How Initial Conditions Can Have Permanent Effects: The Case of the Affordable Care Act

Florian Scheuer
Stanford, Zurich and NBER*

Kent Smetters
Wharton and NBER*

February 2017

Abstract

We document that states that experienced website glitches in the ACA's first year faced higher average costs that persisted to future years. This dynamic is not consistent with the standard model with strategic pricing, which requires non-localized common knowledge, but is consistent with price-taking. Initial conditions can have a permanent effect—including a Pareto-inferior, stable equilibrium—under conditions that we show are plausible in this setting. Changing the fine from a fixed value consistent with current law to a fraction of equilibrium prices increases the likelihood of reaching a Pareto-efficient equilibrium without increasing the equilibrium fine collected.

Keywords: Health Exchanges, Adverse Selection, Insurance Markets, Mandate

*Email addresses: scheuer@stanford.edu, smetters@wharton.upenn.edu. We are grateful to Jeffrey Clemens, Michael Dickstein, Liran Einav, Daniel Gottlieb, Nathan Hendren, Amanda Kowalski, Casey Rothschild and seminar participants at Stanford, Wharton and the NBER Insurance Meeting 2014 for helpful comments

1 Introduction

The Patient Protection and Affordable Care Act (ACA) limits the degree to which insurers can price discriminate based on age and preexisting conditions, while fining people who do not obtain health coverage. Enrolling healthier people is, therefore, widely regarded as important for avoiding high premiums. However, several of the ACA's "initial conditions" might have not been sufficient to encourage healthy individuals to enroll. The initial condition that received the lion's share of attention—and the one that is most plausibly exogenous—was the failure of the websites that managed the enrollment process.[1] As *The Economist* (November 23, 2013) put it:

> Insurers have set their premiums on the assumption that lots of young, healthy people would be compelled to buy their policies. But if it takes dozens of attempts to sign up, the people who do so will be disproportionately the sick and desperate. Insurers could be stuck with a far more expensive pool of customers than they were expecting, and could have no choice but to raise prices next year. That would make Obamacare even less attractive to the young "invincibles" it needs to stay afloat. (p. 15)

In the United States, insurance is regulated at the state level, where enrollment also occurs. The ACA permits individual states to set up their own compliant websites for enrollment, or states could elect to use the federal website. Besides the well-documented problems with the federal website, some states that established their own websites also suffered from website failures. We, therefore, can distinguish between three types of enrollment experiences: (i) "Glitch" states that established their own ACA exchanges but suffered severe technological glitches during the 2013 open enrollment season, immedi-

[1] There were several other initial conditions that likely played a role in reducing enrollment by healthier households. First, the ACA fine (the "shared responsibility penalty") was quite small in 2014, equal to the greater of $95 or 1% of income, growing modestly in the following years. Moreover, these fines can only be levied against positive tax refunds. Second, around half of non-compliant households are exempted from paying the fine, including due to financial hardship (Pauly (2017)). Third, because some households on the individual health insurance market lost their coverage, the government announced on November 14, 2013 that it was allowing individual state insurance commissioners to extend canceled policies by one year, a move widely denounced by insurers as potentially creating adverse selection. Fourth, many young shoppers with new employers also have to separately submit payroll stubs, re-confirm their health exchange status at a later time, and then contact the insurer to make a payment. A potential counterbalancing effect is the fact that, if the initial enrollment deadline is missed, subsequent enrollment is delayed until the next open enrollment season. However, this effect was weakened during the initial year of the ACA implementation since open enrollment occurred twice in 2014, roughly six months apart, in order to make its timing consistent with Medicare's open enrollment season in the following years. Moreover, as shown in Section 2, broad exemptions where provided for "exceptional" cases, which, in practice allowed for considerable enrollment outside of open enrollment.

ately before ACA coverage became available in 2014; (ii) "No Glitch" states that established their own exchanges and had no major technological glitches; and, (iii) "Federal" states that used the federal website, which also had numerous glitches. This classification follows the website review by Dash and Thomas (2014) and was subsequently used by Kowalski (2014).

Using data from the National Association of Insurance Commissioners, we show in Section 2 that "Glitch" and "Federal" states, indeed, suffered from much larger increases in average costs in 2014, the first year of ACA implementation. Remarkably, this pattern then persisted into 2015 and 2016, suggesting that poor initial conditions might have permanent effects. In essence, "Glitch" and "Federal" states appear to have converged, over time, to a "bad" permanent equilibrium, whereas the "No Glitch" states have settled into a better equilibrium. Absent other systematic differences between these states, this outcome requires the existence of multiple equilibria.

The textbook theory of insurance unraveling, however, is not specified in terms of initial conditions but as the equilibrium of a static system of insurance cost and demand equations across risk types (Akerlof (1970), Wilson (1977, 1980)). The main strand of literature has implicitly focused on price-taking with linear demand and cost curves (see, e.g., Cutler and Reber (1998), Einav and Finkelstein (2011), and Hackmann et al. (2015)). Linear curves produce a unique equilibrium that emits a degree of risk sharing ranging from full pooling to no pooling ("unraveling"), or something in between. For example, Handel et al. (2015) analyze a model of the ACA health exchanges with a unique equilibrium and conclude that it may eventually involve limited degrees of risk pooling.[2] But there is no role for initial conditions in choosing that outcome. We present simulation evidence (Section 5) that demonstrates a strong case for nonlinear curves in the context of health insurance.

Another and more recent strand of literature has considered strategic insurers in this setting, potentially with nonlinear cost and demand curves (Einav et al. (2010a)). In Section 3, however, we evoke the well-known result that a model of strategic insurers also cannot produce multiple equilibria, even with nonlinear demand and cost curves. A profitable deviation would always exist at a low level of coverage. As Mas-Colell et al. (1995) emphasize in the setting of labor markets with adverse selection, the strategic model assumes that firms have common knowledge about all market fundamentals, including the *global* shape of the demand and cost curves. We illustrate that this assumption is very strong in our setting. When it fails, insurers will likely resort to less risky, adap-

[2]See also Mahoney and Weyl (forthcoming) who consider the interaction between market power and selection in a model with a unique equilibrium.

tive premium setting that effectively turns them into price-takers, and we provide further statutory reasons that might compel this type of behavior.

In fact, in the numerous media articles discussing the importance of the initial health mix for future prices, we could not find evidence that suggested that an insufficient mix of younger enrollees might actually lead insurers to *reduce* premiums, in order to improve the risk pool, compatible with strategic pricing. Consistently, Cutler and Reber (1998), Monheit et al. (2004), and Clemens (2015) provide evidence of repeated marginal price changes that suggest that insurers do not *a priori* know the entire shape of the demand and cost curves in the market and locally adjust premiums in response to profits or losses they experience.

In Sections 4 and 5, we present a model, which does not constrain the shape of the cost and demand curves, to demonstrate that initial conditions only become relevant when firms are (i) price-takers, and (ii) face at least one nonlinear demand or cost curve. While receiving less attention in the insurance literature, price-taking with nonlinear curves has been a textbook model in the study of adverse selection in labor markets, including Mas-Colell et al. (1995). However, that literature has not emphasized the distinction between stable and unstable equilibria. Only stable equilibria matter for public policy purposes, and with two equilibria, only one of them can be stable. Instead, embedding the nonlinear price-taking model within an inter-temporal framework leads to the intuitive condition that there must be at least three equilibria for the effect of initial conditions to persist over time, one of which must be unstable. While satisfying this condition might seem like a tall order, using data from the Medical Expenditure Survey, we argue that in the context of health insurance there is sufficient nonlinearity to make this outcome a genuine possibility, if not the likely outcome, at realistic levels of risk aversion.

Our baseline analysis does not consider two key policy features of the ACA, the levy of fines on those who do not purchase insurance, and subsidies for many households that do. Fines and subsidies play fairly uninteresting roles from a welfare perspective in the textbook models outlined above. While a fine, for example, can force a higher insurance take-up rate, it is generally not Pareto improving. However, these mechanisms play a more important role in our three-equilibria model, where the "bad" equilibrium with lower coverage is Pareto inefficient. Section 6 shows that fines and subsidies expand the likelihood of arriving at the "good" equilibrium with relatively higher coverage from a given initial condition. The "good" equilibrium is also Pareto efficient.

Our empirical evidence from presented in Section 2, however, suggests that the existing fines and subsidies might not been sufficient in moving the "Glitch" and "Federal" states to the "good" equilibrium. One seemingly obvious fix would be to increase the fine

value. However, this approach may be both inefficient and politically challenging. Instead, Section 6 shows that simply changing the form of the fine—from the ACA's current "absolute" amount that is not based on market premiums to a "relative" amount based on market premiums, as previously taken in Massachusetts—can move these states to the "good" equilibrium. Moreover, this change can be constructed in a way that does not cost non-insured consumers anything more in the good equilibrium than the current absolute fine. A shift from the current "absolute" structure to a relative one can thereby achieve a Pareto improvement.[3]

Section 7 concludes. Proofs are provided in Appendix A. Model extensions are provided in Appendix B.

2 Empirical Patterns

Our model, presented later, adds dynamics to the standard adverse selection model, and demonstrates how adverse selection in the first period after a policy change can competitively persist into future periods. As empirical motivation, this section examines the recent experience of the ACA, which has now been operating for several years. The evidence supports the role of initial conditions predicted by our model.

2.1 The Data

We follow the general data strategy in Kowalski (2014), who examined the initial impact of the ACA at the state level. We extend her analysis to included the full 2014 year and years 2015 and 2016.

State-level data comes from the National Association of Insurance Commissioners (NAIC), as collected by SNL Financial. NAIC collects insurer data on a quarterly basis for enrollment, coverage, premiums, and costs. As with Kowalski (2014), we dropped Massachusetts due to data issues associated with its cross-over from its own individual market exchange system to the ACA. Some health insurers, though, operate across multiple states, and NAIC quarterly data is at the firm level, aggregated across states. For multi-state insurers, therefore, we allocate the insurer-quarter data at the state level using the insurer's annual filings data, which is disaggregated at the state level for regulatory

[3]Hackmann et al. (2015), for example, compute the optimal absolute level of the fine in a model calibrated to the Massachusetts health exchanges. Consistent with the previous literature, their model is static and assumes linear demand and cost curves, thus excluding the possibility of multiple equilibria and any difference between absolute and relative fines.

purposes.[4] In some rare cases, the annual filing for a particular state and year was not captured by NAIC. In those cases, we use the multi-state insurer's Schedule T form, which is filed in each state on a quarterly basis. Relative to annual filings, Schedule T filings have the advantage of being quarterly and distinguishing by state. The disadvantage is that Schedule T filings aggregate the insurer's individual health insurance with its other lines of business, and only include information about premiums (not enrollment, coverage, and costs). Schedule T filings, therefore, are used only as a last resort.

Following Dash and Thomas (2014) and Kowalski (2014), states are divided into three types: (i) "Glitch" states that established their own ACA exchanges but suffered severe technological glitches during the 2013 open enrollment season, immediately before ACA coverage became available in 2014; (ii) "No Glitch" states that established their own exchanges and had no major technological glitches; and, (iii) "Federal" states that used the federal exchange and website, which, as well documented in the national press, suffered from moderate-to-severe glitches during the 2013 open enrollment season.[5]

2.2 Average Costs

Figure 1 shows the weighted average of average costs for individual health plans across "Glitch," "No Glitch," and "Federal" states on a quarterly basis, between March 31, 2014 and June 30, 2016, the last quarter with a stable data release. For each of the three years, bold tick marks highlight the quarter ending in March 31, as the open season for enrollment and new premium rates are set in the previous quarter. March 31, 2014, therefore, represents the first quarter of data impacted by the ACA. March 31, 2015 represents the first quarter of data in the ACA's second year. Similarly, March 31, 2016 represents the first quarter of data in the ACA's third year.

Average costs are computed as total costs in a given state divided by "member months" of coverage provided in that quarter. Because states might differ in average costs for reasons other than technology, the ACA cost experience in each state is normalized relative to its average cost as December 31, 2013, *one quarter* before ACA-impacted data shows up in our data set. The weights for computing the weighted average of the average costs within each of the three state categories are fixed at the relative number of member months in each state as of December 31, 2013.[6]

[4] For 2016, we use the percentages from the 2015 annual filings, since 2016 annual filings are not yet available.

[5] "No Glitch" states included CA, CO, CT, DC, KY, RI, VT, and WA. "Glitch" states include HI, MA, MD, MN, NV, and OR. The remainder of the states are "Federal."

[6] Alternatively, updating weights quarter by quarter, instead of fixing them at their values before the reform, would include the effects of adverse selection from the ACA reform and, therefore, artificially

Figure 1: Average Costs Per Member Month

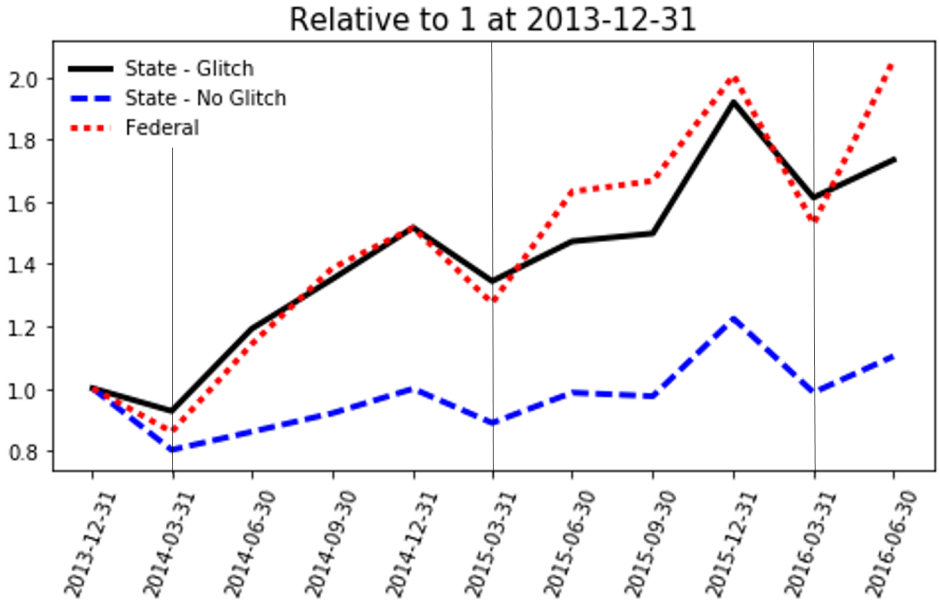

Figure 2: Average Costs Per Member Month

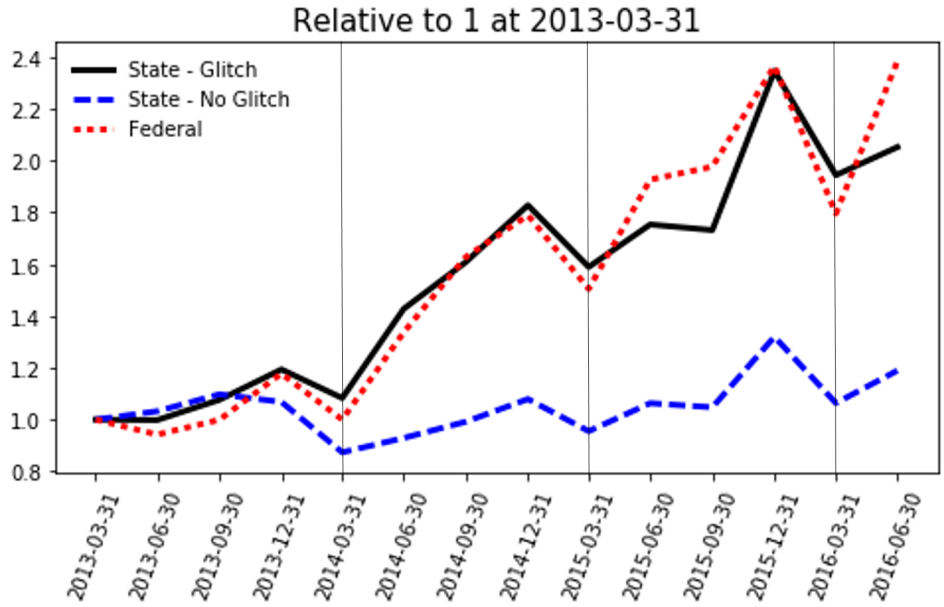

Figure 2 shows the weighted average of average costs for each of the three state categories relative to March 31, 2013, *one year* before the impact of the ACA. Notice that changes in average costs across the three state categories did not vary that much *before* December 31, 2013, suggesting that technological glitches occurring in the fall 2013 where

decrease the weights placed on states experiencing adverse selection after the passage of the ACA.

fairly randomly distributed. Moreover, the subsequent average cost trend lines shown in Figures 1 and 2 are similar, suggesting that the exact comparison date is not a big driver of the analysis. So, in the discussion below, we take the quarter ending on December 31, 2013 as our "base quarter" of comparison, as shown in Figure 1.

