conventional firms does indeed reduce the informational advantage of tech firms. This tends to lower the tech market share. On the other hand, faster learning also increases the informational advantage of incumbent conventional firms over other conventional firms, exacerbating adverse selection *among* conventional firms. Worse adverse selection in the conventional market redounds to the benefit of tech firms, who do not suffer from it.

In the FH model and the naïve myopic model, the latter channel is shut down entirely. In the FH model, it is shut down because we assume that there is an active market for unknown types only once, in period 0, when there are no incumbent firms with informational advantages. In the naïve myopic model, because incumbent conventional firms who have learned their customer's type price exactly as tech firms would, the rate at which the conventional market for unlearned types unravels (via adverse selection as low risks leave the pool) is independent of α. A change in α does not exacerbate adverse selection.

In the savvy myopic model, on the other hand, a higher α does exacerbate adverse selection, as incumbent firms use their informational advantage to retain some additional types who are relatively low risk but who would not have left the market for tech firms (because of the cost disadvantage). The non-monotonic behavior of the tech market share in Figure 1 shows that this adverse selection channel can potentially overwhelm the "faster learning reduces tech firms' informational advantage" channel.

Finally, the OLG model effectively shuts down the "faster learning reduces tech firms' informational advantage" channel. The equilibrium is stationary, and consumers always have the choice of a tech firm or the "unlearned" market—and always at the same price. So, for any customer who original chose a conventional firm, the relevant "outside option" is always an "unlearned" market contract. As such, informed firms are (effectively) not competing with tech firms at all—their greater knowledge therefore only exacerbates adverse selection within the conventional market, and the tech share increases.

Our modeling frameworks are perhaps best seen a capturing costly—but optional—risk classification technologies. Since our framework is about unknown-type informational issues, they are less well suited to modeling unknown action contexts. Some telematics technologies (in-car, in home, or wearable health-monitoring devices e.g.) are at least in part about monitoring *behavior* in order to reduce moral hazard; our approach is less well-adapted to such settings – although we believe similar techniques can be used and some parallel insights are likely to appear in the presence of moral hazard effects.

There are several other directions in which we believe our analysis can be fruitfully extended. In this paper we focus exclusively on the trade-off between the costs of new technologies and the informational advantages thereof. We neglect other potentially important tradeoffs, such as the tradeoff between subjective transparency costs and informational advantages studied in Gemmo et al., (2017). We view the incorporation of moral hazard considerations, transparency aversion, and other complications as important directions for further study.

References

Akerlof, G. (1970). The market for lemons: Qualitative uncertainty and the market mechanism. *Quarterly Journal of Economics*, *84*(3), 488-500.

Dionne, G., Michaud, P. C., and Dahchour, M. (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.

de Garidel-Thoron, T. (2005). Welfare-improving asymmetric information in dynamic insurance markets. *Journal of Political Economy*, 113(1), 121-150.

Doherty, N. and Posey, L. (1998). On the value if a checkup: Adverse selection, moral hazard and the value of information. *Journal of Risk and Insurance* 65(2), 189-211.

Eling, M., Jia, R., and Yao, Y. (2017). Between-group adverse selection: Evidence from group critical illness insurance. *Journal of Risk and Insurance*, 84(2), 771-809.

Filipova-Neumann, L. and Welzel, P. (2010). Reducing asymmetric information in insurance markets: Cars with black boxes. *Telematics and Informatics*, 27(4), 394-403.

Gemmo, I., Browne, M., and Gründl, H., (2017). Transparency aversion and insurance market equilibria. Working paper, http://www.icir.de/fileadmin/user_upload/WP2517.pdf

Hendel, I. (2016). Dynamic selection and reclassification risk: Theory and empirics. Working paper, Northwestern University.

Hoy, M., and Ruse, M. (2005). Regulating genetic information in insurance markets. *Risk Management and Insurance Review*, 8(2), 211-237.

Kofman, P., and Nini, G. P. (2013). Do insurance companies possess an informational monopoly? Empirical evidence from auto insurance. *Journal of Risk and Insurance*, 80(4), 1001-1026.

Kunreuther, H., and Pauly, M. (1985). Market equilibrium with private knowledge: An insurance example. *Journal of Public Economics*, 26(3), 269-288.

Nilssen, T. (2000). Consumer lock-in with asymmetric information. *International Journal of Industrial Organization*, 18(4), 641-666.

Watt, R., and Vazquez, F. J. (1997). Full insurance, Bayesian updated premiums, and adverse selection. *Geneva Papers on Risk and Insurance Theory*, 22(2), 135-150.

Appendix

Proof of proposition FH

We consider two distinct cases: case 1: $p_0 = \hat{p}(q_0^c)$ and case 2: $p_0 = \underline{p} > \hat{p}(q_0^c)$.

