Consumer Beliefs and Learning about Health Insurance Plan Characteristics

Tian Liu

November 16, 2024

Click Here for the Latest Version

Abstract

Models of consumer decision-making in healthcare usually assume that consumers are fully informed about the characteristics of their health insurance plans. In practice, however, health insurance plans are quite complex, which could lead consumers to make sub-optimal healthcare decisions. In this paper, I develop a dynamic structural model of consumer decision-making in healthcare, allowing for the possibility that consumers may initially hold biased beliefs about their effective coinsurance rates (the average percentage of medical expenditures that are paid by consumers), which they may learn about via the rates at which their medical expenditures are reimbursed by their health insurance plans. The model allows for corner solutions in medical expenditures, which affects learning because consumers only receive reimbursements when spending positive amounts. The model also incorporates endogenous health production, which is necessary to separate the spending dynamics arising from learning and those stemming from health transitions. I estimate the model using data from the RAND Health Insurance Experiment, a large-scale experiment conducted in the US. The results suggest that consumers' initial beliefs about their effective coinsurance rates significantly deviate from the true rates, and that consumers may learn from their reimbursement rates, but the infrequency of positive spending and the noisiness of signals slow their learning. I simulate a counterfactual where consumers are assumed to have full information and find that consumers' medical spending and health are substantially distorted under biased beliefs about their plan characteristics, which significantly harms welfare. I also simulate a counterfactual in which consumers receive signals even when they consume no medical services, which shows that, even with a moderate gain in information, consumers could be significantly better off.

1 Introduction

Economists have long been concerned with how to model consumers' medical care decision-making. Researchers usually assume that consumers have perfect information about the characteristics of their health insurance plans (Gilleskie, 1998; Einav et al., 2013; Cronin, 2019). While this assumption may be a natural starting point, health insurance plans are notoriously complicated. Consumers, especially those who need health insurance the most, e.g., the elderly or individuals with low levels of education, may not understand key terms in health insurance plans (Loewenstein et al., 2013). They may also struggle to recall the characteristics of their current plans (Handel and Kolstad, 2015). Therefore, the assumption of perfect information may be unrealistic. Consumers may hold biased beliefs about their plan characteristics.

Consumers' biased beliefs about their plan characteristics have implications for the analysis of healthcare markets. On the one hand, they may significantly distort consumers' utilization of medical services with enduring effects on consumers' health. On the other hand, even if consumers' beliefs are biased, they may also update their beliefs if they observe information deviating from their expectations. However, while there have been studies recording consumers' biased beliefs about their plan characteristics, few papers have considered linking it with consumer learning. This may matter because, if the speed of learning is fast, the potential distortions in healthcare utilization arising from biased beliefs could be significantly reduced.

The RAND Health Insurance Experiment (HIE) offers an ideal context to study this issue. The HIE was a large-scale experiment conducted in the US from 1974 to 1982. Participants were assigned to plans with different coinsurance rates.[1] It has been found that consumers are often uninformed about their coinsurance rates. For example, Loewenstein et al. (2013) find that only 34% of their sample understand the term "coinsurance". Handel and Kolstad (2015) also find that roughly 30% of employees in a big company could correctly recall the coinsurance rates of their plans. Since the HIE contracted with a private insurance company to process claim forms, the participants in the experiment were placed in an environment similar to those in other contexts. Therefore, their beliefs about coinsurance rates may, similarly, be significantly biased.

Consumers may also hold biased beliefs about their coverage. The HIE did not cover all medical services. Therefore, even if a consumer was assigned to a plan with a 0% coinsurance rate (a free plan), their out-of-pocket payments were not always zero. For example, non-prescription drugs were only covered for certain conditions, or large dental expenditures may need prior authorization. In those cases, the HIE would not reimburse the charges. As a result, the participants in the HIE may simultaneously misunderstand their coinsurance rates and which services are covered by their plans.

In this paper, I develop a dynamic structural model of healthcare decision-making, allowing for the possibility that consumers may hold biased beliefs about their *effective coinsurance rates*, i.e., the average percentage of medical expenditures that are actually paid by consumers, but may learn about them via the rates at which their medical expenditures are reimbursed. The possibility of biased beliefs about effective coinsurance rates suggests that there may be scope for learning.

I use the effective coinsurance rate to represent both the coinsurance rate and the coverage for a specific plan because consumers may not be able to distinguish between them in their reimbursement statements. For example, a reimbursement statement may include details for multiple services, potentially hindering consumers' understanding of which portion of their spending is not covered. It may also group services under broad categories, e.g., inpatient

[1] Coinsurance rate refers to the percentage of covered medical expenditures that are paid by consumers.

services or drug expenditures, which combine both covered and non-covered services, thereby making it difficult for a consumer to discern the amount that is not covered without a clear understanding of their coinsurance rates.