Three basic time trends stand out in Figure 1. First, average costs tend to increase in the last quarter of each calendar year and then drop in the first quarter of ACA coverage, ending in March 31. This dynamic is consistent with members rushing to obtain treatment before their current policy expires.

Second, first-year average costs increased substantially more in "Glitch" and "Federal" states relative to "No Glitch" states. By March 31, 2014, average costs were actually lower in all three types of states, by 20% in "No Glitch" states, by 6% in "Glitch" states and by 14% in "Federal" states. The overall reduction is expected since, as just noted, the first quarter tends to have lower costs than the last quarter, and our base quarter is the last quarter of 2013. However, the differential suggests the potential for adverse selection. More importantly, notice that even larger differences begin to emerge during the 2014 calendar year. By December 31, 2014, average costs grew by 51% relative to the base quarter in both "Glitch" and "Federal" states, while remaining unchanged in "No Glitch" states. Intuitively, as noted in quote from *The Economist* in Section 1, a larger hurdle to enroll screens for the sickest members.

Third, Figure 1 shows that the "bad initial condition" in "Glitch" and "Federal" states persisted beyond 2014, carrying into 2015 and then 2016. Relative to the base quarter, by June 30, 2016, average costs rose 73% in "Glitch" states and by 105% in "Federal" states. In contrast, average costs rose by only 10% in "No Glitch" states.

2.3 Average Premiums and Coverage

Figure 3 shows the change of average premiums. Average premiums are calculated symmetrically to average costs, that is, by dividing total premiums collected by member months served. For the quarter ending on March 31, 2014, premiums collected per member were 19% higher in "Glitch" states, 24% higher in "Federal" states and 17% higher in "No Glitch" states. These premiums, which were collected under the ACA, are not directly comparable to premiums collected in the base quarter, prior to the ACA, as many pre-ACA plans offered very basic coverage. However, the similar percent change for all three plans shows that insurers in "Glitch" states and "Federal" states had a limited ability to alter their premiums in response to website glitches during the first year.

Differences in average premiums begin to emerge in future years. For the quarter end-

Figure 3: Average Premiums Per Member Month

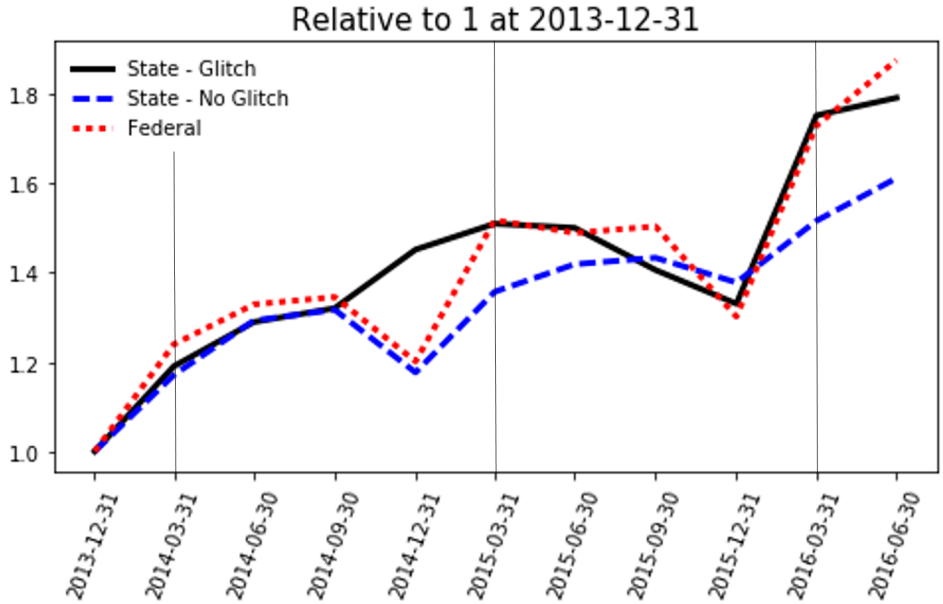

Figure 4: Growth in Member Months

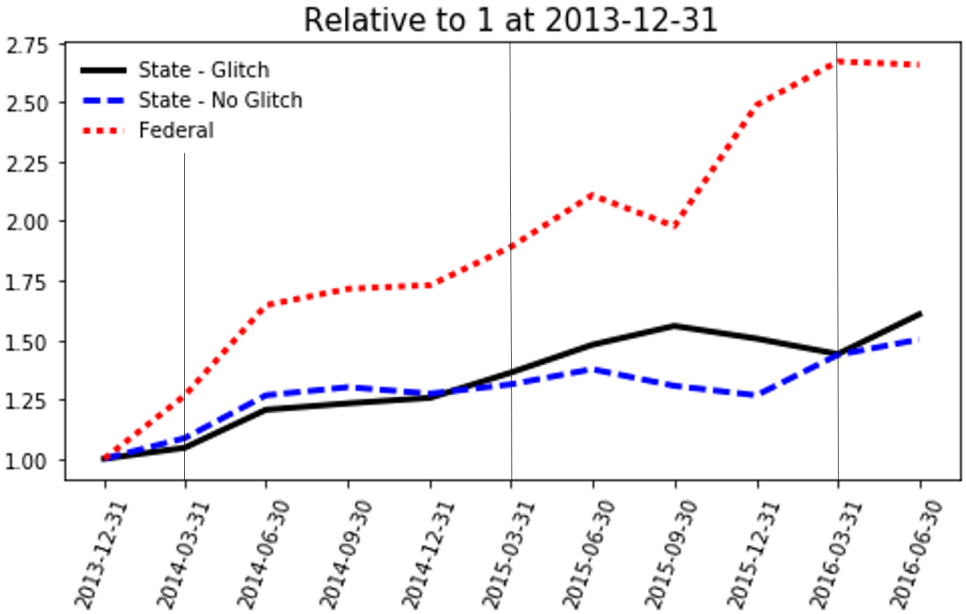

ing on March 31, 2015—the first quarter after which premium rates could be readjusted by insurers—premiums collected per member rose around 51% in both "Glitch" and "Federal" states, but by 36% in "No Glitch" states. By the quarter ending on March 31, 2016, premiums rose 75% in "Glitch" states, 73% in "Federal" states and by 52% in "No Glitch" states. Rate increases are regulated at the state level under the McCarran–Ferguson Act;

rate increases above 10% might also be reviewed by the U.S. Department of Health and Human Services (HHS), if the state does not have an established Effective Rate Review Program. Given the limited data, it is unclear if these review programs have prevented rates in "Glitch" and "Federal" states from increasing even more.[7]

Notice, however, that *in-between* the March 31 tick marks, premiums collected per member month are sometimes similar between the three different types of states. At first glance, these dynamics are surprising because contract prices are fixed during this period and the enrollment period is closed, barring a list of exceptions.

Figure 4, though, shows that member months increased substantially in "Glitch" and "Federal" states, relative to "No Glitch" states, during the closed enrollment period. Apparently, the federal exchange governing the "Federal" states was substantially more generous with allowing people to enroll outside of the open-enrollment period. "Glitch" states also seemed more tolerant as well, especially in 2015.[8] These sharp increases in enrollments outside of the open-enrollment period are consistent with a common complaint made by major ACA insurers, in particular, that they were forced to expand coverage to a large number of "exceptional" cases during the closed-enrollment period, and these cases were sicker on average. Aetna claimed it dropped coverage in 11 of 15 states primarily due to the increase in exceptional cases that lead to losses (USAToday (2016)). UnitedHealth Group completely withdrew from the ACA market for the same stated reason (CNN (2016)), and Humana also pulled all of its coverage (NYTimes (2017)). More recently, for the 2017 season, the Obama Administration announced that it has tightened the rules for signing up for policies outside of open enrollment. On February 15, 2017, President Trump signed an executive order requiring those applying outside of open enrollment to provide documentation proving their eligibility.

Of course, an increase in member months increases the denominator of average premiums, helping to reduce their values in-between the March 31 tick marks. But why did the numerator (premiums collected) not increase proportionally? One likely reason is that people who waited until the closed enrollment period to sign up tend to be poorer. Under federal law, health services can typically be obtained for a period up to 90 days

[7]Premiums were also potentially distorted due to early promises of "risk corridors" during the first three years of ACA coverage, designed to share risks ex-post across insurer experiences. Congress "de-funded" the risk corridors in December, 2014. However, the Obama Administration maintained that Congressional appropriation was not necessary. The issue remains in the court system, although it is generally believed that the Trump Administration will drop the White House's petition.

[8]The generosity of the federal exchange has been well documented, including in the newspaper articles referenced later. Interestingly, all of the "Glitch" states lean heavily Democratic, which could explain their generosity relative to "No Glitch" states.

before premiums are paid.[9] In some cases, premiums were never paid.[10]

3 Model with Strategic Insurers

In the next three sections, we now turn to possible theoretical foundations for the patterns described in the preceding section. For the sake of expositional simplicity, we start by considering a simple insurance model that incorporates many of the key features highlighted in Akerlof (1970), Wilson (1980), Wilson (1977), Einav et al. (2010b), and Einav and Finkelstein (2011). Our most parsimonious model assumes a continuum of risk types, that risk is the only source of consumer heterogeneity, and losses are binary. In Appendix B, we demonstrate that our key results extend to a model setting with discrete risk types, richer forms of heterogeneity and multiple loss sizes.

3.1 Consumers

A unit measure of consumers have wealth $w > 0$ and face a potential loss of size $0 < l < w$ in the presence of limited liability. Consumers only differ in the probability $\pi \in [0, 1]$ of the loss occurring, which is distributed throughout the population by the continuous cumulative distribution function $H(\pi)$ with support $[0, 1]$.[11] Let the random variable Π be H-distributed, and denote a realization of Π by π. Agents are risk-averse with a concave Bernoulli utility function $u(c)$ over consumption, so the expected utility of type π is given by

$$\pi u(w - l) + (1 - \pi)u(w)$$

when there is no insurance.

We assume that individuals can choose from exactly two available insurance contracts that differ exogenously in how much of the loss l they cover. Following Einav et al. (2010b), and without loss of generality, we normalize the low coverage contract to be no

[9] Insurers must cover medical bills for the first 30 days of unpaid premiums. Over the next 60 days, insurers may "pend" payments, although the patient can usually still obtain care (KHN (2014)).

[10] The exact amount of unpaid premiums is hotly contested. A report issued by The Energy and Commerce Committee (2014) of the U.S. Congress argued that only 67% of sign-ups paid premiums by April 15, 2014, two weeks after the first quarter of coverage was provided. However, a pro-ACA group, ObamaCare Facts (2014) asserts that some of the plans started providing services on May 1, 2014, and probably 80% to 90% of premiums were likely eventually paid. Regardless, there is a significant potential for underpayment relative to services provided.

[11] The full support assumption can be viewed as a limiting case where, as in Hendren (2014), even the most extreme risk types, for whom the loss never or always occurs, exist with arbitrarily small but positive density. None of our substantive results depend on this assumption. Appendix B relaxes this and other assumptions, as noted earlier.

insurance at a zero premium, and the high coverage contract to be full insurance at some endogenous premium p. Abstracting from moral hazard, we take l and each individual's risk π as exogenous and, therefore, independent of the insurance choice. The demand for insurance, therefore, is only a function of the price p.

The assumption of fixed coverage levels places our analysis in the spirit of Akerlof (1970) rather than Rothschild and Stiglitz (1976) who endogenize coverage levels as well.[12] As discussed in Einav et al. (2010b) and Einav and Finkelstein (2011), this assumption is a reasonable characterization of many insurance markets. It becomes an even more appropriate assumption for the ACA health exchanges which place regulatory bounds on minimum coverage, despite allowing for a range of plans that differ in copayments made by consumers (see Handel et al., 2015, for a model of insurance markets with two fixed (non-zero) coverage levels, covering 90% and 60% of an individual's cost, respectively). This modeling decision is also most natural to analyze one of the main policy interventions of the ACA, namely the penalty for not having insurance, which affects the demand for health insurance on the extensive margin and which we consider in Section 6.

3.2 Insurers

Following Einav et al. (2010b) in the insurance market and Mas-Colell et al. (1995) in a labor market setting, suppose there are at least two identical, risk-neutral insurers that maximize their respective expected profits by setting premiums in a two-stage Bertrand game. In the first stage, insurers simultaneously announce their premiums. In the second stage, individuals decide whether to purchase insurance and, if so, from which insurer.[13] In Section 4, we demonstrate how the results change when insurers act instead as price-takers, as in Akerlof (1970).

We assume throughout that an individual's risk type π is private information, so insurers cannot offer different premiums to different individuals. Even if insurers could observe risk types, the ACA does not permit pricing based on pre-existing conditions. One can, therefore, think of our analysis as applying to a set of individuals who are otherwise identical in terms of characteristics that insurers are allowed to price, such as smoking status.

[12] See, e.g., Netzer and Scheuer (2013) for a recent treatment.

[13] As usual, to break a tie (since actual currency denominations are technically a countable set to the penny level), if multiple insurers announce exactly the same premium levels, individuals then randomize among them with equal probabilities.

3.3 Strategic Equilibrium

To characterize the set of subgame perfect equilibria of this game, a graphical representation following Einav and Finkelstein (2011) is useful. For any critical buyer $\pi \in [0,1]$, the *average cost* of insuring everyone with risk equal to or greater than π is

$$\Gamma(\pi) \equiv \mathbb{E}[\Pi|\Pi \geq \pi]l. \tag{1}$$

Our assumptions ensure that $\Gamma(\pi)$ is continuous, increasing in π, and satisfies $\Gamma(0) = \mathbb{E}[\Pi]l$ and $\Gamma(1) = l$. In words, when $\pi = 0$ is the critical type, the average cost of the entire population is just the unconditional expected loss. On the other hand, with critical type $\pi = 1$, their losses are certain, and so their expected loss is simply l.

On the demand side, we can define the *willingness to pay* $\Omega(\pi)$ for insurance of each type π implicitly by solving

$$u(w - \Omega) \equiv \pi u(w - l) + (1 - \pi)u(w). \tag{2}$$

Since the right-hand side is decreasing in π, there is a unique solution $\Omega(\pi)$ for each π, which is also continuous, increasing in π, and satisfies $\Omega(0) = 0$, $\Omega(1) = l$.[14] In words, the lowest risk type $\pi = 0$ never experiences a loss and, therefore, has no willingness to pay for insurance. In contrast, the highest risk type $\pi = 1$ experiences the loss l for sure and is, therefore, willing to pay a premium up to l. We can also interpret $\Omega(\pi)$ as an *inverse demand curve*: with a premium $p = \Omega(\pi)$, insurance will demanded by all types higher than π (so that the inverse function $\Omega^{-1}(p)$ identifies the marginal buyer when the premium is p).

For the purpose of this section, we assume that there is common knowledge of the distribution $H(\pi)$, consumer wealth w, the loss amount l, and the form of utility $u(c)$. In other words, the shapes of $\Gamma(\pi)$ (the average cost curve) and $\Omega(\pi)$ (the willingness to pay curve) are common knowledge among insurers.

The following proposition, which is easily adapted from Mas-Colell et al. (1995) and Einav et al. (2010a), shows that there is a unique equilibrium in this model.

Proposition 1. *With strategic insurers, the unique subgame perfect equilibrium outcome of the above two-stage game involves the critical type*

$$\pi_1^* = \min\{\pi \in [0,1]|\Omega(\pi) = \Gamma(\pi).\}.$$

[14] As is standard, we are assuming that the loss size l does not exceed available wealth w. As discussed more in Section 5, in the presence of limited liability, it is possible for actual losses to exceed wealth for some types, and so $\Omega(\pi)$ is only weakly increasing in π and $\Omega(1) < l$.

Figure 5: Equilibrium with Complete Unraveling

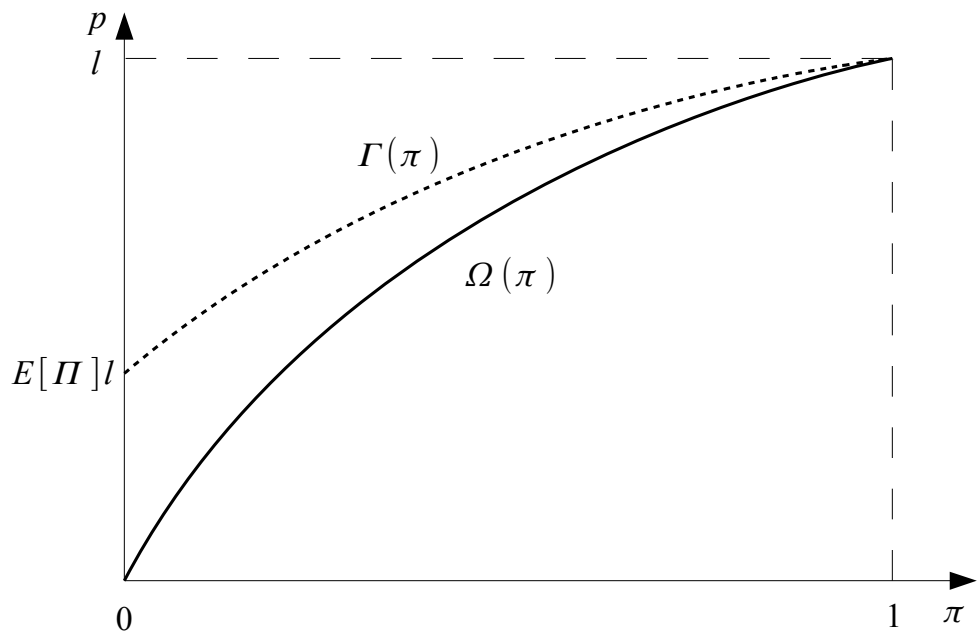

In words, when insurers set premiums strategically, only the intersection of the average cost and the inverse demand curve with the lowest premium and the most people covered is an equilibrium.[15] Hence, there cannot exist multiple equilibria and, therefore, this model does not allow for any role of the type of "initial conditions" discussed in the Introduction. This result holds even when there are multiple intersections of the average cost and inverse demand curves (which, as we will show in the next section, would each correspond to an equilibrium in a price-taking model).