Case 2: in this case, conventional firms sell to everybody in period 0, $p_0 = \underline{p}$, and, along a smooth family of equilibrium, this cutoff will be locally independent of small changes in α, C_0, and C_1. For local changes in α, C_0 and C_1, the period 1 cutoff p_1 solves:

$$\max_{p^\dagger \in [\underline{p}, \overline{p}]} \int_{p^\dagger}^{\overline{p}} \left(p^\dagger L - C_0 - pL\right) f(p) dp \equiv \max_{p^\dagger \in \left[\underline{p}, \overline{p}\right]} \underline{\pi}^1\left(C_0, p^\dagger\right).$$

Since $\underline{\pi}^1$ is independent of α and C_1, so is p_1. We compute $\frac{\partial^2 \underline{\pi}^1}{\partial C_0 \partial p^\dagger} = f\left(p^\dagger\right)$, which is strictly positive (as, otherwise, $\frac{\partial \underline{\pi}^1}{\partial p^\dagger} > 0$ and $p^\dagger$ would not be optimal). By Topkis's Theorem, the cutoff p_1 is increasing in C_0, strictly so if $p_1 > \underline{p}$.

Case 1: For any given q_0^c, it is straightforward to see that raising C_0, C_1, or α raises profits. So q_0^c mut fall in response by the stability requirement in part (iv) of the definition of an FH equilibrium. Since it is easy to see that a rise in $\hat{p}$ (together with a rise in C_0 or C_1) would imply an increase in q_0^c. So $\hat{p}$ must fall.

The period 1 cutoff satisfies:

$$\max_{p^\dagger \in [p_0(\alpha, C_0, C_1), \overline{p}]} \int_{p^\dagger}^{\overline{p}} \left(p^\dagger L - C_0 - pL\right) f(p) dp \equiv \max_{p^\dagger \in [p_0(\alpha, C_0, C_1), \overline{p}]} \underline{\pi}^1\left(C_0, p^\dagger\right).$$

An increase in α or C_1 only affects this problem by lowering the lower bound $p_0(\alpha, C_0, C_1)$. Hence, p_1 is weakly decreasing in α and C_1. By the same argument as in case 2, the $p^\dagger$ solving $\max_{p^\dagger \in [p^*, \overline{p}]} \underline{\pi}^1\left(C_0, p^\dagger\right)$ for any *fixed* p^* is decreasing in C_0, and strictly so unless the optimum is at p^*. Together with the fact that the lower bound $p_0(\alpha, C_0, C_1)$ is strictly decreasing in C_0 implies that p_1 is decreasing in C_0, and strictly so unless $p_1 = p_0$.

Proof of proposition OLG

First, observe that

$$V(q^c,\hat{p}^*;C_0,C_1) = \frac{1}{1-\beta\eta}U(W-q^c) - \left[U\big(W-q^{\tau,n}(\hat{p}^*)\big) + \frac{\beta\eta}{1-\beta\eta}\,U\left(W-q^{\tau,i}(\hat{p}^*)\right)\right]$$
$$= \frac{1}{1-\beta\eta}U(W-q^c) - [U(W-\hat{p}^*L-C_0) + \frac{\beta\eta}{1-\beta\eta}U(W-\hat{p}^*L-C_1)$$

is strictly increasing in C_0 and C_1 and strictly decreasing in $\hat{p}$. Thus, for any given q^c, the $\hat{p}^*(q^c;C_0,C_1)$ satisfying $V(q^c,\hat{p}^*;C_0,C_1)=0$ is strictly decreasing in C_0 and C_1 unless $\hat{p}^*(q^c;C_0,C_1)=\underline{p}$. Given any q^c, $\hat{p}^*(q^c;C_0,C_1)$ is obviously independent of α.

Denoting by $\mathbb{E}\pi_M \equiv \mathbb{E}\big[q^c - Lp|p\in[\hat{p}(q^c),\tilde{p}(q^c)]\big]$ and $\mathbb{E}\pi_H \equiv \mathbb{E}\big[q^c - Lp|p\in[\tilde{p}(q^c),\overline{p}]\big]$, lifetime per-sale firm profits in the "unlearned pool" can be written, using the definition of $Q_M(q^c)$ as:

