Claim Denials: The Health Insurance Market from a Differential Learning Perspective

ECON 1071 / CS 37 Final Project

Jessica Chen, Eli Kirtley, Kasie Leung, Jessica Li, TingTing Yan

April 2025

Abstract: Coverage denials by insurance companies are a growing concern, leaving patients with unexpected medical bills and decreasing trust in insurance. We analyze how insurance companies may exploit information asymmetries by denying claims in a manner that is only gradually learned by consumers. We add an extension to adverse selection by incorporating dynamic Bayesian learning of consumer belief about denial rate. Our analytical toy model shows that insurers can maximize long-term profits by partially denying claims and allowing naive consumers to gradually drop out. Through simulations, we model consumers with heterogeneous risk levels who update their beliefs of claim denial rate based on personal claim experiences. We find that high-risk individuals exit earlier, initially improving the risk pool before low risk individuals also exit. This leads to a non-monotonic path in the evolution of the risk composition of the insured pool. Our results reveal how belief-based dropout creates strategic opportunities for insurers and suggests that temporary market stability can obscure deeper instability. These insights challenge static insurance models and highlight the need for consumer protections and transparency. Future extensions may explore interventions such as public denial rate disclosures or dynamic pricing schemes to stabilize participation and equity in health insurance markets.

1 Introduction

In recent years, coverage denials by health insurers have drawn widespread scrutiny from both policymakers and the public. Patients often find themselves unexpectedly responsible for high medical bills, even after paying premiums and believing they were covered. At the heart of this phenomenon is a growing practice among insurers to deny claims based on vague criteria like ”medical necessity” or pre-existing conditions. While such denials are ostensibly a tool for cost control, they also reshape the incentives and beliefs of insured individuals. Patients who experience denial not only face financial harm but may also revise their expectations of future coverage, ultimately leading them to exit the insurance market altogether.

From a theoretical perspective, this phenomenon poses a challenge to classical models of insurance under asymmetric information. In the canonical Rothschild-Stiglitz framework, adverse selection causes high-risk individuals to self-select into generous contracts, while low-risk individuals drop out, rendering full insurance unviable. But this model assumes that claims, once insured, are always honored. The reality of frequent coverage denials suggests that insurance itself is a risky good. This raises new questions: What happens when patients learn over time that their insurer is less likely to pay out? Can insurers profit by exploiting the gap between perceived and actual claim approval rates? And how does this affect the composition of the insured pool and long-term market dynamics?

We investigate these questions using both analytical and computational models. Our work builds on and extends prior literature, particularly a recent paper by Chade and Schlee (2020). They show how even in the presence of gains from trade, coverage may be denied to high-risk individuals due to provision costs, and markets may exhibit complete pooling or even fail to exist. Our contribution lies in incorporating belief updating and dynamic consumer behavior into this framework. We find that denial-based dropout can lead to surprising results, including temporary improvements in risk pool quality and non-monotonic selection effects.

2 Phenomenon Description

The phenomenon we study is the denial of health insurance claims—specifically, how insurers make profits by denying coverage in ways that are not always transparent or anticipated by consumers. Empirical data suggests that insurers routinely deny a significant fraction of claims, with 19% of in-network and 37% of out-of-network claims denied under the ACA marketplace.

From a societal standpoint, this is highly suboptimal. Health insurance is meant to reduce the financial volatility of health shocks. But if individuals are uncertain whether claims will be approved, or if denial seems arbitrary, the utility of insurance diminishes. This is especially harmful for risk-averse individuals, who are precisely those most likely to value insurance.

Our central claim is that insurers can profit from increasing denial rates precisely because insured consumers do not immediately update their expectations. A key insight is that consumers only learn about denial probabilities through their own claim experiences. High-risk consumers, who are more likely to file early claims, update faster and drop out sooner. Low-risk consumers may remain insured longer, creating a temporary improvement in the insured pool that is reversed over time.

Our analytical model shows how such strategies can produce equilibria where insurers set denial thresholds that maximize expected profit over a consumer's lifetime—balancing claim costs against dropout probabilities. Our simulation then confirms that under Bayesian learning, belief updating leads to declining participation, rising adverse selection, and eventual market erosion.