A simple illustration is provided in Figure 5. Note that, by risk aversion, $\Omega(\pi) \geq \pi l$, so the inverse demand curve must always lie above the diagonal line $p = \pi l$. Obviously, we have $\Gamma(1) = \Omega(1)$, so there always exists an intersection between these two curves and hence an equilibrium. Here, nobody buys insurance except for the very highest risk types with $\pi = 1$, who are just indifferent between buying or not buying when faced with the fair premium $p^* = l$ for this pool. In the situation depicted in Figure 5, this outcome is, in fact, the only intersection between demand and average costs, corresponding to the case of complete unraveling emphasized in Akerlof (1970). Specifically, for any $\pi < 1$, the average cost curve is above the demand curve, so insurers would make losses at any premium $p < l$.

[15]For the remainder of the paper, we confine attention to the generic case where all intersections are proper intersections rather than tangency points of the two curves.

The fact that condition $\Omega(1) = l = \Gamma(1)$ holds in this baseline model is not required for an equilibrium to always exist. In Section 5, we consider the role of limited liability where the willingness to pay at the very highest risk ($\pi = 1$) is lower than the average cost. There always exist a "corner" equilibrium where $p = \Gamma(1)$ and nobody buys insurance. Conversely, as in the model extension in the Appendix, if we had $\Omega(1) > \Gamma(1)$, then there would always exist an (interior) equilibrium, due to the continuity of Γ and Ω and the fact that $\Omega(0) < \Gamma(0)$.

More interestingly, when the $\Omega(\pi)$- and $\Gamma(\pi)$-curves cross multiple times, why does only the intersection with the lowest premium correspond to an equilibrium in the strategic model? To understand this, consider a setting with three intersections. Figure 6 illustrates the mechanics when firms behave strategically. Suppose we are in the worst intersection with critical type $\pi = 1$ and premium $p = l$.[16] *If* all insurers set premium $p = l$, only types $\pi = 1$ demand insurance, and the average cost of this pool is $\Gamma(1) = l$, and so all insurers make zero profits. Hence, as we will formalize below, this outcome is a *competitive* equilibrium in a price-taking model. But it also emits a profitable deviation by any *strategic* insurer. In particular, suppose that an insurer deviates and sets a premium $p_2' < p_2^*$. As drawn in Figure 6, this insurer will capture the entire market with demand from all types $\pi \geq \pi_2'$ (and observe $\pi_2' < \pi_2^*$). Moreover, at π_2', the average cost curve is below the demand curve, so $\Gamma(\pi_2') < p_2'$. Hence, offering the premium p_2' will result in strictly positive profits for the deviating insurer corresponding to the dashed area in Figure 6. The only intersection from where there is no such profitable deviation is the one with marginal buyer π_1^* and the lowest premium p_1^*.

3.4 Is Strategic Pricing Realistic?

Strategic premium setting in a framework of Bertrand-like competition may seem like the more relevant case than price-taking for insurance markets, since many insurers are not atomistic and do actively set premiums taking into account their competitors' and customers' responses to their actions. However, as e.g. Mas-Colell et al. (1995) emphasize in the setting of labor markets with adverse selection, the outcome in Proposition 1 relies on the assumption that firms have common knowledge about all market fundamentals, including the *global* shape of the demand and cost curves Ω and Γ. In contrast, in Sections 4 to 6, we relax this assumption and only assume that insurers know the average cost of those who buy insurance *at the going premium*; they may not know anything about the preferences or risk distribution underlying this equilibrium or have non-localized

[16] The same argument could be made about the intermediate equilibrium with π_2^* and p_2^*.

Figure 6: Equilibrium Uniqueness with Strategic Insurers

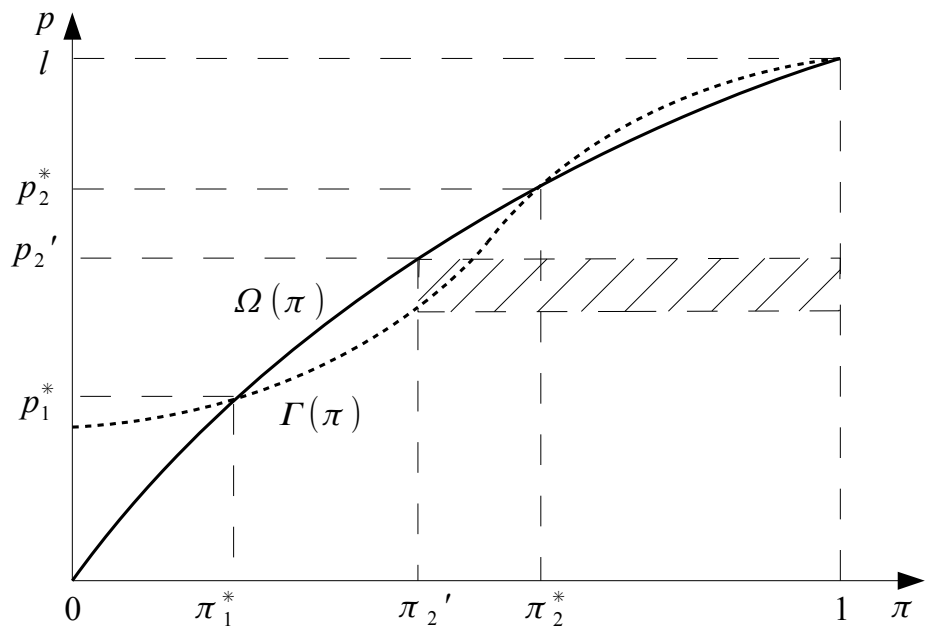

knowledge *away* from current conditions.

The strong information requirements in the strategic model are crucial to obtain equilibrium uniqueness. In particular, a *mistaken* attempt at a profitable deviation could lead to substantial losses. For instance, suppose again we start from the intersection with $\pi^* = 1$ in Figure 6. If an insurer deviates by offering a marginally lower premium $p = l - \varepsilon$, this will lead to losses since $\Gamma(\pi) > \Omega(\pi)$ for π close to one. To make profits, a deviating insurer would have to offer a discretely lower premium $p < p_2^* < l$, but also not too low, since losses would be incurred again if $p < p_1^*$. In Figure 6, the demand and cost curves are drawn such that the interval of profitable premium deviations (p_1^*, p_2^*) is still relatively large. However, this need not be the case. Figure 7 depicts market fundamentals where this interval is very small and far away from the going premium $p = l$. In this case, insurers would actually incur losses for large range of premium cuts in (p_2^*, l), and only make profits if they reduce premiums by a very large (and just the right) amount until $p \in (p_1^*, p_2^*)$. In other words, Proposition 1 requires that insurers have precise knowledge about market conditions potentially far away from the current situation.

Strategic insurers could potentially limit their losses if they could rapidly change prices to try to discover the global shapes of the willingness-to-pay and average cost curves. However, transitory losses from pricing experimentation are magnified by the fact that regulations prevent insurers from changing prices frequently. Once set, ACA

Figure 7: Small and Distant Interval of Profitable Premium Deviations

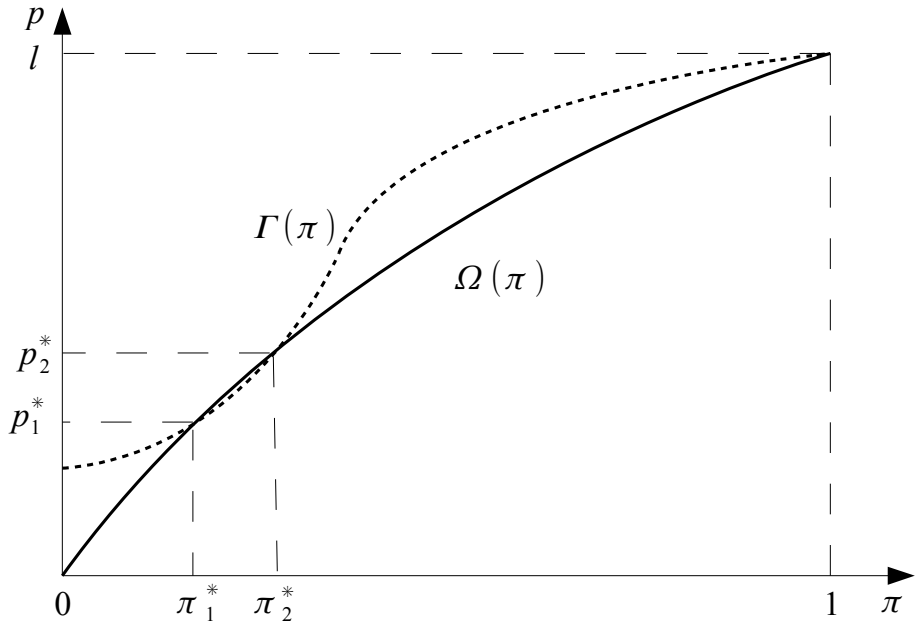

plan premiums are generally locked until the next open enrollment period (Kaiser Family Foundation (2013)).

Even then, sharply increasing premiums after a pricing mistake is challenging. Under the McCarran-Ferguson Act of 1945, individual states typically regulate the business of insurance, and most states already require some steps before rates can be increased (National Conference of State Legislatures 2013).[17] However, because rules vary between states, Title I (Subtitle A, Sec. 1003) of the ACA creates a more uniform standard around rate increases. These rules include requiring states to collect premium information and determine if plans should be excluded from the health exchange based on unjustified premium increases.[18] If an insurer requests a premium increase above 10%, a more detailed explanation must be provided and posted on their and the HHS website. The ACA also makes $250 million available to states to take action against insurers requesting unreasonable rate increases. According to the Centers for Medicare & Medicaid Services (2010), "[t]his funding will help assure consumers in every state that any premium increases requested by their insurance company, regardless of size, is justified."

[17] At least two dozen states require that the insurer receives prior approval from the state insurance commissioner or department before increasing health insurance premiums (National Conference of State Legislatures 2013).

[18] For a few states — Alabama, Louisiana, Missouri, Oklahoma, Texas, and Wyoming — these determinations will be made by the federal government since these states do not have review processes in place.

In the absence of perfect information about the market structure, insurers, therefore, may simply prefer local adjustments to premiums in a backward looking manner, as documented in Cutler and Reber (1998). Such behavior would effectively make them price-takers. Consumer-protection laws intended to protect consumers from frequent and large price increases could undermine experimentation and essentially force insurers into price-taking behavior. We demonstrate in the following sections that equilibrium multiplicity can arise naturally with price-taking insurers, including the potential of getting stuck in a bad equilibrium with low coverage and high prices.[19]

4 Model with Price-Taking Insurers

We now turn to the notion of competitive equilibrium where insurers act as price-takers, as originally proposed by Akerlof (1970). In particular, we ask under what conditions multiple equilibria are possible, and how this potential for multiplicity shapes the role of initial conditions in a dynamic framework. Of course, each consumer's risk level π is still private information. But we can now relax the assumption of common knowledge of the global average cost and demand curves. Instead, we only require that insurers know the average cost of those who buy insurance at the current premium. We assume that there are many identical insurers that take the market premium as given and decide whether to offer insurance at that premium or not.

4.1 Competitive Equilibria

The following definition of a competitive equilibrium is consistent with the informational assumptions outlined above:

Definition 1. *With unobservable risk types π, a competitive equilibrium is a premium p^* and a critical type π^* such that*

$$u(w - p^*) \geq \pi u(w - l) + (1 - \pi)u(w) \;\; \forall \pi \geq \pi^*, \tag{3}$$

$$u(w - p^*) < \pi u(w - l) + (1 - \pi)u(w) \;\; \forall \pi < \pi^* \tag{4}$$

and

$$(1 - H(\pi^*))p^* = \int_{\pi^*}^{1} \pi dH(\pi) l. \tag{5}$$

[19]See also Rothschild (1974) for the classic model on experimentation to learn about demand conditions. These two-armed bandit models also have the typical feature that multiple equilibria can arise.

The first two conditions characterize consumers' demand for insurance, given the equilibrium premium p^*. At that premium, individuals of risk type $\pi \geq \pi^*$ are just indifferent or strictly prefer to buy insurance, whereas all other types $\pi < \pi^*$ prefer to stay uninsured. The third condition then requires insurers to make zero profits at the policy premium p^* on the pool of risk types who demand insurance when the premium is p^*, which includes all types $\pi \geq \pi^*$. In particular, the left-hand side of equation (5) equals the total premium revenue collected from these agents while the right-hand side is equal to their expected losses. This zero profit condition can be simply rewritten as $p^* = \mathbb{E}[\Pi|\Pi \geq \pi^*]l$, i.e. the equilibrium premium must equal the expected loss of the pool of insurance buyers induced to buy the policy.

In terms of the graphical representation introduced in Section 3, a competitive equilibrium, therefore, is simply any π^* such that $\Gamma(\pi^*) = \Omega(\pi^*)$, so that the average cost and willingness-to-pay curves intersect. Hence, the strategic equilibrium from Proposition 1 is also a competitive equilibrium, but there may now exist additional competitive equilibria.

4.2 Equilibrium Multiplicity

Since most of the insurance literature has implicitly focused on either linear demand or cost curves (see, e.g., Cutler and Reber (1998) and Einav and Finkelstein (2011)) or strategic insurers (Section 3 and Einav et al. (2010a)), the possibility of multiple competitive equilibria has not received much attention in this context.[20] As a result, there is no real role for the type of "initial conditions" discussed in Section 1, including a website failure.

Equilibrium multiplicity, however, can arise naturally in the price-setting framework because the average cost and demand curves are upward sloping, but their shapes are otherwise largely unrestricted. A simple example is depicted in Figure 8. There are three competitive equilibria in total, namely the one with unraveling located at $\pi^* = 1$ as well as two additional equilibria with critical buyers π_1^* and π_2^*.

As is well known in other contexts (see e.g., Wilson (1980) for a lemons goods market and Mas-Colell et al. (1995) for a labor market model with unobservable productivities), whenever there are multiple equilibria, they are Pareto ranked: The equilibrium at π_1^* is Pareto better than the equilibrium at π_2^*, which is Pareto better than the equilibrium with complete unraveling at $\pi^* = 1$. For example, compare π_1^* against π_2^*. In the equilibrium at π_1^*, all types $\pi \geq \pi_2^*$ are better off than in the equilibrium at π_2^* because they pay

[20] See also Hackmann et al. (2015), Handel et al. (2015) and Mahoney and Weyl (forthcoming) for recent studies of health exchanges in models with a unique equilibrium.

Figure 8: Multiple Competitive Equilibria

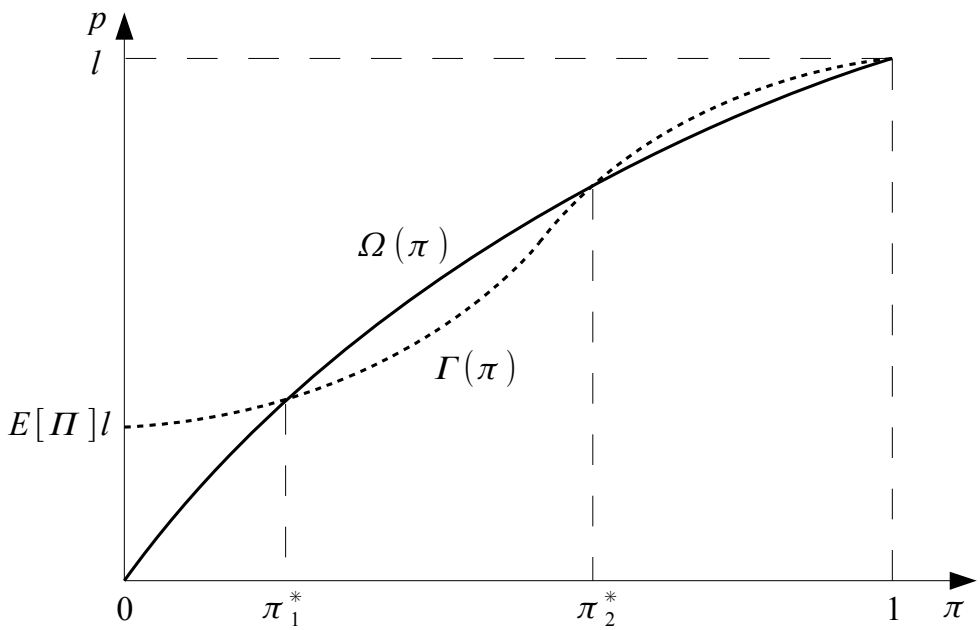

a lower premium for their insurance (i.e., $\Gamma(\pi_1^*) < \Gamma(\pi_2^*)$). Moreover, the types $\pi \in [\pi_1^*, \pi_2^*)$ prefer to buy insurance at π_1^* rather than staying with their endowment, which would be their choice at π_2^*. So, they are also better off. The types $\pi < \pi_1^*$ are indifferent because they do not buy insurance in either case. Moreover, insurers earn zero profits in all equilibria.