Investigating the effects of biased beliefs and learning in a dynamic structural model offers two advantages over reduce-form approaches. First, studying the issue using a reduced-form method typically requires researchers to collect survey data on consumers' beliefs over multiple periods, a task that can be both challenging and costly. With some standard assumptions, a dynamic learning model could estimate consumers' beliefs without belief data. Specifically, the model developed in this paper can be estimated if data on spending, reimbursements and health are available. Such information could be found in various administrative datasets. Second, a structural model could be used for ex-ante policy evaluation, e.g., how much consumers would spend on their medical services when they are fully informed. A reduced-form method may be ill-suited because healthcare decisions made under full information are not observed for consumers with biased beliefs.

The model builds upon the Bayesian learning models in the dynamic demand literature (Erdem and Keane, 1996; Ackerberg, 2003; Crawford and Shum, 2005; Mehta et al., 2017). In each period, a consumer first decides whether to keep their current plan (versus an outside option), whether to spend and how much to spend on medical services, based on their health state and belief about their effective coinsurance rate. If they choose to spend, they will receive a reimbursement statement from their provider, which details the fraction of healthcare spending that is reimbursed by their health insurance plan. They can compute the proportion of out-of-pocket payments (*experienced coinsurance rate*) using this information. Based on the experienced coinsurance rate (signal), they update their belief using a Bayesian formula. Their health may also improve based on the amount of medical spending. If they decide not to spend, they receive no signal about their effective coinsurance rate, and thus, their belief remains the same. This consumer then repeats these steps in the following periods until the experiment ends.

While identifying the parameters that characterize learning is heavily dependent on spending dynamics, the inclusion of endogenous health transitions is essential. For example, if a consumer reduces their spending after incurring spending in the previous period, this may be because this consumer has observed an experienced coinsurance rate (signal) that is higher than what they thought and, thus, revises their belief about the effective coinsurance rate upwards. Because now they believe the fraction of out-of-pocket payments is higher, they reduce their spending in the current period. Therefore, how spending evolves after signals is crucial to the identification of learning. However, this reduction in spending could also result from an improvement in their health, which reduces their demand for medical services. As a result, it is important to account for the health transitions to isolate the spending dynamics arising from learning.

The model is estimated via maximum likelihood using data from the HIE. Participants were randomly assigned to plans that differed along three principal dimensions: coinsurance rates, maximum dollar expenditure (MDE), and years of enrollment.[2] Plans with a uniform

[2]MDE is the out-of-pocket limit during the corresponding contract year. If services were covered by the experiment, consumers' out-of-pocket payments would be zero when their MDEs are reached. I assume that consumers are well-informed about their MDE.

coinsurance rate for all covered services, i.e., the free plan, the 25% coinsurance rate plan, and the 95% coinsurance rate plan, are used in this paper.[3] The HIE provides claims-level information on medical expenditures and reimbursements, as well as detailed health surveys and screening examinations at enrollment and exit. In contrast to purely administrative datasets, even if a person consumes no medical services, their health is still observed, which facilitates the identification of endogenous health production. To mitigate the concern for measurement errors, the paper allows for observed health measures, spending, and reimbursements to be noisy measurements of the latent ones.

The results indicate that there is substantial heterogeneity in consumers' prior beliefs. While roughly 60% of consumers hold beliefs quite close to the true rates, the remaining have initial beliefs that are substantially different. The estimated effective coinsurance rates in the free plan, the 25% coinsurance rate plan, and the 95% coinsurance rate plan are 2.64%, 28.24%, and 88.49%, respectively, but the group of consumers whose initial beliefs are more biased have initial means that are 34.90%, 4.94% and 17.62%, respectively. Consumers also learn very slowly during the experiment. When updating for the first time, the weights consumers put on signals are less than 2% in all three plans.[4] Therefore, consumers' beliefs show limited improvement after 3 years, when the majority of the population had left the experiment.

To quantify the importance of biased beliefs and learning, I simulate a counterfactual where consumers are assumed to have full information about their effective coinsurance rates and compare it with the base case which is simulated using the original estimates. The medical expenditures for consumers with relatively accurate beliefs are similar to those of consumers with full information. Compared to the group of consumers with more biased beliefs, consumers with full information, on average, spend \$41.37 (192.96%) more in the free plan, \$34.22 (53.06%) less in the 25% coinsurance rate plan, and \$30.82 (69.45%) less in the 95% coinsurance rate plan on a four-week basis. As a result, their health, on average, is 25.00% better in the free plan, 14.29% worse in the 25% coinsurance rate plan, and 22.22% worse in the 95% coinsurance rate plan. Consumers' discounted experienced utility at enrollment is also significantly higher with full information: compared with the group of consumers with more biased beliefs, they, on average, receive \$2465.6 more utility in the free plan, \$6820.4 more utility in the 25% coinsurance rate plan and \$40341.8 more utility in the 95% coinsurance rate plan.[5] The results indicate that consumers' biased beliefs could significantly distort their healthcare decisions and health outcomes, which may also deeply harm their welfare.