3 Literature Review

3.1 Rothschild-Stiglitz Model

Proposed in 1976, the Rothschild-Stiglitz model was the first to consider the consequences of imperfect information in insurance markets [4]. In this model, a health event is modeled as a decrease in income, $-c$. The insurance contract is given by the vector (α_1, α_2), where α_1 is the cost of the contract, and α_2 is the payout if a health event occurs. Finally, there exist two groups of buyers, a high-risk group with probability s^H of a health event, and a low-risk group with probability s^L of a health event. If the fraction of high-risk customers is λ, then the expected probability of incident for the firm is $\bar{s} = (1-\lambda)s^L + \lambda s^H$.

The customer always has the option of not purchasing insurance, in which case $U = -pc$. The value of contract for a buyer with known probability p is $U' = -\alpha_1 + p(-d+\alpha_2)$, so they will only buy if $U' - U > 0$, or $\alpha_1 < p\alpha_2$. It is a reasonable assumption that the firm does not know any individual p, since that is private information, but does know the expected $\bar{p}$ over the population, from sources like actuarial tables. Therefore, expected utility of a firm for selling one contract is is $U = \alpha_1 - \bar{p}\alpha_2$. Assuming consumers will buy at $\Delta U = 0$, and the population has a homogeneous $\bar{p}$, then all insurance contracts sold will be described by $(\bar{p}\alpha_2, \alpha_2)$.

However, if there are low-risk buyers and high-risk buyers, then the utility of the low-risk buyers for a contract $(\bar{p}\alpha_2, \alpha_2)$ is given by $U^L = -\bar{p}\alpha_2 + p^L(-d + \alpha_2)$, which is negative, since $\bar{p} > p^L$ by definition. Therefore, they choose to not buy insurance, and they drop out of the market, which means that the new expected probability of an health event is $\bar{p}' = p^H$, and high-risk individuals are offered $(p^H\alpha^H, \alpha^H)$. This prompts the consideration of a separating equilibrium, where high-risk buyers are offered coverage at or $(p^H\alpha^H, \alpha^H)$, such that $U^H = -p^H d$. However, low-risk buyers will not choose this contract, since $\Delta U^L = -p^H\alpha^H + p^L\alpha^H < 0$. Now, say that low-risk buyers are offered contracts, $(p^L\alpha^L, \alpha^L)$, which gives utility $U^L = -p^L\alpha^L + p^L(-\alpha^L + d) = -p^L d$, or fair odds. However, high-risk buyers would also choose this contract, since their utility is given by $\Delta U^H = -p^L\alpha^L + p^H(-d + \alpha^L) + p^H d = -p^L\alpha^L + p^H\alpha^L > 0$. Therefore, Rothschild and Stiglitz conclude that there can exist no separating equilibrium.

In the same way that the sale of used car is a "lemons market," and low-quality cars drive down the sale of high-quality cars, the cost of insuring high-risk individuals drives out the participation of low-risk individuals, and a pooling equilibrium cannot exist. This phenomenon is known as “adverse selection.” However, in modern insurance markets, there are more low-risk individuals than high-risk individuals, and insurance firms increase profit for every contract sold, so the phenomenon of “advantageous selection” occurs, in which firms attempt to exclude high-risk individuals from coverage, so that they can insure low-risk individuals for $(p^L\alpha_2, \alpha_2)$ [2]. This can occur either through outright denial of coverage for pre-existing conditions, or absurdly high premiums.

3.2 Claim Denial & Risk Aversion

One limitation of the Rothschild-Stiglitz Model is that insurance is assumed to be a fixed contract, where α_2 is always paid out. However, Schelesinger and Schulenburg extend the Rothchild-Stiglitz model by assuming that consumers are aware of some fixed rate of not being indemnified by their insurers after a claim [5]. They define 3 distinct probabilities: q_1, which represents no accident occurring, q_2, for an accident occurring and being covered, and q_3, for an accident which is not indemnified. Clearly, we see that $p = q_2 + q_3$, but value of a contract to the customer is $U = -\alpha_1 + q_2\alpha_2$, since claims are only paid out with probability q_2. Therefore, since $q_2 < p$ for $q_3 \neq 0$, we see that less insurance is purchased, since each customer behaves like a lower-risk customer than they are.