Indeed, the "good" equilibrium π_1^*—which offers the lowest premium and entices the most consumers to buy insurance—Pareto dominates all the others.[21] The other equilibria arise because of a *coordination failure*: When only a few individuals purchase insurance, they will be the highest risk types, and so the premium that breaks even for this pool will also be high. At the same time, only the riskiest types find it worthwhile to sign up for insurance because the premium is so high.

In the next section, we will investigate the nature and policy relevance of equilibrium multiplicity in more detail, both in terms of the conditions on primitives that are give rise to it, and in terms of their empirical plausibility.

[21] In fact, the good equilibrium is constrained Pareto efficient under the restriction to full insurance contracts.

5 Equilibrium Multiplicity and Initial Conditions

The previous literature examining multiple equilibria has not explicitly distinguished between stable and unstable equilibria. For policy purposes, only stable equilibria are material. Moreover, initial conditions do not matter with multiple equilibria unless at least two are stable. Generating two stable equilibria requires having at least three equilibria in total. Initial conditions do not matter if there are only two equilibria. In this section, we will provide sufficient conditions on fundamentals under which multiple stable equilibria exist, and explore their relevance quantitatively.

5.1 Introducing Dynamics

To study the circumstances under which initial conditions could affect which competitive equilibrium is reached, a dynamic version of this static model is required. The most straightforward way to introduce dynamics is to assume that, in each period, premiums reflect the average cost of the pool of individuals who purchase insurance, thereby allowing insurers to always break even. Then, given this premium, consumers decide whether to enroll for insurance the next period. Cutler and Reber (1998), Monheit et al. (2004) and Clemens (2015) provide evidence for this pattern of price and demand adjustments.

These dynamics can be conveniently illustrated graphically using the same type of diagram as before. Recall that, for any critical type π_t in period t, we can read the premium from the average cost curve by setting $p_t = \Gamma(\pi_t)$. The consumers' reaction in $t+1$, therefore, can then be read off the demand curve to obtain a new marginal buyer $\pi_{t+1} = \Omega^{-1}(p_t)$, and so forth. This leads to the recursion $\pi_{t+1} = \Omega^{-1}(\Gamma(\pi_t))$ for the evolution of marginal buyers, as illustrated in Figure 9. It immediately implies that π increases (i.e. there is unraveling where the premium increases and fewer consumers sign up for insurance) whenever $\Gamma(\pi) > \Omega(\pi)$ while π falls otherwise (more consumers demand insurance, so the premium falls).

We can see that, of the three competitive equilibria here, only two are stable, whereas the intermediate one with the marginal buyer π_2^* is unstable. Which competitive equilibrium is eventually reached depends on the initial value of π, as formalized in the following proposition:

Proposition 2. *(i) Initial conditions $\pi \in [0,1)$ matter for which competitive equilibrium is reached only if there exist at least three competitive equilibria.*
(ii) When there are exactly three equilibria with critical types $\pi_1^* < \pi_2^* < 1$, the intermediate equilibrium with marginal buyer π_2^* is generically unstable while the other two are stable.

Figure 9: Dynamics and Equilibrium Stability

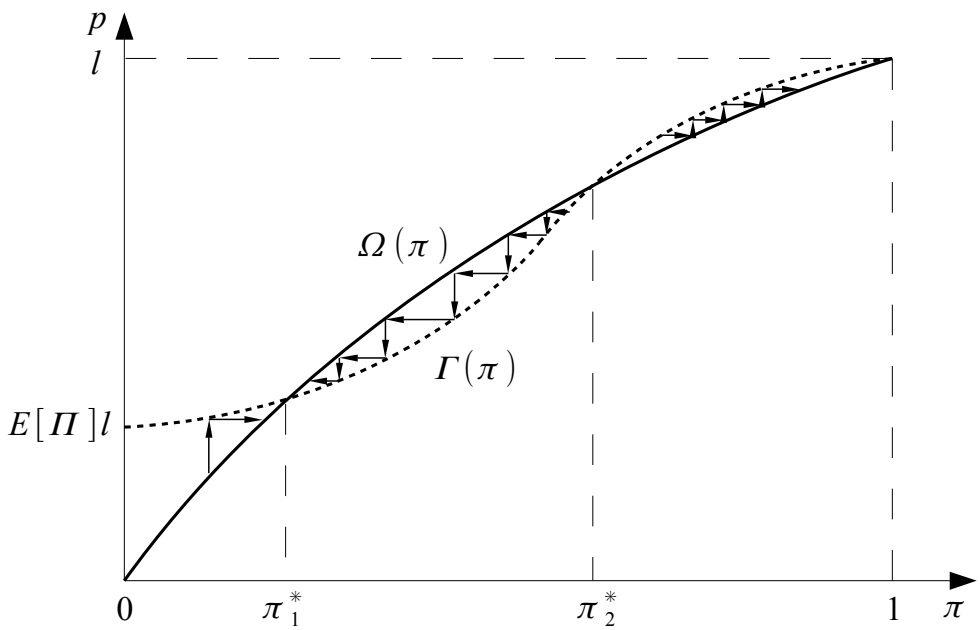

(iii) In this case, π_2^* is the critical threshold for initial conditions: for any initial $\pi > \pi_2^*$, there is unraveling to the "bad" stable equilibrium where $\pi^* = 1$. For any $\pi < \pi_2^*$, the "good" stable equilibrium with critical type π_1^* is reached.

Proof. See Appendix. □

While the possibility of equilibrium multiplicity per se is not surprising in this model, Proposition 2 formalizes the less immediate result that at least three competitive equilibria are required to obtain at least two *stable* equilibria, which are the only type of equilibria that are relevant for policy purposes. Figure 10 illustrates why the existence of just two equilibria is not enough for initial conditions to matter. In this case, the best equilibrium is always the (unique) globally stable equilibrium and so convergence to the bad equilibrium, which is *unstable*, cannot occur. As discussed earlier, Proposition 2 (i) does not depend on the fact that the condition $\Omega(1) = \Gamma(1)$ implies that $\pi = 1$ is always an equilibrium. In particular, we would have a "corner" equilibrium at $\pi = 1$ if $\Omega(1) < \Gamma(1)$, or an interior equilibrium when $\Omega(1) > \Gamma(1)$. Instead, at least one intermediate, unstable equilibrium is always required in order to produce two stable equilibria and, hence, for initial conditions to matter.

With exactly three equilibria as depicted in Figure 9, unraveling occurs starting from initial conditions to the left of π_1^* (a partial unraveling) and to the right of π_2^* (a full un-

Figure 10: Globally Stable Equilibrium with Two Equilibria

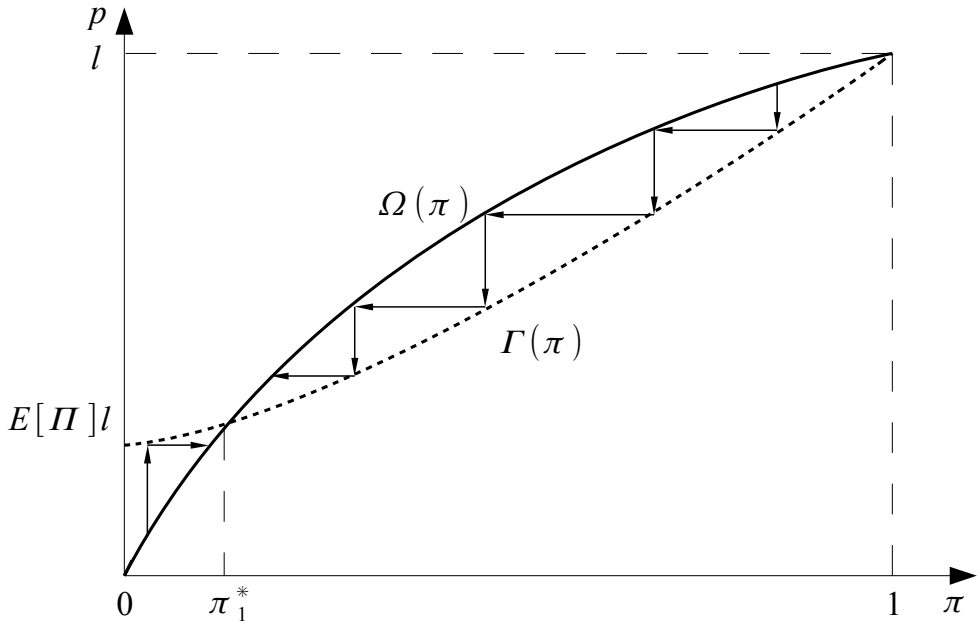

raveling), whereas the dynamics imply falling premiums and more individuals enrolling otherwise. Evidence for such dynamics have been documented in states which, before the ACA, placed restrictions on adjusting premiums based on age and preexisting conditions. Writing about the New Jersey Individual Health Coverage Program (IHCP) the began in 1993, Monheit et al. (2004) found dynamics similar to those shown in Figure 9 for values of $\pi > \pi_2^*$ (or $\pi < \pi_1^*$). In particular, between the end of 1995 and the end of 2001, enrollment fell from 186,130 individuals to just 84,968, with premiums rising by 200% to 300%. Three other states — Kentucky, New York and Vermont — tried health care reforms with similar consequences (Cohn (2012)). Clemens (2015) provides a comprehensive analysis of the effect of the introduction of community rating regulations in these states and finds that the fraction of uninsured gradually increased by around 70%, from 18% to 31%, in the three years following the reforms. Cutler and Reber (1998) provide evidence for gradual unraveling of high coverage plans in a setting with employer provided insurance.

As widely reported in the popular media reports, the initial conditions described in Section 1 would quite reasonably discourage lower-risk consumers from enrolling in the health exchanges relative to higher-risk consumers. In the context of our model in this section, we would expect consumers with large values of π to be the first to enroll, potentially trapping the system in the bad stable equilibrium. Of course, by Proposition 2, these mechanics only matter if multiple equilibria actually exist in the first place, a topic

to which we now turn.

5.2 Sufficient Conditions for Multiple Stable Equilibria

We now provide the sufficient conditions under which multiple stable equilibria must exist. In particular, if risk-aversion is not too high, then the bad equilibrium at $\pi^* = 1$ exists and is stable. Moreover, if the distribution of risk types is sufficiently concentrated, then there must exist at least one other stable equilibrium with $\pi^* < 1$. The following proposition formalizes these conditions.

Proposition 3. *Suppose that*

1. $$\frac{u(w) - u(w-l)}{l} > \frac{u'(w-l)}{2}, \tag{6}$$

   which holds whenever u is not too concave or l is sufficiently small, and

2. *$H(\pi)$ is sufficiently concentrated, i.e. there exists some interval $[a, a+\Delta]$ with $0 < a < a + \Delta < 1$ such that $\int_{[a,a+\Delta]} dH(\pi) \geq 1 - \varepsilon$ and both Δ and ε are sufficiently small.*

Then there exist at least three competitive equilibria, including a stable one at $\pi^* = 1$, another stable one at some $\pi^* < 1$ and an unstable one in between.

Proof. See Appendix. □

Intuitively, the inequality in (6) puts an upper bound on the willingness to pay for insurance, ensuring that the inverse demand curve is steeper (and, hence, located below) the average cost curve to the left of $\pi = 1$. As a result, the worst equilibrium is locally stable. Condition 2., in turn, implies that the average cost curve is S-shaped and, in particular, must be located below the inverse demand curve for some interior π. Because $\Omega(0) = 0 < \Gamma(0)$ and by continuity of the $\Omega(\pi)$- and $\Gamma(\pi)$-curves, there must exist another stable equilibrium with $0 < \pi < 1$.

While we have abstracted from limited liability concerns so far by assuming $l < w$, this issue will emerge in some of our calibrations below, as well as in the heterogeneous loss-levels extension in Appendix B. The following corollary of Proposition 3 shows that limited liability can also ensure equilibrium multiplicity:

Corollary 1. *Suppose individuals are liable only for losses up to $L < l$, where L is not too small. Then, condition 2. in Proposition 3 alone is sufficient for the existence of at least three competitive equilibria, including a stable one at $\pi^* = 1$, another stable one at some $\pi^* < 1$ and an unstable one in between.*

Proof. See Appendix. □

The intuition is that, similar to bounded risk aversion, binding limited liability constraints reduce the willingness to pay for insurance for high risk types, thus ensuring the existence of a stable bad equilibrium where nobody buys insurance. It is also worth emphasizing that these conditions are sufficient but not necessary, i.e. multiple stable equilibria may arise even in the absence of these conditions.

Taken together, these conditions suggest that multiple stable equilibria could easily—if not quite likely—arise with realistic parameter values for the average cost and willingness-to-pay functions. The next subsection presents simulation evidence supporting multiplicity under a risk distribution calibrated from U.S. data.

5.3 A Calibration based on the Medical Expenditure Panel Survey

To examine the potential for the type of multiple equilibria shown in Figure 8, we now present a simple quantification of the model. We start with a simple baseline scenario where household types differ in their probability of an identical loss. For robustness, we then consider several variations, including different wealth constructions as well as focusing on self-employed and uninsured households. Appendix B presents a calibration based on the case where household types face the same loss probability across different loss amounts. The key lessons are the same in both sets of calculations.

5.3.1 Baseline Scenario

In our baseline scenario, we use data from the 2010 Medical Expenditure Panel Survey (MEPS) for the pre-Medicare population (ages 18 - 64) to calibrate the average cost (AC) and demand curves.[22] The calculations, therefore, correspond to the pre-ACA population to avoid conflating the potential impacts from the ACA itself.

Figure 11 shows the average cost curve as well as the willingness-to-pay curves at different levels of risk aversion. The horizontal axis corresponds to the top X% percent of spenders, where X (the "rank") is the shown value.[23] The vertical axis is denominated in

[22] As verification, we used data kindly provided by Cohen and Uberoi (2013), and we also directly did our own analysis on the 2010 MEPS data for robustness.

[23] Notice that the horizontal axis in Figure 11, denoted in X%, has the same ordering as the horizontal axis in the previous figures, denoted in π. Rightward movements in both imply greater risk. However, the support itself is now bounded above zero. In particular, the left-most point of 100% in Figure 11 now corresponds to a willingness to pay that is greater than zero since the average person in the bottom 50% of spenders now faces a chance of loss greater than zero. In contrast, in the previous figures, the left-most point of $\pi = 0$ corresponded to a person who faced no chance of loss.

dollars. The average cost curve simply sorts medical spenders by percentile. For example, the mean health expenditure per person in the top 100% of the population (i.e., the entire population mean) is equal to \$3,844, increasing to \$7,476 for population in the top 50%, and climbing to \$38,147 for the top 5%.

Of course, an important question is how much of this (ex-post) heterogeneity corresponds to (ex-ante) private information of individuals. On one extreme, the cost distribution could entirely result from ex-post risk, where all individuals have the same expected costs and, therefore, no private information. In this case, the AC curve would be flat and there would be no adverse selection. On the other extreme, all of the distribution could be driven by heterogeneous individuals with private information about their (deterministic) health expenditures. Instead, we take an intermediate stance that is closest to our formal model. In particular, we distinguish quantiles including the top $X\%$ of spenders, with $X \in \{5, 10, 20, 25, 30, 50, 100\}$ and assume that individuals only have private information about which of these bins they belong to.[24]

We calibrate the marginal willingness to pay for insurance for each of the shown percentiles as follows. First, we assume throughout a constant relative risk aversion utility function $u(c) = c^{1-\alpha}/(1-\alpha)$, where α is the level of risk aversion. Second, for this calibration, the constant loss value l is derived from equation (1) by using the average cost of the top 5% of the population from the MEPS and setting with $\pi_{5\%} = 1$ for them. Third, given this fixed loss value, the (marginal) value of π is then calculated recursively (from the top) at each value of $X\%$ by solving equation (1) for π.[25] (Hence, the value of π increases as the shown value of $X\%$ decreases.) Finally, for a given value of α, the demand curve is then calculated by solving equation (2) for the value of Ω for each value of π, and hence $X\%$. The value of wealth w in equation (2) is initially set equal to the median net worth found in the 2010 Survey of Consumer Finances (Board of Governors (2012)), which assumes that the probability of a loss is independent of the household's wealth.