I also simulate a counterfactual where consumers receive signals even when they consume no medical services. In all three plans, consumers are better off when receiving more information. Consumers in the 95% coinsurance rate plan with more biased beliefs receive the highest gain in information due to their relatively higher rate of learning and more frequent

[3]The plan with a 50% coinsurance rate is dropped because there would be too few observations (6 individuals) after sample restrictions.

[4]Under the Bayesian updating framework, a consumer's mean of belief after observing a signal would be a weighted average of the signal and the previous mean.

[5]Experienced utility refers to the actual utility that consumers experience when their effective coinsurance rates are realized. The ex-ante utility may not be appropriate to measure consumers' welfare because their beliefs may be biased, and their ex-ante utility may significantly deviate from what they actually experience.

learning opportunities: After 3 years, the average difference between their means of beliefs and the true rate is -62.39% in the base case, but it increases to -44.89% with more frequent signals. Although their beliefs still differ significantly from the true rate, they substantially improve their decision-making by receiving more frequent signals: compared to consumers in the base case, they, on average, reduce their spending by $7.97 during the experiment and increase their discounted experienced utility at enrollment by $15632.4. This improvement indicates that consumers may be significantly better off even with a moderate gain in information.

This paper contributes to several literatures. First, I present additional evidence that consumers may not optimize their healthcare decisions and, therefore, have bounded rationality in healthcare markets. Previous literature, particularly studies on the Medicare Part D program, usually demonstrates the presence of consumers' bounded rationality by examining whether consumers choose their health insurance plans optimally (Abaluck and Gruber, 2011, 2016; Ketcham et al., 2012, 2015; Ho et al., 2017; Kaufmann et al., 2018; Handel et al., 2024).[6] However, the bounded rationality discussed in these studies may stem from consumers being overwhelmed by an excessive number of plan options. In contrast, my research focuses on a different source of bounded rationality, i.e., biased beliefs about the characteristics of their current plans, and shows that some consumers are far from being rational.

Second, this paper presents novel evidence that consumers may learn slowly about the characteristics of their current plans via their experiences. Some of the above studies on consumers' bounded rationality also discuss learning in the context of plan choices and have not uniformly found evidence of its presence (Ketcham et al., 2012; Kling et al., 2012; Abaluck and Gruber, 2016; Kaufmann et al., 2018). Two other papers present evidence of learning in consumers' current plans. Hodor (2021) uses the same dataset and finds that if families share the same health insurance plans, they may increase their medical expenditures after the healthcare utilization of other family members. She interprets this increase as evidence of learning about health insurance or the effectiveness of healthcare but cannot distinguish between them. Anderson et al. (2024) find that a bill's arrival causes households to reduce their spending, indicating that households may learn their medical expenditures from bills. Contrary to theirs, I find that consumers may learn very slowly about their plan characteristics.

Third, I extend the literature that estimates structural models of consumers' medical care decision-making. The existing literature typically assumes that consumers are well-informed about the characteristics of their health insurance plans (Gilleskie, 1998; Einav et al., 2013; Cronin, 2019). Two papers leverage survey data on consumers' beliefs about health insurance to quantify their welfare losses resulting from information frictions, without considering the possibility of learning (Handel and Kolstad, 2015; Boyer et al., 2020). To the best of my knowledge, Anderson et al. (2024) is the only paper that estimates a model where households could learn about their out-of-pocket cost without survey data on beliefs. They assume that households may overestimate or underestimate their actual out-of-pocket costs by a constant rate before receiving the bill, and households learn about the rate via

[6]Medicare Part D program provides prescription drug coverage for elders (age 65 or older) and younger people with disabilities in the US.

their out-of-pocket costs. Their approach assumes a common prior across all plans and that health evolves exogenously, partly due to their larger number of plans and the lack of available health information. Compared to theirs, my paper introduces a different source of biased beliefs: the effective coinsurance rate, which varies across plans. My model also incorporates endogenous health transitions, which may help separate the spending dynamics arising from learning from those driven by health changes. Furthermore, this model allows for heterogeneous priors within plans, which could help deepen our understanding of how biased beliefs influence healthcare decisions and welfare. For example, some consumers might have relatively unbiased beliefs and, as a result, may be less impacted by these biases.