In the Rothschild-Stigliz model, it's clear that risk adversion leads to an increase in consumption, since insuring against $-d$, which is presumed to be a large loss, is more important than not paying α_1. However, Schelesinger and Schulenburg assert that if even if buying insurance is a risky act, due to the possibility of claim denial, if a risk-adverse buyer's utility is given by $V = g \circ U$, where g is a concave functions, they will still purchase more insurance.

Importantly, Schelesinger and Schulenburg assume that the rate of claim denial is not decided by insurance providers, but the result of circumstances such as insolvency on the part of the insurer, or ambiguity as to which conditions are covered. In modern insurance markets, insurers are able to deny claims based on treatments being “not medically necessary.” Clearly, this implies that absent of regulation, insurance providers can maximize profits by increasing rates of claim denial. Combined with the observation that risk-aversion leads to higher consumption of insurance, insurance companies could raise rates of claim denial until they have consumed the risk-averse buyer's perceived surplus. Empirically, this is observed by insurers of qualified health plans denying 19% of in-network claims, and 37% of out-of-network claims [3].

More recently, Chade and Schlee (2020) [1] reconceptualize insurance as a lemons market. They argue that provision costs—such as administrative processing or fixed costs of underwriting—can create outcomes where insurers deny coverage to those perceived as the worst risks. Their comparative statics theorem shows that once there is no trade at some level of perceived risk, trade fails entirely for worse risks. This formalizes the intuition that insurers deny coverage selectively, targeting the individuals from whom they expect to lose money. They also demonstrate that such markets can exhibit pooling equilibria where all types receive the same contract, challenging the traditional prediction of complete sorting.

Our contribution differs in two major ways. First, we explicitly model claim denial not as an ex-ante contract feature but as a stochastic post-contract event, whose probability is gradually learned by consumers. Second, we simulate how dropout dynamics unfold in a multi-period setting, where belief updating occurs through private experience. This adds a behavioral layer to the lemons logic and highlights the path-dependent nature of market erosion. The interaction of Bayesian learning with claim denial behavior reveals selection patterns that are nonlinear and time-varying—a nuance that static models cannot capture.

Taken together, our work contributes to a growing literature that seeks to understand the real-world frictions in insurance markets—beyond the classical assumptions of full contract execution and instantaneous learning. It supports the emerging view that adverse selection is often shaped not just by hidden types, but also by hidden features of the contract itself, such as the reliability of payout.

4 Insurer-Side Approach

4.1 Motivation

In this work, we want to understand the effects of claim denial rates set by insurance providers, as they are learned by patients through signals. Firstly, we will find the analytical solution to a toy model to show the existence of a non-trivial, finite equilibrium for claim denial. Then, we will model additional factors, such as risk aversion and Bayesian learning by numerical simulation.

4.1.1 Model Parameters

In this model, we consider one monopolistic insurance provider, which sets the insurance contract and one patient, which can choose to stop purchasing insurance at any year. The profit of the insurance provider in given by $U_I = \mathbb{E}_{k,c}[p - C]$, where C is the distribution of healthcare costs covered by the provider, and k is the distribution of years that the patient purchases healthcare. Each patient acts symmetrically and independently, so the maximization of expected profit over time generalizes to the n patient case.

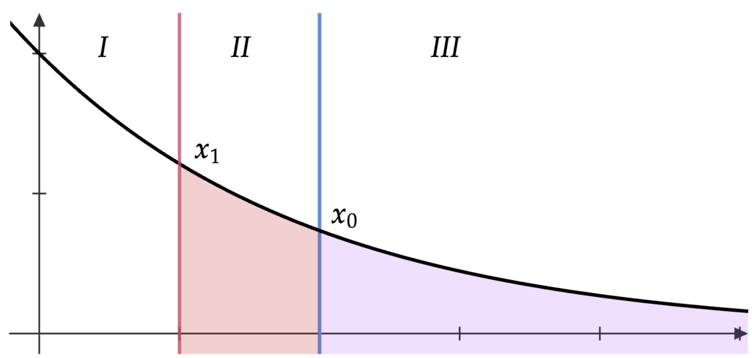

Figure 1: A patient's healthcare claim is drawn from the distribution $x \sim \lambda e^{-\lambda x}$. In region I, where $x < x_1$, the consumer is fully compensated for their claim, as expected. In region II, where $x_0 \leq x \leq x_1$, the consumer's claim is *unexpectedly* denied. In region III, where $x_1 < x$, the consumer is not compensated for their claim, as expected.