$$\pi^c(q^c) = Q_M(q^c)\left(1+\frac{\alpha\eta\beta}{1-\eta\beta}\right)\mathbb{E}\pi_M + \big(1-Q_M(q^c)\big)\mathbb{E}\pi_H$$
$$= Q_M(q^c)\left(\frac{1-\eta+\alpha\eta}{1-\eta}\right)\frac{\left(\frac{1-\eta\beta+\alpha\eta\beta}{1-\eta\beta}\right)}{\left(\frac{1-\eta+\alpha\eta}{1-\eta}\right)}\mathbb{E}\pi_M + \big(1-Q_M(q^c)\big)\mathbb{E}\pi_H$$
$$= \big(1-Q_M(q^c)\big)\frac{1-F\big(\hat{p}(q^c)\big)}{1-F\big(\tilde{p}(q^c)\big)}$$
$$\times\left[\frac{\Big(F\big(\tilde{p}(q^c)\big)-F\big(\hat{p}(q^c)\big)\Big)}{1-F\big(\hat{p}(q^c)\big)}\frac{\left(\frac{1-\eta\beta+\alpha\eta\beta}{1-\eta\beta}\right)}{\left(\frac{1-\eta+\alpha\eta}{1-\eta}\right)}\mathbb{E}\pi_M + \frac{1-F\big(\tilde{p}(q^c)\big)}{1-F\big(\hat{p}(q^c)\big)}\mathbb{E}\pi_H\right].$$

Hence, $\pi^c(q^c)$ is independent of C_0 and C_1 for any given q^c *except* through the effect on $\hat{p}^*(q^c;C_0,C_1)$, and it is easy to see that $\pi^c(q^c)$ is strictly decreasing in $\hat{p}^*(q^c;C_0,C_1)$. Together with the preceding paragraph, this implies

Observation 1: $\pi^c(q^c)$ is increasing in C_0 and C_1, strictly so unless $\hat{p}^*(q^c;C_0,C_1)=\underline{p}$.

Finally, observe that the expression:

$$\frac{\left(\frac{1-\eta\beta+\alpha\eta\beta}{1-\eta\beta}\right)}{\left(\frac{1-\eta+\alpha\eta}{1-\eta}\right)}$$

is strictly decreasing in α (except in the limit as $\beta\to 1$, in which case it is independent of α. It follows that (starting from an equilibrium where profits are zero), $\pi^c(q^c)$ is strictly decreasing in α as long as $\tilde{p}(q^c) > \hat{p}(q^c)$. And $\tilde{p}(q^c) > \hat{p}(q^c)$ obviously holds in equilibrium if $\alpha > 0$: if $\tilde{p}(q^c) = \hat{p}(q^c)$, conventional firms never retain any customers, which means that they must break even on the average risk type above $\hat{p}(q^c)$. But then would be strictly profitable to sell to learned types with type $\hat{p}(q^c)$. This yields:

Observation 2: For any q^c, $\pi^c(q^c)$ is decreasing in α, strictly so if $\beta < 1$.

Observations 1 and 2, together with the fact that $\pi^c(q^c) = 0$ and $\frac{d\pi^c(q^c)}{dq^c}$ is increasing in any equilibrium, implies that the equilibrium q^c is increasing in C_0 and C_1 (strictly so if $\underline{p} \neq \hat{p}^*(q^c)$) and decreasing in α (strictly so if $\beta < 1$).

Observing that $\hat{p}^*(q^c)$ is decreasing in q^c (strictly so if $\underline{p} \neq \hat{p}^*(q^c)$) completes the proof.

It is worth commenting briefly on the intuition behind the α result. The expression for π^c above implies that $\pi^c = 0$ if and only if

$$\left[\frac{\left(F\big(\tilde{p}(q^c)\big) - F\big(\hat{p}(q^c)\big)\right)}{1 - F\big(\hat{p}(q^c)\big)} \frac{\left(\frac{1-\eta\beta+\alpha\eta\beta}{1-\eta\beta}\right)}{\left(\frac{1-\eta+\alpha\eta}{1-\eta}\right)} \mathbb{E}\pi_M + \frac{1 - F\big(\tilde{p}(q^c)\big)}{1 - F\big(\hat{p}(q^c)\big)} \mathbb{E}\pi_H\right] = 0.$$

When $\beta = 1,$ this is a weighted average of the profits from "retained" types and "non-retained" types, where the weights are the *population* shares of these types. To see why, note that all types above $\hat{p}(q^c)$ are buying from conventional firms at all times. When there is no time discounting, the cross-sectional profits *per period* must be zero. Time discounting does not (directly) the distribution of types over time—it still matches the cross sectional distribution. But from an individual firm's point of view a sale loses money immediately, and will make that profit up later if they learn that they had sold to a lower risk type. With time discounting, though, being later in time is less valuable. So if the *cross sectional* profits were zero, then the per-firm discounted profits from a sale would be negative. Prices must therefore be higher to break even when there is discounting. Moreover, the faster is the learning, the more back-loaded is the profit stream, and the higher the prices have to be. This is manifesting itself in the preceding formula in the term

$$\frac{\left(\frac{1-\eta\beta+\alpha\eta\beta}{1-\eta\beta}\right)}{\left(\frac{1-\eta+\alpha\eta}{1-\eta}\right)},$$

The numerator of which is, effectively, the *present discounted* "number of periods" of sales to profitable learned types. The denominator is the *undiscounted* "number of periods". The gap between these two grows with α.