Fourth, this paper contributes to the dynamic demand literature on learning models. The majority of these papers assume that consumers have rational expectations (Erdem and Keane, 1996; Crawford and Shum, 2005; Arcidiacono et al., 2024).[7] The assumption of rational expectations may not be realistic in my setting because there is ample evidence that consumers may hold biased beliefs about their plan characteristics. Therefore, this paper allows for heterogeneous initial beliefs about effective coinsurance rates that may deviate from the true rates. As noted by Ching et al. (2013), it is difficult to identify complex learning models with rich specifications of consumer behaviour because it may be difficult to understand what data patterns help identify the parameters or the likelihood is so flat that it is not practical to estimate it. Therefore, it may be difficult to identify all parameters in a complex learning model without the rational expectation assumption. In this paper, the heterogeneity in consumers' priors can be identified because the signals (or experienced coinsurance rates) are noisily observed, i.e., signals can be computed using reimbursements and medical spending, both of which are assumed to be noisily observed. Nosily observing signals helps to identify the distributions of signals. How consumers with identical information respond to different noisy signals also helps to identify the means of beliefs, which significantly facilitates the identification of the parameters characterizing learning.

Finally, this paper can be viewed as an empirical application of Grossman (1972), which constructs a lifetime utility-maximizing model of demand for health. In his model, health is both a consumption good and an investment good that can be enhanced through the utilization of healthcare resources, thereby resulting in greater future utility. All empirical papers following Grossman (1972) use one or two particular health measures to evaluate individual health without considering the possibility of measurement errors (Gilleskie, 1998, 2010; Darden, 2017; Cronin, 2019). If health is noisily measured, estimates of health production functions may not be consistently estimated. Thus, I use factor analysis to estimate a distribution of latent health to mitigate the concern for measurement errors. Factor analysis has been widely used in the economics of education, especially in studies of cognitive and noncognitive abilities, to reduce dimensionality and allow for measurement errors (Cunha and Heckman, 2008; Cunha et al., 2010), but is seldom used in the health literature (Liu et al., 2023; White, 2023).

The paper is organized as follows. Section 2 details the data. Section 3 provides the details of the model. Section 4 describes the empirical strategy. Section 5 presents the

[7]The rational expectation assumption means that the mean of initial beliefs is the same as the true mean of the population.

estimation results and the model fit. Section 6 presents the counterfactuals to measure the importance of biased beliefs and learning. Section 7 is the conclusion.

2 Data

2.1 Description of Data

The data come from the RAND Health Insurance Experiment (HIE), a large-scale experiment in the US conducted from 1974 to 1982. The goal of the experiment was to assess how cost-sharing affected the utilization of medical services and health outcomes. The HIE provides detailed data on assigned health insurance plans, demographic characteristics, individual health, and medical services utilization, which is very suitable for this study.

In the HIE, each family was randomly assigned to a plan that differed in three principal dimensions: coinsurance rates, MDEs (out-of-pocket limits), and years of enrollment. This paper only uses the plans with a uniform coinsurance rate for covered services.[8] Their coinsurance rates were as follows: 0% (free care), 25%, and 95%. These plans were divided further into three possible MDE levels: 5%, 10%, or 15% of family income with a maximum of $1,000 or $750. Years of enrollment could be either three or five years. Families were offered participation incentive payments (Participation Incentive) if they decided to join and stay in the experiment. To test whether families were sensitive to their participation incentive payments, the experiment offered the Super Participation Incentive, a bonus participation incentive payment in the next-to-last year of the assigned enrollment term.

Several other restrictions are imposed on the sample. First, the analysis includes individuals with no family members. Adding people with families significantly complicates the analysis of learning.[9] Second, if a person left the experiment not because of attrition, i.e., death, suspension, or termination, then I keep only the observations in periods one year before they left the experiment.[10] Third, I remove periods when the MDE is not reached, but actual rates of out-of-pocket payments are lower than the coinsurance rate, which may be due to administrative mistakes.[11]

Before the HIE assigned plans to each family, they used a screening questionnaire or a baseline interview that gathered some basic demographic information to establish eligibility for enrollment, which provides a list of individual characteristics, including age and education

[8] Having plans with two coinsurance rates significantly complicates the analysis of learning. For example, if the coinsurance rates for dental services and mental services are different, consumers may learn about the coinsurance rate of dental services from consuming mental services.

[9] For example, a person may learn from their family members' medical bills, suggesting strategic decisions could be present within the family.

[10] Common reasons for termination outright are: (1) accepting health care under Medicaid, Medicare disability or a similar program; (2) being uncooperative or unlocatable. Common reasons for suspension are: out of the country, military duty, and confinement to a penal or other government supporting institutions, where health care options are very different. These people may have very different outside options than others and, thus, are excluded from the sample.