Each year, the patient's healthcare needs are drawn from the exponential distribution

$$\Pr(X = x) = f(x) = \lambda e^{-\lambda x} \qquad \Pr(X \leq x) = F(x) = 1 - e^{-\lambda x}$$

which reflects the assumption that worse health events occur less frequently. The insurance contract is governed by (p, x_0), where p is the yearly premium, and x_0 is threshold up to which the insurer will pay all claims. By the memoryless property of the exponential distribution, we can set $p = 1/\lambda$ without loss of generality.

4.1.2 Insurance Utility

Now, consider a *deceitful*, profit-maximizing insurance provider which only compensates claims up to $x_1 < x_0$. To simplify, we will assume that a when a naive consumer files a claim in region II, where $x_1 \leq x \leq x_0$ and is denied, they will learn that they have been deceived and immediately cease purchasing insurance. In any given year, this occurs with

$$\Pr(x \in \text{II}) = \Pr(x_1 \leq x \leq x_0) = F(x_0) - F(x_1) = -e^{-\lambda x_0} + e^{-\lambda x_1}$$

Then, if a patient purchases insurance $k+1$ times, then in the first k turns, they must have submitted a claim $x \notin \text{II}$, which occurs with probability $1 - q = 1 + e^{-\lambda x_0} - e^{-\lambda x_1}$. The $k+1$th turn must be the one where the consumer is unexpectedly denied, so the insurance profits $U_I^{(k)} = 1/\lambda$. Therefore, we can calculate the firm's expected profit as

$$U_I = \frac{1}{\lambda} + \sum_{k=0}^{t} (1-q)^k q \cdot k(p - \mathbb{E}_c[C|x \notin \text{II}])$$

since a firm earns kp from a patient who files k claims where $x \notin \text{II}$, but it costs $k\mathbb{E}_c[C|x \notin \text{II}]$ to compensate those claims. If $x \in \text{III}$, then there is no cost to compensating the claim, so by the law of total probability, we compute

$$\mathbb{E}_c[C|x \notin \text{II}] = \Pr(x \in \text{I}|x \notin \text{II}) \cdot \mathbb{E}_c[C|x \in \text{I}]$$

Then, by the law of conditional probability, we have

$$\Pr(x \in \text{I}|x \notin \text{II}) = \frac{\Pr(x \in \text{I})}{1 - \Pr(x \notin \text{II})} = \frac{e^{-\lambda x_1}}{1 + e^{-\lambda x_0} - e^{-\lambda x_1}}$$

since the firm only has costs if the patient's claim is in region I. The expectation of insuring a claim $x \leq x_1$ is given by

$$\mathbb{E}_c[C|x \in \text{I}] = \int_0^{x_1} x\lambda e^{-\lambda x} dx = \frac{1}{\lambda}(1 - e^{-\lambda x_1}) - x_1 e^{-\lambda x_1}$$

4.1.3 Profit Maximization

For all x_0, the insurance firm makes a profit regardless of what region the consumer's claim lies in, since the premium $p = 1/\lambda$ is the expected cost to the insurer for a policy that covers *all* claims.[1] Therefore, in the limit where $t \to \infty$, or there are infinite turns, the sum U_I diverges at $x_1 = x_0$. We can interpret this as $\Pr(x \in \text{II}) = 0$, so the consumer purchase insurance indefinitely, leading to infinite profits for the firm. However, no firm maximizes profit over an infinite timescale, so we analyze the behavior of the finite case.