Consider first the case of $\alpha = 3$. Notice that the willingness-to-pay curve is always smaller than the AC curve (i.e., $\Omega(X\%) < \Gamma(X\%)$ except at the top rank—the smallest shown value of $X\%$—where both curves join). This outcome corresponds to the full unraveling case shown previously in Figure 5. Intuitively, at this comparatively small level of risk aversion, agents with a smaller loss probability π (located at *larger* values of $X\%$ on the horizontal axis) are willing to forgo insurance, pushing up its average cost, thereby leading to unraveling as the value of $X\%$ gets smaller. Now, consider the case of $\alpha = 5$.

[24] We have performed robustness checks with even fewer bins and found similar results.

[25] For instance, to compute the probability $\pi_{10\%}$ of the loss for the top 10-5% of spenders, we solve $(0.5\pi_{10\%} + 0.5\pi_{5\%})l = AC_{10\%}$, where we take the average costs $AC_{10\%}$ and $l = AC_{5\%}$ from the MEPS data and set $\pi_{5\%} = 1$. $\pi_{20\%}$ is then obtained from solving $(0.5\pi_{20\%} + 0.25\pi_{10\%} + 0.25\pi_{5\%})l = AC_{20\%}$, etc.

Figure 11: Willingness to Pay and Average Costs: Median Net Worth

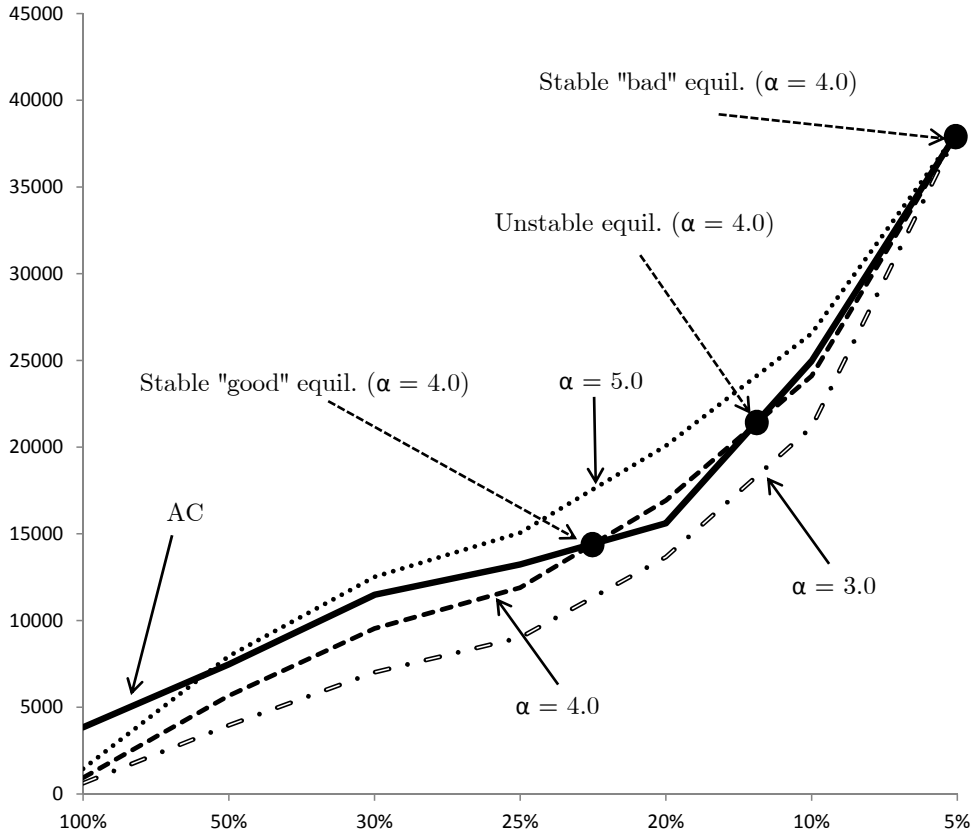

Source: 2010 Medical Expenditure Panel Survey.
Explanation: Wealth equals net worth (assets less liabilities), including net housing wealth.

In this case, the willingness-to-pay curve intersects the AC curve just once before again joining the AC curve at the smallest value of $X\%$. Since at the good equilibrium (close to the top 50% of spenders), the demand curve intersects the cost curve from below, we know from Section 5 that it is stable, whereas the equilibrium at the top 5% is unstable. Starting from any initial condition, the dynamics will bring the market to the stable equilibrium located at the larger rank, with more than half of the population being covered, consistent with a comparatively large level of risk aversion.

Finally, consider the in-between case where $\alpha = 4$. Notice that the willingness-to-pay and AC curves now intersect at three places, at the shown stable "good" and "bad" equilibria and an unstable intermediate equilibrium. Notice that the sharply rising AC curve as rank $X\%$ grows smaller plays a critical role in causing these multiple intersections. As we know from Proposition 3, multiple equilibria are more likely to be produced as losses become more concentrated. Indeed, health costs are much more concentrated than many

Figure 12: Willingness to Pay and Average Costs: Median Liquid Assets

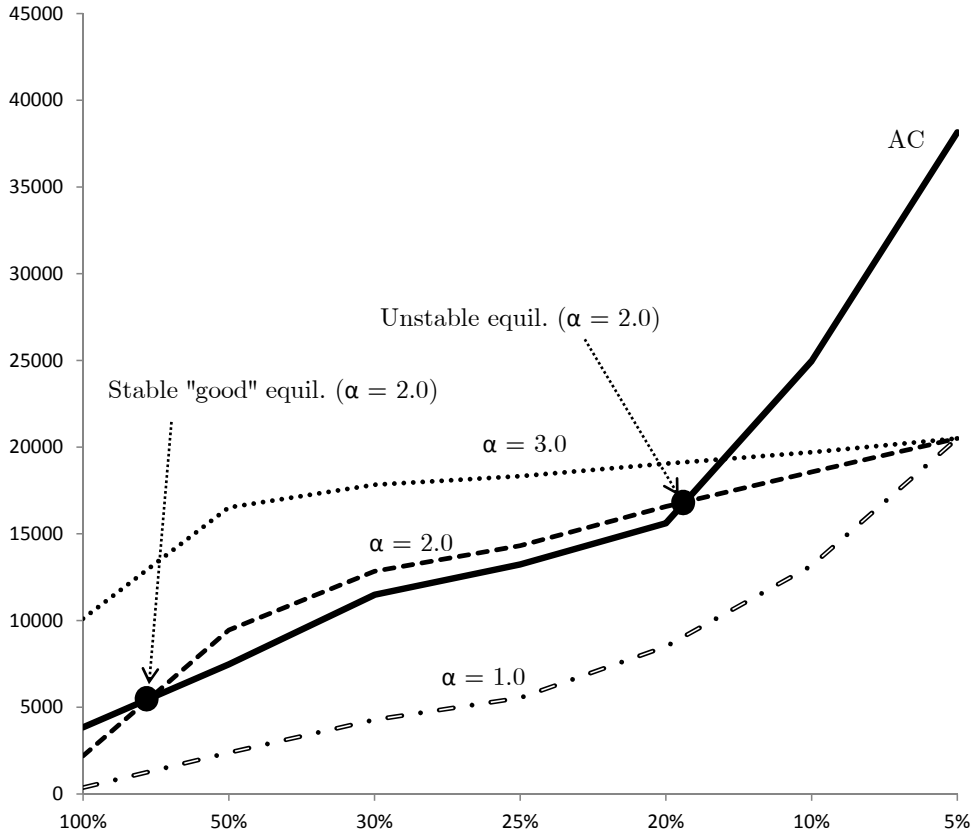

Source: 2010 Medical Expenditure Panel Survey.
Explanation: Wealth equals liquid assets only.

other types of insurable losses. By Proposition 2, initial conditions matter: if we start from a point to the right of the unstable equilibrium (covering somewhere between the top 10 and 20% of the population in terms of spending), the dynamics converge to the worst equilibrium with only the top 5% covered. Otherwise, we eventually reach the best equilibrium between the top 20 and 25% quantiles.

5.3.2 Robustness: Median Liquid Assets

To check the robustness of our results to various assumptions, Figure 12 repeats the same calculations assuming that wealth w is now set equal to the median *liquid assets* reported in the 2010 Survey of Consumer Finances. (Hence, the AC curve remains unchanged.) Liquid assets are a potentially more accurate measure of the relevant amount of wealth when illiquid assets, mainly housing, cannot be legally confiscated to pay for medical bills. Because the value of liquid assets is smaller than the median net worth, we can

consider relatively smaller values of α in our comparisons.

The reason why the highest willingness to pay no longer also joins the AC curve at its highest point is due to *limited liability*. The calibrated loss amount l now exceeds the wealth level w, and so the maximum potential loss is capped at w.[26] By Corollary 1, we know that the presence of limited liability generally enhances the likelihood of equilibrium multiplicity by reducing the slope of the willingness-to-pay curve in the same neighborhood where the slope of the AC curve is increasing.

Notice that the smallest value of α, now set equal to 1, produces a willingness-to-pay curve that is always below the AC curve, corresponding to the case of a single (corner) equilibrium with full unraveling at 0% (not shown). But we get three equilibria for the other two values of α. In particular, the largest value of α, now set equal to 3, produces just one intersection, corresponding to an unstable equilibrium. But this unstable equilibrium falls in-between two *corner* stable equilibria: a stable corner equilibrium at 100% where everyone buys insurance (where the willingness to pay exceeds the average cost of the entire population) and another stable corner equilibria at 0% (not shown) where nobody buys insurance. For the in-between value of α, now set equal to 2, we have a stable interior equilibrium, followed by an unstable equilibrium at a larger value of X%, followed by a corner stable equilibria at 0% (not shown) where nobody buys insurance.

5.3.3 Robustness: Assets and Probability Loss Probability Increase with Age

By focusing on median wealth for all cost quantiles, our estimates, however, have not accounted for the fact that both the size of wealth and the probability of loss tend to increase in age. For additional robustness, Figure 13 shows the effect of assuming that rank now grows linearly in age, where 18-year olds are now effectively located at the 100% mark on the horizontal axis while 64-year olds are located at the 5% mark. We can now also use the median values of liquid assets at each age from the 2010 Survey of Consumer Finances. Notice that allowing for this relationship has very little impact on our results. Relative to Figure 12, the effects of limited liability are now absent because older people tend to have both more assets and a higher probability of loss.

5.3.4 Robustness: Self-Employed

Figure 14 repeats the baseline scenario but now only includes people who identified as self-employed in the 2010 MEPS. This subpopulation might be a more relevant than the

[26] In the actual simulations, we cap the loss at w less \$1,000. Not only does this threshold avoid "almost" infinite marginal utility states, it roughly corresponds with Medicaid qualifications as well.

Figure 13: Willingness to Pay and Average Costs: Risk Increasing by Age

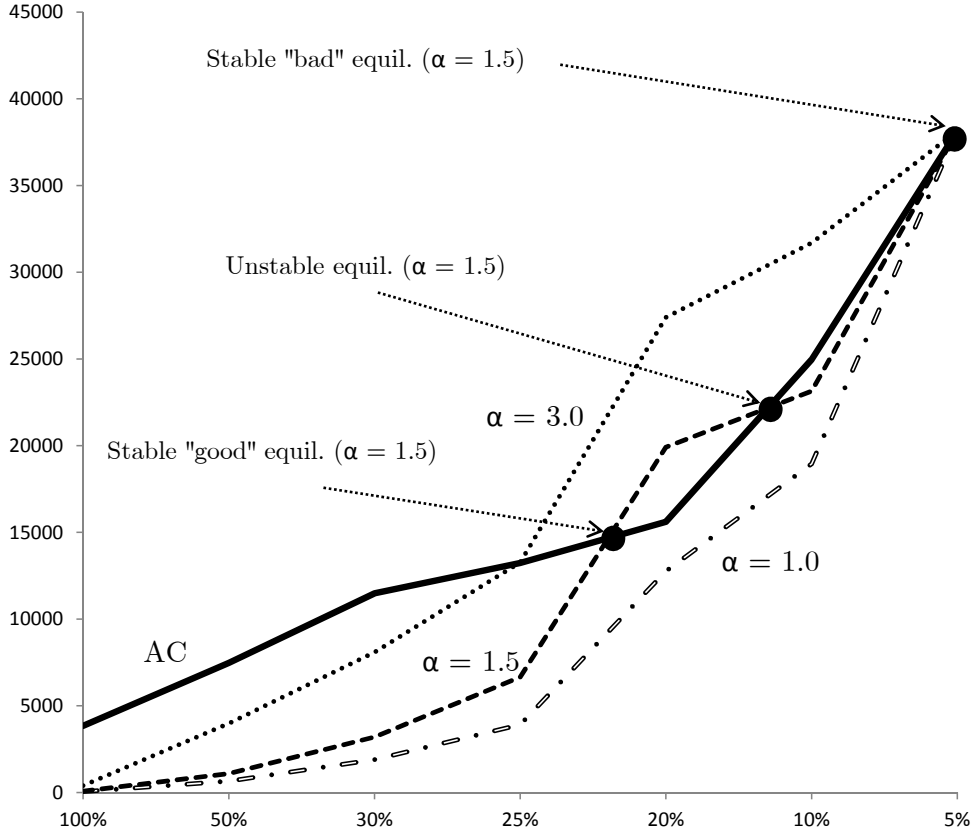

Source: 2010 Medical Expenditure Panel Survey.
Explanation: Wealth equals liquid assets only.

general population for the ACA experience, since these individuals do not have access to employer-provided health insurance and, therefore, are more likely to demand insurance on the exchanges. The results from the baseline scenario are basically the same, except that a smaller risk aversion, $\alpha = 3$, can now produce multiple equilibria, compared to $\alpha = 4$ in the baseline scenario.

5.3.5 Robustness: Uninsured

Figure 15 repeats the baseline scenario but now only includes people who were identified as uninsured in the 2010 MEPS. Again, this was the population targeted by the ACA health exchanges. The wealth of uninsured households is set equal to half of median net worth, consistent previous studies showing that uninsured households hold less wealth (e.g., Bernard et al. (2009)). The results from the baseline scenario are, again, basically the same, except that an even smaller risk aversion, $\alpha = 1.5$, can now produce multiple

Figure 14: Willingness to Pay and Average Costs: Self-Employed

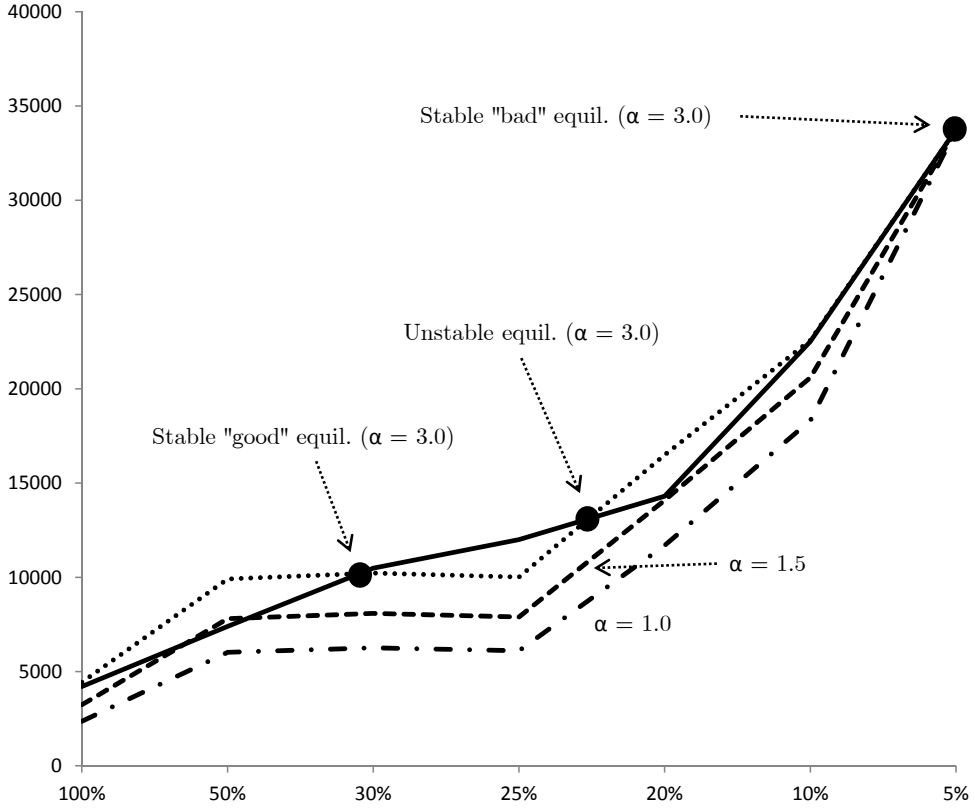

Source: 2010 Medical Expenditure Panel Survey.
Explanation: Wealth equals Median Net Worth.

equilibria.

5.3.6 Robustness: Fixed Probability, Variable Losses

Appendix B presents additional estimates that relax the assumption that the size of loss is fixed across different risk types. In particular, the probability of loss is now held fixed across households, and a risk type is now defined in terms of the heterogenous size of loss. As before, a low level of risk aversion α leads to unraveling. Moreover, the in-between level of risk aversion produces multiple stable equilibria. However, unlike before, even large values of risk aversion now produce multiple stable equilibria. The generalization provided in Appendix B also demonstrates that similar quantitative exercises could be performed while allowing for richer forms of heterogeneity.