[11] An example of potential administrative mistakes is as follows: the insurance company may simultaneously receive multiple claims from one consumer, but reimburse a late bill first because the claims are not ordered by date. This may cause some early bills to be fully reimbursed because the MDE is reached after reimbursing the late bill.

years. The HIE also collected a comprehensive set of health measures. All participants who exited the experiment normally, i.e., did not attrit or were not ineligible, were asked to fill out a questionnaire at both enrollment and exit. Individuals who left earlier usually filled out their questionnaires only at enrollment. There were also two medical screening examinations implemented in the experiment. A subset of families (50 to 75 percent, varying by site) were randomly chosen to finish a medical screening examination at enrollment, and all families were asked to finish an examination at exit.

To evaluate individual health, I construct two health measures using a large number of raw variables from the HIE. The first measure is the number of conditions from screening examinations.[12] The HIE has already discretized the condition status in screening examinations. I encode the condition status as binary and sum them up to create my first health measurement. The conditions included in the number of conditions from screening examinations are anemia, COAD, diabetes, hearing loss, hypercholesterolemia, hypertension, far vision, and thyroid. The second measure is the number of disorder concerns collected from questionnaires. Most conditions have the following questions in the survey: "During the past 30 days, has this condition concerned or worried you?". I encode this value as one if an individual has any concern regarding this condition, and zero if this individual has no concern or is never diagnosed by a doctor to have this condition. The conditions included in the number of disorder concerns are eyesight, glaucoma, hearing, hay fever, gum, thyroid, joint, shortness of breath, blood pressure, chest and heart, bronchitis, tuberculosis, stomach, kidney, cholesterol, anemia, diabetes, cancer, hemorrhoids, hernia, and varicose veins.

Table 1 presents the summary statistics for individual characteristics and health at enrollment and exit. Eventually, 174 people are included in my sample. At enrollment, their average years of education was 12.74 and their average age was about 40. Most people were offered a 3-year plan. The majority had at least one disorder concern or one condition diagnosed from their screening examinations. Therefore, the means of the two health measurements are all above one at both enrollment and exit.

[12] The sum of conditions or diseases has been widely used in previous literature (Cronin, 2019; Hosseini et al., 2022).

Table 1: Summary Statistics for People Included in the Sample

	Mean	Std. Dev.
Years of education at enrollment	12.74	3.13
Age at enrollment	40.46	14.70
MDE (Out-of-pocket payments, measured in dollars)	244.47	354.88
Participation Incentive (measured in dollars)	254.55	328.43
Years of enrollment	3.49	0.87
Super Participation Incentive [1]	0.37	0.49
Num. of disorder concerns at enrollment	2.06	1.85
Num. of conditions in screening examinations at enrollment	1.17	1.09
Num. of disorder concerns at exit	2.01	2.02
Num. of conditions in screening examinations at exit	1.42	1.24
Num. of people	174	

[1] Whether the person was offered a bonus participation incentive payment in the next-to-last year of the assigned enrollment term. The Super Participation Incentive was offered only to people who were assigned to the 95% coinsurance rate plan.

The experiment also provided line-item data for medical expenditures. Each instance of a billed service, drug, or supply on a claim form is called a "line item". Line-item records were organized into 14 files according to the type of medical or dental services rendered. For each line item, the date of service, line-item charge, and reimbursement amount were recorded. The date of leaving the experiment for each person was also recorded. Based on the line-item data and leaving date, I can reconstruct everyone's medical expenditure history during the experiment.

Table 2 presents the summary statistics for spending on a four-week basis. There are 99 people in the free plan, 30 in the 25% coinsurance rate plan, and 45 in the 95% coinsurance rate plan. Since each period is four weeks, the total observations for spending are 4252 for the free plan, 1061 for the 25% coinsurance rate, and 1200 for the 95% coinsurance. People with the free plan tend to spend more and also more frequently than people with the other two plans, which is reasonable as they face the lowest marginal financial cost.

2.2 Uncertainty about Experienced Coinsurance Rates

Consumers' experienced coinsurance rates may not equal the coinsurance rates, e.g., consumers might need to pay even if they are in the 0% coinsurance rate plan (the free plan). This is because some types of services may not be covered by the HIE. The reasons for

Table 2: Summary Statistics for Spending

Plan Type	Pr. of Zero Spending	Positive Spending[1]	Indiv. Obs.	Person-Period Obs.
Free	0.501	160.525 (700.53)	99	4252
25% coin. rate	0.620	108.419 (213.264)	30	1061
95% coin. rate	0.780	106.797 (371.122)	45	1200

[1] Positive spending refers to the average spending when consumers spend a positive amount. Standard deviations are presented in parentheses. Each period lasts 4 weeks.

noncoverage include: (1) Some charges were not covered if not proved to be medically necessary. For example, charges for private rooms were only covered when they were certified as medically necessary or no semiprivate room was available. Another example would be that non-prescription drugs were only covered for certain conditions such as chronic allergic conditions or chronic respiratory diseases. They were covered up to $100 per person for each condition each year. Most importantly, there must be physicians certifying that the condition existed before payment was made. (2) Charges that were compensable under worker's compensation, employers' liability laws, or automobile accident insurance policies were not covered until they had been exhausted. (3) Large expenditures may need prior authorization. For example, prior authorization was needed for any treatment plan for dental issues exceeding $500 (except emergency care) and for the replacement of crowns, bridges, or dentures. (4) Coverage for some services may be limited to a certain amount, e.g., only one eye examination was covered per year, and only one pair of corrective lenses per year was covered. (5) Some other charges may not be covered, e.g., cosmetic surgery except for damage arising from accidental injury, most orthodontics, and most cosmetic dental services.