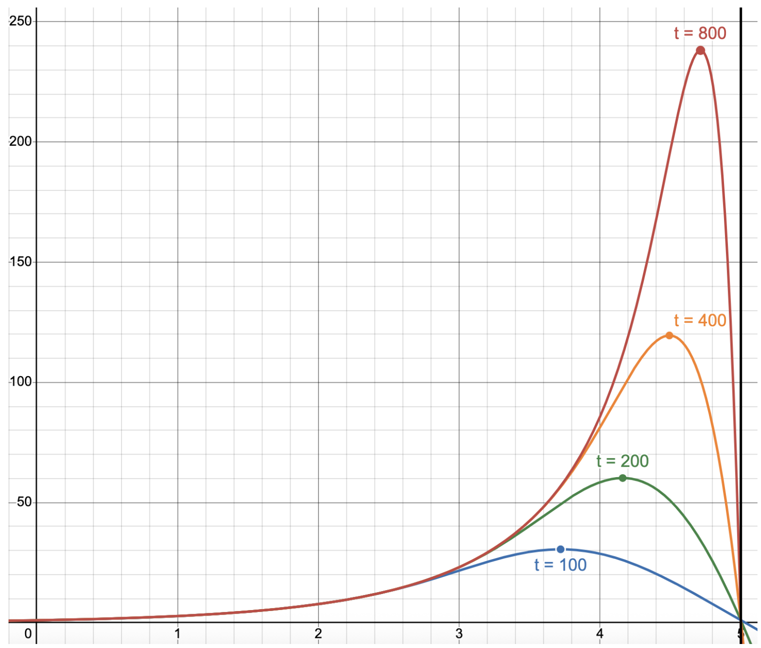

Figure 2: U_I as a function of x_1 for $t = \{100, 200, 400, 800\}$ when $x_0 = 5$ and $\lambda = 1$

In this toy model, we observe the existence of a global maximum value for $x_1 \in (0, x_0)$, where the insurance firm is neither incentivized to deny all claims, nor to compensate all of them. Additionally, we observe that as t increases, so does the optimal value of x_1, which asymptotically approaches x_0, the solution to the infinite case. Mathematically, this phenomenon occurs because avoiding the payment of large claims $x > x_1$ is more profitable than retaining customers indefinitely. Furthermore, when we consider heterogenous risk type, a patient filing a large claim may be an indicator of latent health issues, in which case the firm may not even want to insure them. To explore this phenomenon, we turn to a numerical simulation.

[1]Why would a consumer purchase insurance where the firm always makes a profit? This is a difficult question, since unlike other goods, insurance has a common monetary value. However, we can consider a risk-adverse consumer, or that insurance has the bargaining power to provide healthcare at a lower cost than a consumer would otherwise be able to obtain it.

5 Comsumer-Side Approach

5.1 Model Parameters

To explore the effects of differentiation between patients, we consider a toy model with two patient types in a multi-period setting. In any given period, a *high-risk consumer* is sick with probability s_H and a *low-risk consumer* is sick with probability s_L. This is the only parameter that differentiates these two patient types. If a consumer is sick, then they incur a monetary loss $-c$.

As in Rothschild-Stiglitz, insurance is denoted by the tuple (p, c). When a health loss occurs, the insurer pays claim c with probability $1 - d$, which fully compensates the consumer. However, a claim is denied withprobability d, in which case the consumer incurs the full loss c in addition to paying the premium p.

5.2 Consumer Behavior and Decision Process

Consumers are risk-averse and evaluate outcomes using a concave utility function. We adopt the **Constant Absolute Risk Aversion (CARA)** utility:

$$u(x) = -\exp(-\alpha x),$$

where x is the net monetary outcome. $\alpha > 0$ is the risk-aversion parameter, and higher values indicate stronger risk aversion. This formulation captures diminishing marginal utility and is well-suited for modeling risk in loss scenarios.

5.2.1 Expected Utility Comparison

Let s_i denote the sickness probability for a consumer of type i (with $s_i \in \{s_H, s_L\}$), and let d denote the consumer's *perceived* claim denial rate. Then:

1. **If Insured:** The consumer pays the premium p. In the event of sickness, the claim is approved with probability $1 - d$ and the consumer receives full compensation c, so the net loss is $-p$. However, the claim is denied with probability d, which leads the customer to incur an additional loss c, for a net loss of $-(p+c)$. Therefore, the expected utility when insured is:

$$U_{\text{insured}} = (1 - s_i)\, u(-p) + s_i \left[(1 - d)\, u(-p) + d\, u(-p - c)\right].$$

Substituting the CARA expression for utility, we have

$$U_{\text{insured}} = -\exp(-\alpha p)\,[1 - s_i\,d + s_i\,d\,\exp(-\alpha c)]\,.$$