Figure 15: Willingness to Pay and Average Costs: Uninsured

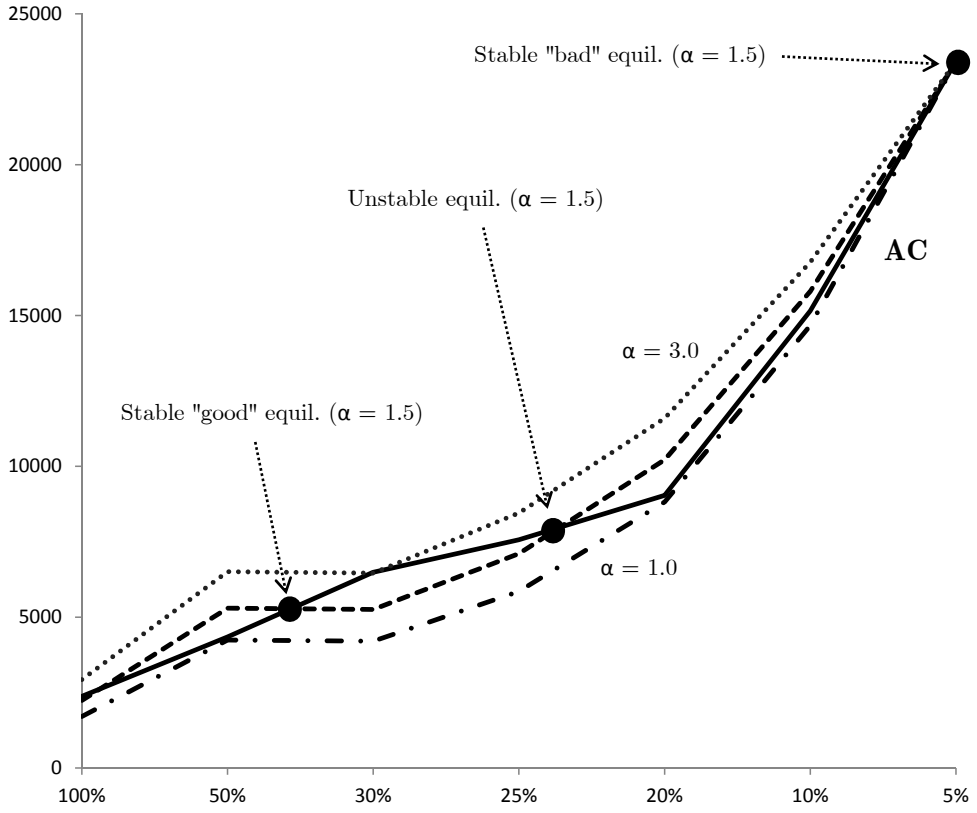

Source: 2010 Medical Expenditure Panel Survey.
Explanation: Wealth equals One-Half Median Net Worth.

5.3.7 Summary

Overall, therefore, the empirical evidence is consistent with the potential for multiple equilibria, especially for moderate values of risk aversion and when loss sizes are not fixed (Appendix B). To be sure, the calculations in this section focus on the pre-ACA population and do not include the impact of absolute fines for non-participation or the impact of subsidies to lower-income households who do participate. However, the next section shows that fines and subsidies shift up the inverse demand curve (potentially with some rotation due to means testing), thereby lowering the required level of risk aversion that is needed to achieve multiple equilibria. On one hand, fines and subsidies can help reduce the likelihood of multiple, stable equilibria that would have otherwise existed without these provisions. On the other hand, fines and subsidies can also *induce* multiple, stable equilibria in a setting where multiple equilibria would have otherwise not existed, for example, at lower levels of risk aversion. The net effect of fines and subsidies on

multiple, stable equilibria, therefore, is ambiguous.

Future work can simulate a life-cycle model with more detail, including the interaction of the ACA with other social insurance. For example, using a detailed micro-simulation model, Zewde (2017) shows that ACA subsidies are dominated by bankruptcy protection for many uninsured households. Consistently, a majority of uninsured households eligible for ACA subsidies have chosen not to participate (Levitt et al. (2016)).

6 Extension: Fines and Subsidies

Given the previous experience in the states, the mandate—in reality, the associated *fine* that gives the mandate its force—has been widely viewed as one of the important ingredients for the ACA to succeed. Indeed, the mandate was the focus of the challenge to the ACA heard by the Supreme Court in *National Federation of Independent Business v. Sebelius*. Chandra et al. (2011) present evidence that the phase-in of the mandate in the Massachusetts plan encouraged healthier consumers to enroll, and Hackmann et al. (2015) estimate the socially optimal level of the penalty that enforces the mandate in a model with linear cost and demand curves.

In this section, we show that the actual *form* of the fine, and not just its level, plays an important role in the presence of equilibrium multiplicity. The ACA imposes an absolute fine—a fixed dollar amount or a percentage of income, whichever is greater—whereas the Massachusetts reform in 2006 created a *relative* fine that was a function of market premium prices. In addition to the fine, the ACA also makes subsidies available to households with lower income.

6.1 Absolute Fine and Subsidy

We begin by demonstrating how the introduction of an *absolute* fine f for not having insurance as well as a subsidy s for having insurance can be captured in our graphical framework. Of course, neither the fine nor the subsidy affects the average cost curve $\Gamma(\pi)$. However, they affect the construction of the willingness to pay for insurance, now denoted as $\hat{\Omega}(\pi)$, through the modified indifference condition

$$u(w - \hat{\Omega} + s) = \pi u(w - f - l) + (1 - \pi)u(w - f). \tag{7}$$

Notice that the subsidy s and the fine f both *shift up* the inverse demand curve $\hat{\Omega}(\pi)$, in fact, in a parallel manner in the case of a subsidy. Notice that potentially numerous

Figure 16: Enforcement of a Mandate through a Fine f

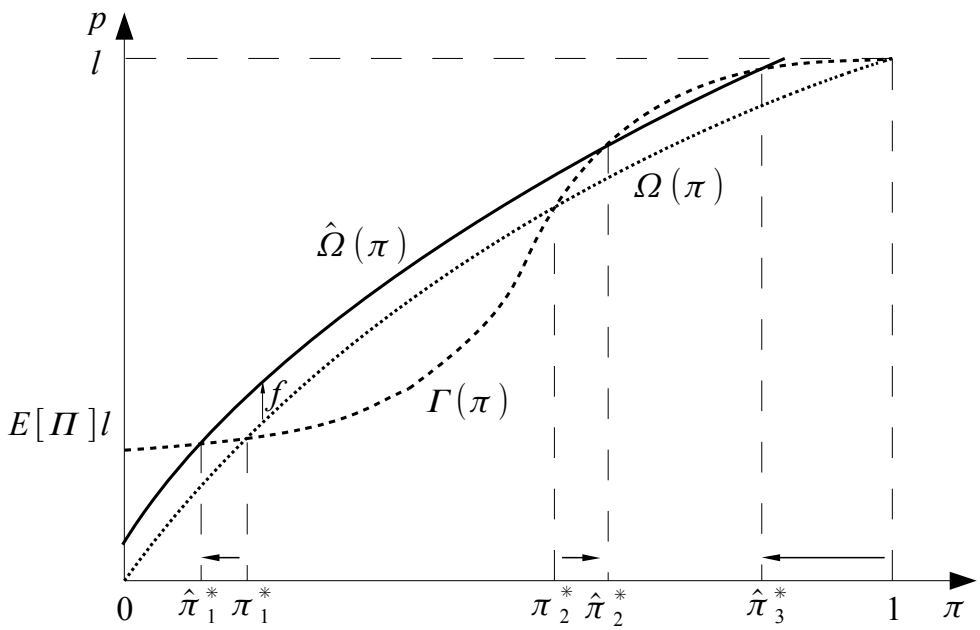

combinations of fines and subsidies can, therefore, actually *induce* multiple, stable equilibria in settings where they would have otherwise not existed, including in many of the calibration exercises in Section 5.2 with more modest levels of risk aversion. Given its greater complexity and its historical focus in ACA-related debates, we, therefore, focus on the construction of the fine f.

Unsurprisingly, an absolute fine could also give rise to better equilibria with more individuals insured. In the example previously shown in Figure 5, where only the complete unraveling equilibrium exists, shifting up the $\hat{\Omega}$-curve will induce the emergence of an equilibria where a positive mass of individuals get coverage. For the example shown in Figure 8 with three equilibria, Figure 16 illustrates how shifting up the $\hat{\Omega}$-curve can shrink the range of initial values $(\hat{\pi}_2^*, 1]$ for which unraveling to the bad stable equilibrium π_3^* occurs. It also shifts both the good stable equilibrium π_1^* and the bad stable equilibrium π_3^* to the left and, hence, leads to a greater number of individuals being covered at a lower premium.

6.2 Speed of Adjustment

The speed of adjustment to a positive fine is material after a reform like the ACA that shifts the willingness-to-pay schedule $\Omega(\pi)$ and/or the average cost curve $\Gamma(\pi)$. The

pre-reform equilibrium, denoted as π_{pre}, is, of course, no longer an equilibrium after the policy change. If π_{pre} is larger than the *post-reform* value of the unstable equilibrium with *no fine*, π_2^*, then the market will unravel toward the new bad stable equilibrium after the reform. Let f_{post} denote the minimum value of the fine that is necessary to shift the market from the new bad stable equilibrium to the new good equilibrium. Similarly, let f_{pre} denote the minimum value of the fine that is necessary to shift the market from the pre-reform equilibrium π_{pre} to the new good equilibrium. It is easy to see from Figure 16 that $f_{post} > f_{pre}$. In other words, a delay in implementing a fine after a policy change could require a larger fine to obtain the good equilibrium relative to a fine that is introduced with the reform itself. More generally, an incremental increase in the value of the fine over time—as done in the ACA—could result in a larger fine being necessary to eventually shift the market to the good equilibrium.

6.3 A Pareto Improving Shift to a Relative Fine

Of course, we can always set the fine to be large enough such that there exists a unique equilibrium where everyone buys insurance and there is no risk of unraveling. However, this outcome may be both inefficient and politically challenging. Indeed, the peculiar nature of the fine's construction under the ACA—namely, its assessment only on tax filers who are owed a refund along with generous exclusions for financial hardship—reflects the sensitivity that Congress felt it faced in creating a fine that causes too much hardship.

A more interesting question, therefore, is whether there exists another fine mechanism that eliminates the possibility of equilibrium multiplicity *without* imposing a higher fine *in the best equilibrium* π_1^*. As Figure 16 makes clear, it is actually not necessary to impose a higher fine *everywhere* in order to eliminate the bad stable equilibria. A large fine value is only necessary in situations where few individuals enroll for insurance, that is, where the critical value of π and, hence, the premium are both large.

A *relative* fine that is tied to the actual equilibrium premium in the market achieves exactly this outcome. Using the dynamics developed in Section 5, suppose that in each period t, the fine that must be paid by uninsured consumers is set equal to kp_t, where $k > 0$ is some constant that can be interpreted as the percentage of the current premium p_t.[27] Since $p_t = \Gamma(\pi_t)$, the resulting willingness to pay for insurance, now denoted as $\tilde{\Omega}(\pi)$, for each risk level π is then defined implicitly by

$$u(w - \tilde{\Omega}) = \pi u(w - l - k\Gamma(\pi)) + (1 - \pi)u(w - k\Gamma(\pi)).^{28} \tag{8}$$

[27] In the case of Massachusetts, the corresponding value of k would roughly equal $1/2$.

[28] Notice that, since $\Gamma(\pi)$ is increasing, the right-hand side of (8) is still decreasing in π, and so $\tilde{\Omega}(\pi)$

Figure 17: Relative versus absolute fine with $k\Gamma(\hat{\pi}_1^*) = f$

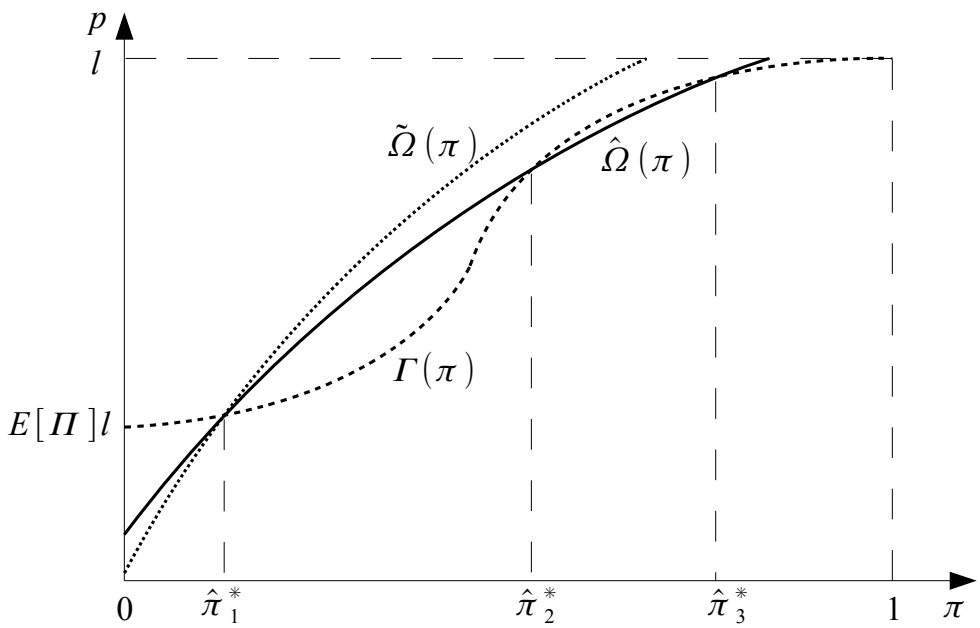

The benefit of the relative fine is that the fine value — and, hence, consumers' demand for insurance — automatically increase as the market unravels towards a bad stable equilibrium. This outcome occurs *even if* we choose k such that $k\Gamma(\hat{\pi}_1^*) = f$, so the relative and the absolute fines take exactly the same value in the best equilibrium $\hat{\pi}_1^*$. This is illustrated in Figure 17, which shows how the inverse demand curve under the relative fine $\tilde{\Omega}(\pi)$ is a counter-clockwise rotation at point $\hat{\pi}_1^*$ relative to the inverse demand curve under the absolute fine $\hat{\Omega}(\pi)$. Proposition 4 formalizes the advantages of a relative fine compared to an absolute fine with this normalization.

Proposition 4. *Let $\hat{\pi}_1^* < 1$ be the best equilibrium under an absolute fine $f > 0$, and set the relative fine such that $k\Gamma(\hat{\pi}_1^*) = f$ (i.e., equal fine values at the best equilibrium). Then:*
(i) for any number $N \geq 1$ of equilibria, the worst equilibrium $\tilde{\pi}_N^*$ under the relative fine is Pareto better (more coverage at a lower price) than the worst equilibrium under the absolute fine $\hat{\pi}_N^*$, i.e. $\tilde{\pi}_N^* \leq \hat{\pi}_N^*$.
(ii) The best equilibrium under the absolute and relative fine are identical, i.e. $\hat{\pi}_1^* = \tilde{\pi}_1^*$.
(iii) The interval of initial conditions $[0, \tilde{\pi}_1)$ from which we converge to the best equilibrium $\tilde{\pi}_1^*$ under the relative fine is larger than the range of initial conditions $[0, \hat{\pi}_1)$ from which we converge to the best equilibrium $\hat{\pi}_1^* = \tilde{\pi}_1^*$ under the absolute fine, i.e. $\tilde{\pi}_1 \geq \hat{\pi}_1$.

remains well-defined and increasing.

Proof. See Appendix. □

In sum, a re-construction of the fine toward a relative basis is more likely to expand coverage, by moving the market to the good stable equilibrium, without costing non-insured consumers anything more in the good equilibrium. However, even if the bad stable equilibrium does emerge (which may still be possible under a small enough relative fine), its "badness" is also reduced (more coverage at a lower price). A *shift* from an absolute fine, which exists under the ACA, to a properly constructed relative fine, which previously existed in the Massachusetts plan, is Pareto improving.[29]

While the normalization $k\Gamma(\hat{\pi}_1^*) = f$ requires knowledge of the best equilibrium on the part of the policymaker, part (i) of Proposition 4 can be shown to go through under the weaker condition that $k\Gamma(\pi) = f$ for some $\pi < \hat{\pi}_N^*$. In words, whenever the relative fine is set such that it coincides with the absolute fine for some sufficiently low premium level (below the one corresponding to the worst equilibrium), the worst equilibrium under the relative fine is better than the worst equilibrium under the absolute fine. Moreover, when the relative and absolute fine are equal under a premium level sufficiently close to the best equilibrium, parts (ii) and (iii) also go through in an approximate sense. Hence, for practical purposes, the relative fine dominates the absolute fine under weaker informational requirements than suggested by the exact result in the proposition.