Table 3 presents the summary statistics for experienced coinsurance rates when the MDE is not reached. In all three plans, the means of their experienced coinsurance rates differ from the coinsurance rates. The standard deviations are significant in the free plan and the 25% coinsurance rate plan, suggesting that it is common that experienced coinsurance rates deviate from coinsurance rates. I also test whether the means of experienced coinsurance rates are equal to their coinsurance rates for all the plans. The results are displayed in the third column of Table 3. For all three plans, the hypothesis that the means of experienced coinsurance rates are equal to their coinsurance rates can be rejected. Therefore, even if consumers can understand the term "coinsurance", they may still hold biased beliefs and be uncertain about their experienced coinsurance rates, because they may not be aware of which services are covered by the experiment.

2.3 Preliminary Evidence of Learning

Before presenting the model, I test for evidence of learning. Intuitively, if learning does not exist, then consumers with the same individual characteristics, plans and health would spend the same amount regardless of how many signals they have observed. Specifically, I estimate the regressions of the log(spending+1) on the number of observed signals to test whether learning exists in the HIE. The number of observed signals equals the number of periods that consumers have non-zero spending and their MDEs are not reached because consumers are assumed to observe a signal during periods in which they choose to spend and

Table 3: Summary Statistics for Experienced Coin. Rates when the MDE is not Reached

Plan Type	Mean	Std.	P-Value (Test Mean Diff. =0)[1]
Free	0.025	0.111	0.00
25% coin. rate	0.278	0.085	0.00
95% coin. rate	0.952	0.008	0.00

[1] Test whether the means of experienced coinsurance rates are equal to the corresponding coinsurance rates. For example, in the free plan, I test whether the mean of experienced coinsurance rates equals zero.

Table 4: Regression of log(spending+1) on the Number of Signals Observed

	free log(spending+1)	25% coin. rate log(spending+1)	95% coin. rate log(spending+1)
Num. of signals	0.090**	0.087**	0.101**
	(0.004)	(0.012)	(0.012)
Num. of person-period obs.	4250	1060	1195

Standard errors in parentheses.

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$

Coefficients on the other controls are presented in Table 12 in the Appendix.

have not yet reached their out-of-pocket limits.[13] If learning does not exist, the coefficient on the number of signals should be insignificant. The regressions are estimated separately for different plans, as people in different plans may have different beliefs about their effective coinsurance rates. Individual characteristics are added as controls for preference and initial health. Consumers' health during the experiment is not observed, which makes it difficult to interpret the coefficient on the number of signals. For example, consumers who observe more signals may be those who are sicker and, thus, demand more medical services. Therefore, the coefficient on the number of signals may result from changes in health rather than learning. To alleviate this issue, I add past spending as another control. Elapsed periods are also added as a control for time trends.

Table 4 presents the regression results. For all three plans, the coefficients on the number of signals are positive and significant. This indicates that consumers who observe more signals tend to spend more, regardless of which plans they are in, which may suggest that consumers have initial beliefs that lie above the truth. As they accumulate more knowledge about their plans, they find that their plans actually cost them less than they thought and, thus, choose to spend more afterward. However, because health is not perfectly controlled, it is important to construct a structural model that controls the health transitions to offer more convincing evidence of learning.

[13] The periods in which consumer reach their MDE may not provide informative signals because the reimbursements are not based on the coinsurance rates.

3 Model

This section presents my model of learning, where a consumer makes healthcare decisions to maximize their expected discounted lifetime utility based on their health and beliefs about the effective coinsurance rate. Time is discrete. I present the model for one consumer i, so I suppress that index in this section.

At enrollment, a consumer is randomly assigned to an experimental plan, p. During the experiment, this consumer is forward-looking and makes decisions based on their expected discounted lifetime utility. Specifically, the timing of the model in each period is as follows:

1. Learns the state, e.g., health, beliefs about the effective coinsurance rate and current MDE
2. Chooses whether they want to leave the experiment, whether to spend and how much to spend given their state
3. If they choose to spend and their MDE is not reached, they will receive a signal (the experienced coinsurance rate) and update their belief based on that signal. Otherwise, their belief remains the same.[14]

This consumer repeats the above steps until he attrits or the experiment ends.