2. **If Uninsured:** A healthy consumer incurs no cost, whereas if sick the consumer bears the full cost c:

$$U_{\text{uninsured}} = (1 - s_i)\,u(0) + s_i\,u(-c) = -\Big[(1 - s_i) + s_i\,\exp(-\alpha c)\Big].$$

5.2.2 Consumer Decision Rule

A consumer chooses to remain insured if $U_{\text{insured}} \geq U_{\text{uninsured}}$, or

$$-\exp(-\alpha p)\Big[1 - s_i\,d + s_i\,d\,\exp(-\alpha c)\Big] \geq -\Big[(1 - s_i) + s_i\,\exp(-\alpha c)\Big].$$

This condition implicitly defines a threshold premium p^* for a given risk type. The proportion of consumers who remain insured is given by

$$\text{Fraction}_{\text{stay}} = \frac{1}{n}\sum_{i=1}^{n}\text{Stay}_i \qquad \text{Stay}_i = \begin{cases} 1, & \text{if } U_{\text{insured}} \geq U_{\text{uninsured}}, \\ 0, & \text{otherwise.} \end{cases}$$

5.3 Insurer's Profit Function

In a period, t, let N_t be the number of insured consumers, and θ_t be the proportion of high-risk, *insured* consumers. Under complete coverage with claim denial, the insurer pays c for each approved claim, which occurs with probability $1 - d$. Therefore, the expected cost per insured consumer is:

$$\text{cost}(s) = (1 - d)s_i c$$

and the average probability of sickness over the insured pool gives

$$\bar{s} = \theta_t\,s_H + (1 - \theta_t)\,s_L$$

Thus, the insurer's total profit in a period t is:

$$\text{Profit} = N_t p - \mathbb{E}_s(\text{cost}) = N_t p - N_t(1 - d)\Big[\theta_t s_H + (1 - \theta_t)s_L\Big]c$$

5.4 Differential Bayesian Consumer Learning

Each consumer initially holds a prior belief about the claim denial rate d that is concentrated near zero. We model this prior by a Beta distribution:

$$d \sim \text{Beta}(\alpha, \beta),$$

with parameters chosen so that

$$\mathbb{E}[d_{\text{initial}}] = \frac{\alpha_0}{\alpha_0 + \beta_0} \approx 0.$$

Before period 1, setting $\alpha_0 = 0.001$ and $\beta_0 = 1$ yields a prior mean of approximately 0.001.

Consumers update their beliefs about d based solely on their own claim outcomes:

- If a consumer's claim is **denied**, they increment α by 1.
- If the claim is **approved**, they increment β by 1.

Thus, after each period, the consumer's updated estimate of d is:

$$\mathbb{E}[d \mid \text{claim filed period}_i] = \frac{\alpha_{\text{i-1}} + I_{\text{denial}_i}}{\alpha_{\text{i-1}} + \beta_{\text{i-1}} + 1}$$

5.5 Hypothesized Dynamic Effects

- **Declining N_t:** Customers begin dropping out when they experience a negative health outcome at the end of each period AND gets their claim denied. Receiving a claim denial increases an individual's perceived d, making insurance less attractive.
- **Rising θ_t:** The insured pool becomes increasingly composed of high-risk consumers. Low risk consumers experience fewer negative health events and thus although they have the same claim denial rate d, their absolute dropout rate $s_L d$ is smaller than $s_H d$
- **Profit Erosion:** With fewer insured consumers and a riskier composition, the insurer's profit per period,

  $$\text{Profit}_t = N_t\, p - N_t\,(1-d)\Big[\theta_t\, s_H + (1-\theta_t)\, s_L\Big] c,$$

  decreases and may eventually become negative.

5.6 Simulations and Key Findings

To investigate how adverse selection and claim denial impact the dynamics in an insurance market, we construct a multi-period simulation where consumers repeatedly make insurance decisions while updating their beliefs. Here, we set 20% of consumers to be high-risk (with a 30% probability of getting sick) and 80% of consumers to be low-risk (with a 10% probability of getting sick). If a consumer has insurance and falls ill, they will file a claim. If that claim is accepted (with the denial rate set to 20%), the insurer pays out \$120 per approved claim—the same cost faced by a sick and uninsured consumer. When a claim is filed, the consumer updates their belief about the insurer's denial rate using Bayesian learning, which in turn influences their future decision to stay insured.