6.4 Relative to No Fine

It is worth noting that simply introducing a fine for non-participation typically does not lead to a Pareto improvement relative to *no fine*. The reason is that there are usually individuals with sufficiently low π who, in any equilibrium without a fine, prefer to demand no insurance (in particular, this is always true in our baseline model with full support for π). With a fine in place, they will either remain uninsured and pay the fine or, if the fine is high enough, buy insurance at a premium that is higher than their original willingness to pay. In either case, they will be worse off. Hence, a fine or mandate can increase coverage but typically not in a Pareto improving way.[30]

[29]The same is true for subsidies that are a function of average market premiums rather than just income-dependent. Indeed, the construction of subsidies in the ACA indirectly implies some degree premium dependence because, for instance, people whose income is between 100% and 133% of the Federal Poverty Level receive subsidies so that their premium contribution amounts to no more than 2% of their income. Our analysis reveals that these features are highly desirable beyond their purely redistributive benefits.

[30]This argument goes through unaffected when the revenue from the fine is returned lump-sum to all individuals, or when the fine on non-participants is replaced by a subsidy for participation that is financed by a lump-sum tax.

Even when relaxing the assumption that π has full support on $[0, 1]$, as in Appendix B (as well as in the model calibrated to the MEPS data considered earlier), a Pareto improvement from introducing a fine is possible only in the presence of multiple equilibria, when the fine induces a shift from a bad equilibrium with low coverage to a good equilibrium where in fact *everyone* gets covered, and everyone being covered is an equilibrium *even without the fine*. Only this outcome guarantees that nobody ends up paying the fine and even the lowest risk types actually prefer buying insurance when everyone does so, so they are better off compared to the bad equilibrium (see e.g. Figure 18 in Appendix B). Explicitly accounting for equilibrium multiplicity, therefore, crucially underlies standard arguments for Pareto improving mandates or fines in the context of adverse selection. Nonetheless, *given* the existence of an absolute fine, an appropriate shift to a relative fine is Pareto improving.

7 Conclusion

This paper documents empirically that states that suffered poor exogenous "initial conditions" in the ACA marketplace, in the form of website failures at launch, appear to be stuck in a permanent position with higher average costs. This result is not consistent with the standard insurance model with strategic pricing. Nor is this result consistent with price-taking models with linear average cost and inverse demand curves, which are ubiquitous in the literature. We then characterize when the "initial conditions" of a new insurance market could have permanent consequences. Initial conditions can be material if (i) insurers are competitive price-takers, and if (ii) there exist at least three competitive equilibria. Existing evidence from previous health care reforms at the state level and from some employer-based plans suggest that insurers, indeed, update their prices consistent with the price-taking model.

While some previous papers have noted the possibility of multiple equilibria, this paper appears to be first to formalize the conditions required to have multiple *stable* equilibria, which are the only equilibria that are relevant for policy purposes. We also provide some suggestive empirical evidence using the Medical Expenditure Panel Survey that the presence of three equilibria is indeed consistent with moderate to low levels of risk aversion. Multiple equilibria are also more likely to emerge when losses are very concentrated (as is the case of health care) and in the presence of limited liability.

Moreover, the presence of subsidies and fines have an ambiguous impact on the presence of multiple, stable equilibria. While these provisions can reduce the impact of multiple equilibria that otherwise would exist, these provisions can also induce multiple

equilibria in settings where multiple equilibria would otherwise not exist. Moreover, we demonstate that equilibrium multiplicity is a necessary condition for a fine to achieve a Pareto improvement. The ACA's fine is currently constructed as an absolute amount, equal to the greater of a fixed dollar amount or a fixed fraction of income. In contrast, the 2006 Massachusetts plan, on which the ACA is modeled, levied a fine on a relative basis, equal to a fraction of the equilibrium premium. The relative fine, therefore, grows with the amount of adverse selection. We show that changing the fine from an absolute to a relative amount—normalized to be equal in the desired, good equilibrium—increases the range of initial conditions consistent with reaching the good equilibrium, while also reducing the severity of the bad equilibrium, if it still exists. The shift, therefore, would be Pareto improving.

Future work can explore related policy questions. For example, there could be good reasons for limitations to the frequency of price changes and the amount of increases, such as consumer protection. However, a potential unintended consequence is that they also further discourage price discovery, thereby increasing the potential for reaching a Pareto dominated equilibrium.

References

AKERLOF, G. (1970): "The Market for Lemons: Qualitative Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84, 488–500.

BERNARD, D. M., J. S. BANTHIN, AND W. E. ENCINOSA (2009): "Wealth, income, and the affordability of health insurance," *Health Affairs*, 28, 887–896.

BOARD OF GOVERNORS (2012): "Changes in U.S. Family Finances from 2007 to 2010: Evidence from the Survey of Consumer Finances." *Federal Reserve Bulletin*, 98, 1 – 80.

CENTERS FOR MEDICARE & MEDICAID SERVICES (2010): "Shining a Light on Health Insurance Rate Increases." Tech. rep., https://www.cms.gov/CCIIO/Resources/Fact-Sheets-and-FAQs/ratereview.html.

CHANDRA, A., J. GRUBER, AND R. MCKNIGHT (2011): "The Importance of the Individual Mandate - Evidence from Massachusetts.'," *New England Journal of Medicine*, 364, 293–295.

CLEMENS, J. (2015): "Regulatory redistribution in the market for health insurance," *American Economic Journal: Applied Economics*, 7, 109–134.

CNN (2016): "UnitedHealth expects to lose nearly $1 billion on Obamacare," http://money.cnn.com/2016/01/19/news/economy/unitedhealth-obamacare/index.html/.

COHEN, S. B. AND N. UBEROI (2013): "Differentials in the Concentration in the Level of Health Expenditures across Population Subgroups in the U.S., 2010," *Medical Expenditure Panel Survey*, Brief 421, 1 – 9.

COHN, J. (2012): "The New Jersey Experience: Do Insurance Reforms Unravel Without An Individual Mandate?" Tech. rep., http://www.kaiserhealthnews.org/Stories/2012/March/21/nj-ind-mandate-case-study.aspx.

CUTLER, D. AND S. REBER (1998): "Paying for Health Insurance: The Tradeoff Between Competition and Adverse Selection," *Quarterly Journal of Economics*, 113, 433–466.

DASH, S. J. AND A. THOMAS (2014): "New State-Based Marketplaces Unlikely in 2015, but Technology Challenges Create More Shades of Gray," *The Commonwealth Fund Blog, May*, 1.

EINAV, L. AND A. FINKELSTEIN (2011): "Selection in Insurance Markets: Theory and Empirics in Pictures," *Journal of Economic Perspectives*, 25, 115–138.

EINAV, L., A. FINKELSTEIN, AND M. CULLEN (2010a): "Estimating Welfare in Insurance Markets Using Variation in Prices," *Quarterly Journal of Economics*, 125, 877–921.

EINAV, L., A. FINKELSTEIN, AND J. LEVIN (2010b): "Beyond Testing: Estimating Empirical Models of Insurance Markets," *Annual Review of Economics*, 2, 311–336.

HACKMANN, M. B., J. T. KOLSTAD, AND A. E. KOWALSKI (2015): "Adverse selection and an individual mandate: When theory meets practice," *The American economic review*, 105, 1030–1066.

HANDEL, B., I. HENDEL, AND M. D. WHINSTON (2015): "Equilibria in health exchanges: Adverse selection versus reclassification risk," *Econometrica*, 83, 1261–1313.

HENDREN, N. (2014): "Unravelling vs Unravelling: A Memo on Competitive Equilibriums and Trade in Insurance Markets," *The Geneva Risk and Insurance Review*, 39, 176–183.

KAISER FAMILY FOUNDATION (2013): "Frequently Asked Questions about Health Care Reform," Tech. rep., Link: http://kff.org/health-reform/faq/health-reform-frequently-asked-questions/.

KHN (2014): "Doctors Say Obamacare Rule Will Stick Them With Unpaid Bills," Tech. rep., http://khn.org/news/doctors-say-obamacare-rule-will-stick-them-with-unpaid-bills/.

KOWALSKI, A. (2014): "The Early Impact of the ACA State-by-State," *Brookings Papers on Economic Activity*.

LEVITT, L., G. CLAXTON, A. DAMICO, AND C. COX (2016): "Assessing ACA Marketplace Enrollment," *Kaiser Family Foundation*.

MAHONEY, N. AND E. G. WEYL (forthcoming): "Imperfect Competition in Selection Markets," *Review of Economics and Statistics*.

MAS-COLELL, A., M. WHINSTON, AND J. GREEN (1995): *Microeconomic Theory*, Oxford: Oxford University Press.

MONHEIT, A., J. CANTOR, M. KOLLER, AND K. FOX (2004): "Community Rating And Sustainable Individual Health Insurance Markets In New Jersey," *Health Affairs*, 23, 167–175.

NETZER, N. AND F. SCHEUER (2013): "A Game-Theoretic Foundation of Competitive Equilibria with Adverse Selection," *International Economic Review, forthcoming*.

NYTIMES (2017): "Humana Plans to Pull Out of Obamacares Insurance Exchanges," https://www.nytimes.com/2017/02/14/health/humana-plans-to-pull-out-of-obamacares-insurance-exchanges.html.

OBAMACARE FACTS (2014): "ObamaCare Myths: Myths About the Affordable Care Act," Tech. rep.

PAULY, M. (2017): "Exchanges Can Survive End of Individual Mandate," http://ldi.upenn.edu/healthpolicysense/exchanges-can-survive-end-individual-mandate.

ROTHSCHILD, M. (1974): "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 9, 185–202.

ROTHSCHILD, M. AND J. STIGLITZ (1976): "Equilibrium in Competitive Insurance Markets: An Essay in the Economics of Incomplete Information," *Quarterly Journal of Economics*, 90, 629–649.

THE ENERGY AND COMMERCE COMMITTEE (2014): "Committee Learns Whos Paid for Obamacare: As of April 15, Only 67 Percent of Enrollees in Federal Marketplace Had Paid First Months Premium," Tech. rep., https://energycommerce.house.gov/news-center/press-releases/committee-learns-who-s-paid-obamacare-april-15-only-67-percent-enrollees.

USATODAY (2016): "Aetna's exit deals blow to Obamacare, patients," http://www.usatoday.com/story/money/2016/08/16/aetna-obamacare-affordable-care-act-exchanges/88825798/.

WILSON, C. (1977): "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory*, 12, 167–207.

——— (1980): "The nature of Equilibrium in Markets with Adverse Selection," *Bell Journal of Economics*, 11, 108–130.

ZEWDE, N. (2017): "A Micro Simulation Study of the Financial Benefits of Marketplace Insurance for Previously Uninsured Adults," Tech. rep.

A Appendix: Proofs

A.1 Proof of Proposition 2

(i). We show that initial conditions cannot matter when there are only one or two competitive equilibria. If $\pi^* = 1$ is the only equilibrium, we must have $\Gamma(\pi) > \Omega(\pi)$ for all $\pi < 1$. Otherwise, by continuity of Γ and Ω and since $\Gamma(0) = \mathbb{E}[\Pi]l > \Omega(0) = 0$, there would have to exist at least one intersection of Γ and Ω at some $\pi < 1$ and hence another equilibrium. Since $\Gamma(\pi) > \Omega(\pi)$ for all $\pi < 1$, the dynamics imply unraveling to $\pi^* = 1$ for any initial π and therefore initial conditions do not matter.

If there are two equilibria with $\pi^* = 1$ and some $\pi_1^* < 1$, it must hold that $\Gamma(\pi) > \Omega(\pi)$ for all $\pi \in [0, \pi_1^*)$ and $\Gamma(\pi) < \Omega(\pi)$ for all $\pi \in (\pi_1^*, 1)$ by an analogous argument as above. Hence, for any $\pi \in [0, 1)$, we converge to the constrained efficient equilibrium $\pi_1^* < 1$. As a result, initial conditions again do not matter for the equilibrium that is eventually reached except in the non-generic case where the initial $\pi = 1$.

(ii) and *(iii).* Note first that, for any number N of equilibria $\pi_1^* < ... < \pi_{N-1}^* < 1$, the best equilibrium π_1^* must be stable generically. This is because $\Gamma(\pi) > \Omega(\pi)$ for all $\pi \in [0, \pi_1^*)$ and, since there is a proper intersection of Γ and Ω generically, $\Gamma(\pi) < \Omega(\pi)$ for all $\pi \in (\pi_1^*, \pi_1)$ and $\pi_1 > \pi_1^*$ sufficiently close to π_1^*. With exactly 3 competitive equilibria $\pi_1^* < \pi_2^* < 1$, this implies that in fact $\pi_1 = \pi_2^*$, and since again there is a proper intersection of Γ and Ω at π_2^* generically, $\Gamma(\pi) > \Omega(\pi)$ for all $(\pi_2^*, 1)$, as illustrated in Figure 9. Hence, the intermediate equilibrium π_2^* is unstable and the other two are stable. Moreover, we converge to π_1^* for any initial $\pi < \pi_2^*$ and to $\pi^* = 1$ for any $\pi > \pi_2^*$.

A.2 Proof of Proposition 3

Solving the definition of $\Omega(\pi)$ in (2) yields

$$\Omega(\pi) = w - c(\pi u(w - l) + (1 - \pi)u(w)),$$

where $c(u)$ denotes the inverse function of $u(c)$. This implies

$$\Omega'(\pi) = \frac{u(w) - u(w - l)}{u'(c(\pi u(w - l) + (1 - \pi)u(w)))} > 0,$$

where we used $c'(u) = 1/u'(c(u))$. Evaluating at $\pi = 1$ delivers

$$\Omega'(1) = \frac{u(w) - u(w - l)}{u'(w - l)} > 0.$$

On the other hand, differentiating the definition of $\Gamma(\pi) = \mathbb{E}[\Pi|\Pi \geq \pi]l$ from (1) yields

$$\Gamma'(\pi) = \frac{h(\pi)}{1 - H(\pi)}(\mathbb{E}[\Pi|\Pi \geq \pi] - \pi)l \geq 0.$$

Using L'Hospital's rule, we obtain

$$\Gamma'(1) \equiv \lim_{\pi \to 1} \Gamma'(\pi) = \frac{l}{2}.$$

Hence, we have $\Gamma'(1) < \Omega'(1)$ if condition (i) in the proposition is satisfied. Together with $\Gamma(1) = \Omega(1) = l$, this implies $\Gamma(\pi) > \Omega(\pi)$ in an interval $[1 - \epsilon, 1)$ for some $\epsilon > 0$. Hence, there exists a stable competitive equilibrium at $\pi^* = 1$. Note that condition (6) is satisfied whenever l is sufficiently small because

$$\lim_{l \to 0} \frac{u(w) - u(w - l)}{l} = u'(w) > 0.$$

It also satisfied if u is not too concave since with $u'' = 0$ we have

$$\frac{u(w) - u(w - l)}{l} = u'(w - l) > \frac{u'(w - l)}{2}.$$

Next, under condition (ii) in the proposition, as $\Delta, \varepsilon \to 0$, we have $\mathbb{E}[\Pi] \to a$ and hence $\Gamma(\pi) = \mathbb{E}[\Pi|\Pi \geq \pi]l \to a \; \forall \, \pi \leq a$. In particular, this implies $\Gamma(a) \to a < \Omega(a)$ since $\Omega(\pi) > \pi$ for all $\pi \in (0, 1)$ by risk aversion. Because $\Omega(0) = 0 < \Gamma(0) \approx la$, this implies that there must exist a stable equilibrium with $\Omega(\pi^*) = \Gamma(\pi^*)$ for some $\pi^* \in (0, a)$. Finally, taken together with the above result that $\Omega(\pi) < \Gamma(\pi)$ in the left-neighborhood of $\pi = 1$, this implies that there must exist at least one unstable equilibrium in $(a, 1)$.

A.3 Proof of Corollary 1

When individuals are only liable for losses up to $L < l$, their demand for insurance is determined by

$$u(w - \Omega) = \pi u(w - L) + \pi u(w),$$

so $\Omega(1) = L < \Gamma(1) = l$. By the same argument as in the first part of the proof of Proposition 3, this ensures the existence of a stable "corner" equilibrium at $\pi = 1$, i.e. where noone buys insurance (and the willingness to pay of even the highest risk types is below their average costs). When L is not too small, a sufficiently concentrated distribution $H(\pi)$ ensures the existence of the other equilibria by the same argument as in the second

part of the proof of Proposition 3.

A.4 Proof of Proposition 4

Note first that, for any absolute fine $f > 0$, $\hat{\Omega}(1) = l + f > \Gamma(1) = l$, so for any number N of equilibria, both the best equilibrium $\hat{\pi}_1^*$ and the worst equilibrium $\hat{\pi}_N^*$ under f must satisfy $\hat{\pi}_1^* \leq \hat{\pi}_N^* < 1$. We next observe that, comparing the definitions (7) and (8) and using the normalization that $k\Gamma(\hat{\pi}_1^*) = f$ and the fact that $\Gamma(\pi)$ is increasing in π, $\tilde{\Omega}(\pi) < \hat{\Omega}(\pi)$ for all $\pi < \hat{\pi}_1^*$ and $\tilde{\Omega}(\pi) > \hat{\Omega}(\pi)$ for all $\pi > \hat{\pi}_1^*$. We use this repeatedly to prove claims *(i)* to *(iii)* in the proposition.