3.1 The Health Production Function

A consumer's health evolves stochastically during the experiment. Denote health in period t as H_t. A consumer's health in the next period, H_{t+1}, is determined by their current health H_t, their current medical spending m_t, their age at enrollment age_1, and a stochastic health shock ϵ^H_{t+1}. I use a Cobb-Douglas specification for the health production function to capture the non-linearity of the health transition process, e.g., people with different health levels may have their health depreciate at different rates and their benefits of receiving medical care may also differ. The health production function is as follows:

$$\log H_{t+1} = \gamma^H_1 + \gamma^H_2 \log H_t + \gamma^H_3 \log(m_t + 1) + \gamma^H_4 log(age_1) + \epsilon^H_{t+1}$$

where ϵ^H_{t+1} is a mean-zero IID shock that follows a normal distribution. Medical spending, m_t, cannot enter the production function directly with a log function because $\log(m_t)$ is meaningless when $m_t = 0$. I therefore use $\log(m_t + 1)$ to represent the effect of medical care on health, which allows the effect of medical spending to be zero when people spend nothing.

3.2 The Utility

The per-period utility in period t with an experienced coinsurance rate θ_t and no out-of-pocket limits follows the following form:

[14]The periods in which consumer reach their MDE may not provide informative signals because the reimbursements are not based on the coinsurance rates.

$$U_t(\Omega_t, m_t) = \frac{H_t^{1-r} - 1}{1-r} - \exp(r^c \theta_t m_t) + \alpha^{PI} PI_t - \mathbb{1}_{(m_t>0)} c^f + \mathbb{1}_{(m_t>0)} \epsilon_t^f + \mathbb{1}_{(m_t=0)} \epsilon_t^0$$

Ω_t is a vector of state variables at the beginning of period t that includes health, beliefs about effective coinsurance rates, current MDE and individual characteristics. r is the risk aversion coefficient for health, which allows consumers to derive higher utility from health that carries less risk.[15] r^c is the risk aversion coefficient for the out-of-pocket cost of medical services. It allows for the possibility that consumers may prefer financial decisions with less uncertainty. θ_t is the experienced coinsurance rate in period t, and thus, $\theta_t m_t$ is the out-of-pocket cost in period t.[16] PI_t is the participation incentive in period t that the HIE paid to encourage families to join and stay in the experiment.[17] Because PI_t is measured in dollars, dividing the per-period utility by α^{PI} returns a utility that is measured in dollar units. c^f is the non-pecuniary fixed cost of using medical services, which is added to fit the prevalence of corner solutions. This is also motivated by the presence of unobserved costs of using medical services, e.g., the costs of travelling to a doctor's office. ϵ_t^f and ϵ_t^0 are the unobserved utility (to the econometrician) if using any medical services and the unobserved utility (to the econometrician) if not using any medical services, respectively, both of which follow the type 1 extreme value distributions.

Consumers who were assigned to the 25% coinsurance rate plan or the 95% coinsurance rate plan may face a non-zero MDE. I assume the distribution of experienced coinsurance rates when the MDE is reached follows the true distribution of the experienced coinsurance rates in the free plan. Consumers observe no signals in the periods when they reach their MDEs because the corresponding reimbursements are not based on their coinsurance rates and, thus, do not provide informative signals for the effective coinsurance rates.

Since the experimental plans were only available for either 3 or 5 years, values when the experiment ends should also be captured. Denote the end period of the experiment as T. The value function when the experiment ends is as follows:

$$V_{T+1}^O(\Omega_{T+1}) = \alpha_1^o \frac{H_{T+1}^{1-r^o} - 1}{1-r^o} + \alpha^{PI} PI_{T+1}$$

The value function when the experiment ends consists of two parts. The first one is the utility of health when the experiment ends. r_o is the risk aversion coefficient for health when the experiment ends, which may differ from r. The second one is the value of the Participation Incentive. The experiment offers an incentive payment to encourage people to

[15]The utility of health follows a CRRA specification because when $log(H_t)$ follows a normal distribution, the expected value of health has a closed-form expression, which is computationally convenient. A CRRA specification also yields a wider range for the utility. Because health levels are set to be positive, a CRRA utility of health spans from $-\infty$ to 0, whereas a CARA utility of health, i.e., $-\exp(-rH_t)$, is confined to the interval from -1 to 0.

[16]The utility of out-of-pocket cost follows a CARA specification because when beliefs about effective coinsurance rates follow normal distributions, the expected value of the utility has a closed-form expression.