We implemented these parameters in a multi-period simulation across a wide range of premiums to understand breaking points: $p \in \{10, 11, 12, 13, 14, 15, 20, 25, 73, 74, 95\}$. The simulations are averaged across 1000 independent runs per premium to ensure robustness.

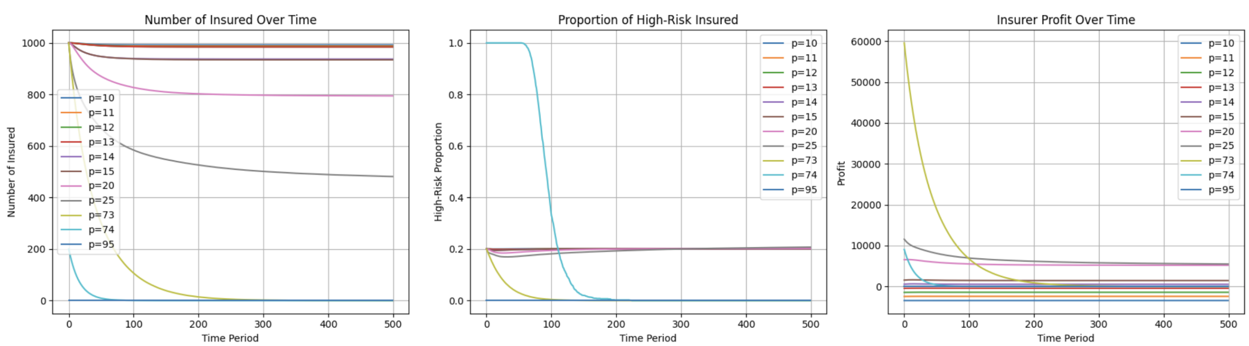

Figure 3: Market variables under different premiums.

We notice some clear patterns begin to emerge:

- **Declining N_t:** As expected, lower premiums lead to higher and more stable participation. At $p = 10$, a very large majority of all consumers stay insured across 30 time periods. Conversely, at $p = 73$, the number of insured individuals declines significantly as many consumers, particularly low-risk ones, exit the market in response to high prices and evolving beliefs. At $p = 74$, we observe that it is not worth buying insurance for low-risk consumers and that at $p = 95$, it is not worth buying insurance for any consumer.

- **Risk Pool Composition θ_t:** From the second graph, we observe once again the fully high-risk composition of insurers at $p = 74$ and no high-risk composition of insurers at $p = 95$. Most of the mid-range premiums seem to stay around the starting high-risk proportion of 0.2 (though we will inspect this further below). Interestingly, we see that at $p = 73$—the highest premium where all consumers buy insurance—, θ_t declines. This can be explained by the high premium acting as a filter. After just one or two denied claims, the perceived value of insurance drops dramatically for these high-risk consumers who are more likely to file claims early-on. As a result, they will continuously exit the market as soon as they are denied.

- **Profit Erosion:** Profits generally rise with premium level though profit growth shrinks as consumers leave. At low premiums (like $p = 10$), insurers struggle to cover expected payouts, leading to near-zero or negative profit. Intermediate premiums like $p = 20$ and $p = 25$ offer a more balanced trade-off—moderate profits with a relatively stable insured population. Just below these mid-range premiums, a premium of $p = 74$ where only high-risk consumers buy insurance will result in worse initial profit and steeper profit decline. We also observe how at $p = 73$, insurance companies will make the most profit but at the steepest profit decline.

Zooming in on mid-range premium levels where all types of consumers buy insurance, we see an interesting phenomenon across all the premiums in the middle graph of Figure 3.