(i). Since $\hat{\Omega}(1) > \Gamma(1)$ under $f > 0$, the worst equilibrium $\hat{\pi}_N^* < 1$ must be such that $\hat{\Omega}(\pi) > \Gamma(\pi)$ for all $\pi > \hat{\pi}_N^*$. Since $\hat{\pi}_N^* \geq \hat{\pi}_1^*$, the above result that $\tilde{\Omega}(\pi) > \hat{\Omega}(\pi)$ for all $\pi > \hat{\pi}_1^*$ a fortiori implies $\tilde{\Omega}(\pi) > \hat{\Omega}(\pi) > \Gamma(\pi)$ for all $\pi > \hat{\pi}_N^*$. This immediately rules out $\tilde{\pi}_N^* > \hat{\pi}_N^*$.

(ii) and (iii). Note first that the best equilibrium $\hat{\pi}_1^*$ is always such that $\hat{\Omega}(\pi) < \Gamma(\pi)$ for all $\pi < \hat{\pi}_1^*$. Moreover, since we observed that $\tilde{\Omega}(\pi) < \hat{\Omega}(\pi)$ for all $\pi < \hat{\pi}_1^*$, we also have $\tilde{\Omega}(\pi) < \Gamma(\pi)$ for all $\pi < \hat{\pi}_1^* \leq \tilde{\pi}_1^*$. Hence, the range of initial values from which we converge to the best equilibrium always takes the form of an interval with lower bound zero and upper bound $\pi_1 \geq \hat{\pi}_1^*$. Since $\tilde{\Omega}(\pi) > \hat{\Omega}(\pi)$ for all $\pi > \hat{\pi}_1^*$ and $\tilde{\Omega}(\hat{\pi}_1^*) = \hat{\Omega}(\hat{\pi}_1^*)$, we also have $\tilde{\pi}_1^* = \hat{\pi}_1^*$, as claimed in *(ii)*.

Suppose first that the best equilibrium $\hat{\pi}_1^* < 1$ is the unique equilibrium under the absolute fine f, so $\hat{\Omega}(\pi) > \Gamma(\pi)$ for all $\pi > \hat{\pi}_1^*$ and vice versa. Then by the above observation that $\tilde{\Omega}(\pi) > \hat{\Omega}(\pi)$ for all $\pi > \hat{\pi}_1^*$ and vice versa, this immediately implies that $\tilde{\pi}_1^* = \hat{\pi}_1^*$ is also the unique equilibrium under the relative fine. It also implies that, in both cases, the best equilibrium is globally stable, so we converge to it for any initial conditions and thus $\tilde{\pi}_1 = \hat{\pi}_1 = 1$.

Otherwise, since the best equilibrium $\hat{\pi}_1^*$ generically corresponds to a proper intersection of $\hat{\Omega}(\pi)$ and $\Gamma(\pi)$ and $\hat{\Omega}(\pi) < \Gamma(\pi)$ for all $\pi < \hat{\pi}_1^*$, me must have $\hat{\Omega}(\pi) > \Gamma(\pi)$ for some interval $(\hat{\pi}_1^*, \hat{\pi}_1)$ with $\hat{\pi}_1 > \hat{\pi}_1^*$. Hence, under the absolute fine, we converge to $\hat{\pi}_1^*$ for any initial π in the interval $[0, \hat{\pi}_1)$. Then the above observation that $\tilde{\Omega}(\pi) > \hat{\Omega}(\pi)$ for all $\pi > \hat{\pi}_1^*$ immediately implies $\tilde{\Omega}(\pi) > \Gamma(\pi)$ for some interval $(\hat{\pi}_1^*, \tilde{\pi}_1)$ with $\tilde{\pi}_1 > \hat{\pi}_1$. The range of initial values for π for which we converge to the best equilibrium under the relative fine is therefore $[0, \tilde{\pi}_1)$ with $\tilde{\pi}_1 > \hat{\pi}_1$.

Figure 18: Equilibrium Multiplicity with Three Types

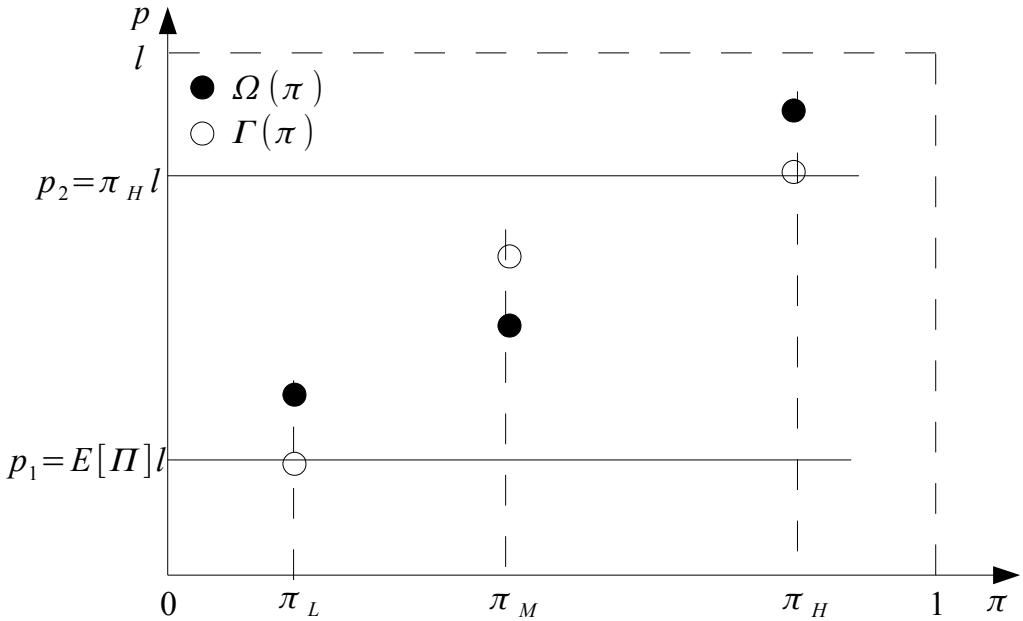

B Appendix: Generalizing the Price-Taking Model

This Appendix generalizes the model of Section 4 to allow for the presence of discrete risk types, a richer amount of heterogeneity between consumers and more general variation in the size of losses.

B.1 Allowing for Discrete Risk Types

We now show multiple equilibria can emerge even when we relax the assumption that the distribution of types $H(\pi)$ is continuous with full support on $[0, 1]$. For example, consider a case with three risk types, $0 < \pi_L < \pi_M < \pi_H < 1$, of low ($L$), medium ($M$) and high ($H$) risk, respectively. Their willingness to pay for insurance $\Omega(\pi)$ is depicted as black dots in Figure 18. The empty circles represent the average costs of insuring the corresponding pools, and so $\Gamma(\pi_H)$ is the cost of only insuring the high risk type H, $\Gamma(\pi_M)$ is the average cost of insuring both the medium M and high risk H types, and $\Gamma(\pi_L) = \mathbb{E}[\Pi]l$ is the average cost of insuring all three risk types.

We have chosen these values such that there are two competitive equilibria: one good equilibrium in which everyone is insured at premium $p_1 = \mathbb{E}[\Pi]l$, and another bad equilibrium in which only the high risk type is insured at a higher premium $p_2 = \pi_H l > p_1$.[31]

[31] With discrete types, competitive equilibria involve points with $\Omega(\pi) \geq \Gamma(\pi)$ rather than necessarily

There is no equilibrium where only the medium and high types π_M and π_H are insured, because the average cost $\Gamma(\pi_M)$ for that pool is higher than the willingness to pay of the medium type $\Omega(\pi_M)$, so the medium risk type would not buy insurance at premium $\Gamma(\pi_M)$ and the dynamics would unravel to the bad equilibrium.

Let us connect Figure 18 to the corresponding figures that we drew for the case of a continuum of types in Sections 4 and 5. Filling up the space between the three discrete types naturally leads to Figure 19. We see that, with continuous types and this pattern of curves, there are in fact three equilibria: a stable bad equilibrium, where only types $\pi \geq \pi_2$ are insured (with $\pi_M < \pi_2 < \pi_H$), an unstable interior equilibrium with critical type π_1 between π_L and π_H, and a stable corner equilibrium where everyone with $\pi \geq \pi_L$ is insured. Notice that Figure 19 is very similar to Figure 8 shown in Section 5. In particular, for initial conditions to matter, the existence of an unstable equilibrium is still required, and the average cost and willingness to pay curves need to intersect at least twice in the interior. Moreover, the marginal buyer π_1 in the unstable equilibrium represents the critical value of initial conditions that determines whether the good or bad equilibrium is reached eventually. The only difference between the discrete and the continuous cases is that the highest risk type in the discrete case may lay within the support shown for the continuous case.

B.2 Richer Forms of Consumer Heterogeneity and Multiple Loss Sizes

It is also straightforward to extend our analysis in Section 4 to allow for richer forms of consumer heterogeneity and multiple sizes of losses. Let the population be indexed by the continuous variable $\theta \in [0, 1]$ with distribution $F(\theta)$. Suppose there are S possible loss levels $l_s(\theta)$ indexed by s, which may differ across θ. The probability that type θ suffers a loss of size s is denoted by $\pi_s(\theta)$, where, of course, $\sum_{s=1}^{S} \pi_s(\theta) = 1 \; \forall \theta$. The expected loss for type θ is, therefore,

$$\sum_{s=1}^{S} \pi_s(\theta) l_s(\theta).$$

We can normalize the population type index θ so that the expected costs are increasing in θ. In particular, let us take θ as the *quantiles* of the average cost distribution, so that

$$\Gamma(\theta) = \int_{\theta}^{1} \sum_{s=1}^{S} \pi_s(\theta') l_s(\theta') dF(\theta') \Big/ (1 - F(\theta))$$

$\Omega(\pi) = \Gamma(\pi)$. However, Definition 1 still applies. For instance, in the competitive equilibrium with premium p_2, we have $\Omega(\pi_L) < \Omega(\pi_M) < p_2 < \Omega(\pi_H)$, so that only the highest type π_H demands insurance. Moreover, $p_2 = \Gamma(\pi_H) = \pi_H l$, and insurers make zero profits.

Figure 19: Connection between continuous and discrete type model

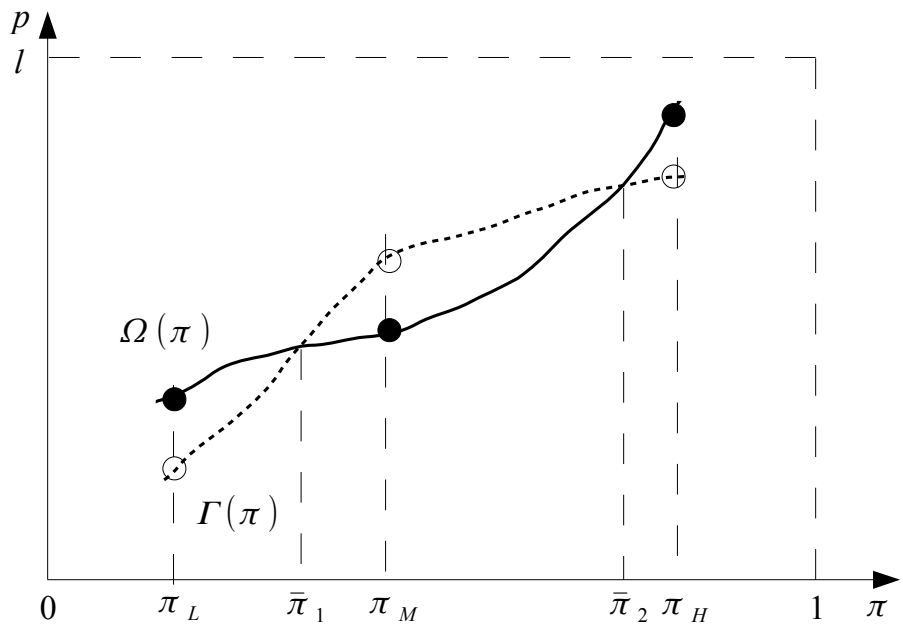

is the average cost of the most costly $1-\theta$ share of the population, and $F(\theta)=\theta$. Clearly, $\Gamma(\theta)$ is still increasing in θ as before.

Correspondingly, we can capture the consumers' willingness to pay for insurance for those individuals who are located at the θ-quantile of the cost distribution. Formally, for each quantile θ, let $\Omega(\theta)$ be given by the highest value of Ω such that

$$u(w(\theta)-\Omega;\theta)=\sum_{s=1}^{S}\pi_s(\theta)u(w(\theta)-l_s(\theta);\theta).$$

Note that we can allow for both wealth levels $w(\theta)$ and preferences (notably risk-aversion) $u(c;\theta)$ to vary across quantiles of the cost distribution; for instance, higher expected cost individuals may on average be wealthier (since older) or more risk-averse (they see the doctor more often).

As long as $\Omega(\theta)$ remains increasing — and, hence, higher expected cost individuals on average have a higher willingness to pay for insurance — our entire analysis from before is maintained: a competitive equilibrium corresponds to a quantile θ where $\Gamma(\theta)=\Omega(\theta)$. We can also employ the same graphical approach as before, the only difference being that the π-axis turns into an θ-axis of quantiles of the cost distribution. At an equilibrium with critical quantile θ^*, the share of the population purchasing insurance is given by $1-\theta^*$, and so $1-\theta$ can also be interpreted as quantity of insurance as in Einav et al. (2010a).

Figure 20: Willingness to Pay with Median Liquid Assets and Constant Loss Probability

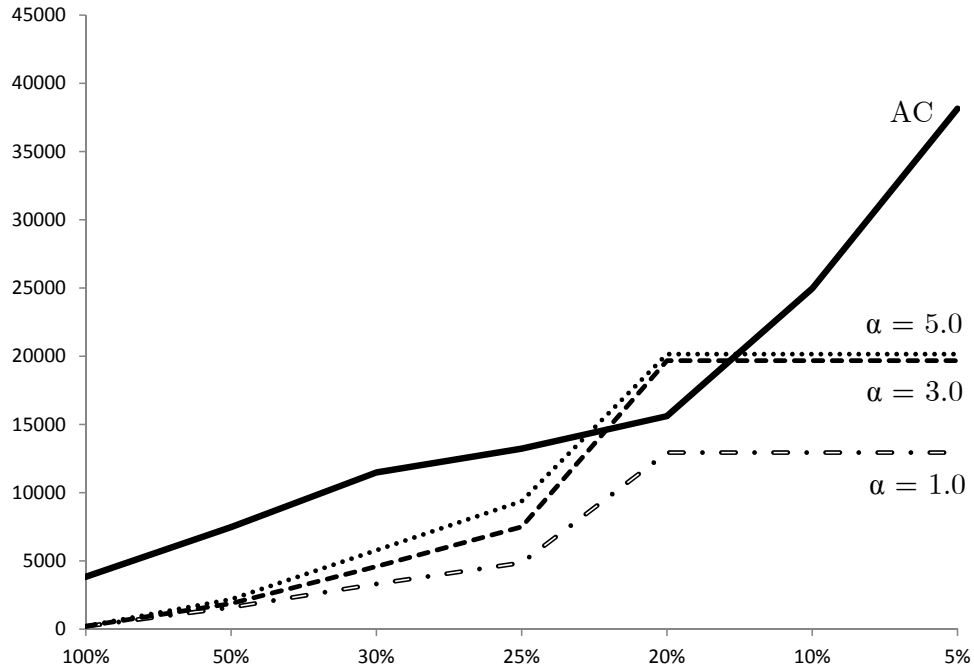

Source: Medical Expenditure Panel Survey (Cohen and Uberoi 2013).
Explanation: Wealth equals net worth (assets less liabilities), including net housing wealth.

Figure 20 shows the empirical evidence from the Medical Expenditure Survey Panel and the Survey of Consumer finances where the probability of loss π is fixed (at 0.3) but the size of loss $l_s(\theta)$ is now allowed to vary across the types. As before, the horizontal axis corresponds to the top X% percent of spenders, where X (the "rank") is the shown value. Now, however, the variation in spending comes from differences in loss amounts rather than probabilities. (Given the fixed value of π, a recursive algorithm parallel to that discussed in the text is used to impute the losses across the different values of X%.) As before, the relatively small value of $\alpha = 1$ leads to unraveling. However, both the in-between and large values of α lead to multiple stable equilibria: one at first intersection of willingness-to-pay and average cost lines and a second at the corner case where $X = 5$, where the willingness-to-pay is below the average cost. The driving force is, again, limited liability. As X% gets small, the value of losses must grow in order to match spending levels in the MEPS. As a result, the willingness to pay is capped for a wider range of types at smaller values of X% corresponding to larger losses. Increasing the level of risk aversion, therefore, has very little impact on the demand for insurance in this range.