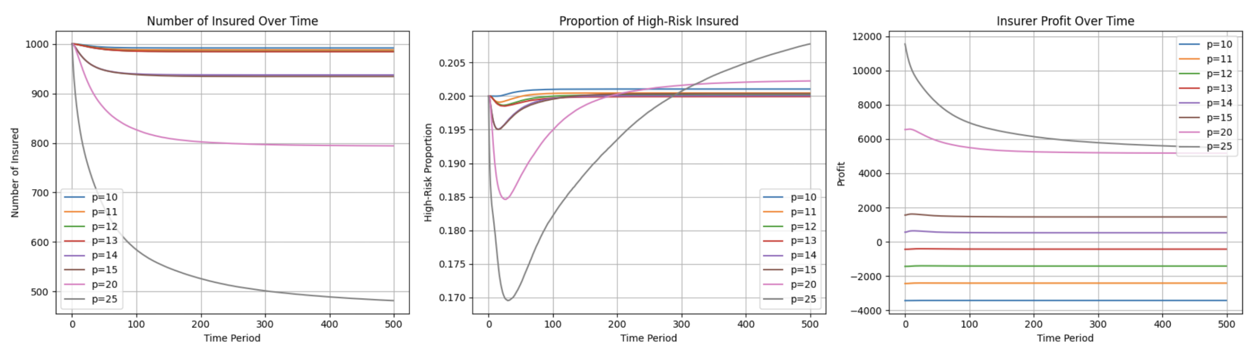

Figure 4: Market variables under reasonable premiums.

Looking at the proportion of high-risk insured consumers against time period, we see that θ_t initially drops before increasing. This behavior contradicts the typical expectation in models of adverse selection, where we assume that as time progresses, low-risk individuals drop out faster, and the insured pool becomes increasingly composed of high-risk consumers. In this model, high-risk consumers are more likely to drop out in the early periods, driving the initial dip in θ_t.

We hypothesize that this early dropout among high-risk individuals is driven by their greater likelihood of experiencing a claim event early in the simulation. Since claims can be denied, these early experiences lead high-risk individuals to revise their beliefs about the claim denial rate more drastically, and will choose to dropout. In contrast, low-risk consumers may remain insured longer simply because they have not yet filed a claim and thus have not encountered a negative event that would cause them to update their beliefs.

This dynamic shows how adverse selection may not be monotonic or immediate—especially in a market with differential privacy. Differential Bayesian learning-based dropout can introduce a nonlinear pattern in the composition of the insured pool, with early exits skewed toward those with higher prior risk, temporarily improving the pool quality before the usual adverse selection trend takes over.

6 Conclusion

Our project reveals how modeling insurance decisions as a Bayesian, multi-period learning process exposes hidden dynamics that traditional models overlook. By allowing consumers to update beliefs based on personal experience—particularly the outcome of claim filings—we uncover a non-monotonic path in the evolution of the insured pool's risk composition. While classical models predict a steady rise in risk concentration under adverse selection, we found that the proportion of high-risk consumers may initially decline, as they disproportionately exit the market due to earlier negative experiences. This creates a temporary window where the pool improves before long-run adverse selection dynamics take over. This result is important for understanding the true time path of insurance markets since it suggests that short-term pool improvements may mask long-term instability, and that pricing strategies should consider the belief trajectories of consumers in addition to static incentives.

Future work could explore:

- Interventions like public claim approval statistics or consumer education campaigns to slow harmful belief updating.
- Dynamic pricing or bonus incentives to retain early high-risk consumers and reduce churn.
- Insurer-side learning and strategy adaptation in response to belief-driven dropout patterns.

References

[1] Hector Chade and Edward E. Schlee. "Insurance as a Lemons Market: Coverage Denials and Pooling". In: *Journal of Economic Theory* 190 (2020), p. 105085.

[2] Liran Einav and Amy Finkelstein. "Selection in insurance markets: Theory and Empirics in pictures". In: *The journal of economic perspectives: a journal of the American Economic Association* (2011). URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC3248809/.

[3] Michelle Long Justin Lo and Jan 2025. "Claims denials and appeals in ACA Marketplace plans in 2023". In: *KFF* (Feb. 2025). URL: https://www.kff.org/private-insurance/issue-brief/claims-denials-and-appeals-in-aca-marketplace-plans-in-2023/#:~:text=Key%20Takeaways,by%20insurer%20and%20by%20state..

[4] Michael Rothschild and Joseph Stiglitz. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information". In: *The Quarterly Journal of Economics* 90.4 (1976), pp. 629–649. ISSN: 00335533, 15314650. URL: http://www.jstor.org/stable/1885326.

[5] Harris Schlesinger and J.-Matthias Graf v. d. Schulenburg. "Risk Aversion and the Purchase of Risky Insurance". In: *Journal of Economics* 47.3 (1987), pp. 309–314. ISSN: 09318658, 16177134. URL: http://www.jstor.org/stable/41793862.