[17]The PI_t was paid every four weeks to reduce the incentive to leave in the middle of a contract year. The PI could vary by year. However, the HIE only provides the data for the PI in the first year. I assume that PI_t does not vary over the years for simplicity.

finish the experiment. Therefore, the Participation Incentive when the experiment ends is different from the Participation Incentive during the experiment.[18]

While attrition is rare in the data, i.e., only 4.60% of consumers choose to attrit during the experiment, it may help us understand learning. For example, as time goes on and people get more signals, they may be more certain about their effective coinsurance rates and, thus, are less likely to attrit. Therefore, the values of attrition are also modelled in this paper. The value function of attrition is assumed to be different from the value function when the experiment ends because people who choose to attrit may have different outside options compared to people who choose to stay. The value function of attrition is as follows:

$$V_t^A(\Omega_t) = \alpha_1^A + \alpha_H^A H_t + \alpha_{age}^A age_t + \alpha_{age^2}^A age_t^2 + \alpha_{hs}^A high\_school + \epsilon_t^a$$

where ϵ_t^a follows the type 1 extreme value distribution. The values of attrition are influenced by current health H_t, current age age_t and an indicator variable *high_school* that equals one if this person completes at least 12 years of education and equals zero otherwise. age_t and *high_school* are included in the values of attrition because a simple probit regression of the attrition rates on the age at enrollment and the high school diploma shows that these two variables are important determinants of attrition.[19]

3.3 Learning

The learning framework is a standard Bayesian updating process. I assume there are K types of consumers who only differ in their initial beliefs about the effective coinsurance rate. Denote θ^p as the effective coinsurance rate for plan p. For Type k consumers with plan p, their initial beliefs about θ^p follow a normal distribution:

$$\theta^p \sim N(\mu_1^{k,p}, (\sigma_1^p)^2).$$

In other words, Type k consumers in period 1 believe that θ^p follows a normal distribution with the mean equal to $\mu_1^{k,p}$ and the standard deviation equal to σ_1^p.[20]

I also assume that all types of consumers have the same beliefs about the experienced coinsurance rates, s_t^p. Consumers know that s_t^p are noisy signals of the effective coinsurance rate and follow a normal distribution with the mean equal to the true mean θ^p:

$$s_t^p \sim N(\theta^p, (\sigma_s^p)^2).$$

If a consumer chooses a positive amount of spending in period t and their MDE is not reached, their belief about θ^p in the next period is summarized as follows:

$$\mu_{t+1}^p = \frac{(\sigma_t^p)^2}{(\sigma_t^p)^2 + (\sigma_s^p)^2} s_t^p + \frac{(\sigma_s^p)^2}{(\sigma_t^p)^2 + (\sigma_s^p)^2} \mu_t^p$$

[18]The details are in Appendix A.

[19]The regression on years of education does not show significant correlations between education and attrition. Therefore, the values of attrition are influenced by the high school diploma rather than the years of education.

[20]Having beliefs to be normal means that the experienced coinsurance rates sometimes may fall below zero or go above one, which seems to be problematic. However, it is unclear whether using other distributions could approximate the distributions of beliefs any better. If the model fits the observed spending and health well, then using normal distributions may not be a bad idea.

$$(\sigma^p_{t+1})^2 = \frac{1}{1/(\sigma^p_t)^2 + 1/(\sigma^p_s)^2}.$$

3.4 The optimization problem

Denote the value of spending m_t on medical services given the state Ω_t as $V^m_t(\Omega_t, m_t)$. Denote A_t as the attrition choice in period t. $A_t = 1$ means that the consumer chooses to attrit, and $A_t = 0$ means that the consumer chooses to stay, i.e., they choose whether to spend and how much to spend. If the experiment does not end immediately, a consumer in the next period will choose whether to leave the experiment, whether to spend and how much to spend to maximize their expected discounted lifetime utility. Because this consumer is uncertain about which experienced coinsurance rate they would receive, their utility in the current period is integrated over the distribution of signals. $V^m_t(\Omega_t, m_t)$ when $t < T$ is written as:

$$\begin{aligned} V^m_t(\Omega_t, m_t) =& E_t U_t(\Omega_t, m_t) + \beta E_t[Pr(A_{t+1} = 1)V^A_{t+1}(\Omega_{t+1}) \\ &+ Pr(m_{t+1} = 0)V^m_{t+1}(\Omega_{t+1}, 0) + Pr(m_{t+1} > 0)V^m_{t+1}(\Omega_{t+1}, m^*_{t+1})] \end{aligned}$$

where $m^*_{t+1} = argmax_{m_{t+1}\in(0,+\infty)}[V^m_{t+1}(\Omega_{t+1}, m_{t+1})]$ and β is the discount factor.

When the experiment ends in the next period, i.e., $t = T$, the consumer has no choice but to leave the experiment. The value function of spending m_t is:

$$V^m_T(\Omega_T, m_T) = E_T U_T(\Omega_T, m_T) + \beta E_T V^O_{T+1}(\Omega_{T+1})$$

At the beginning of each period, the consumer decides whether to stay, whether to spend and how much to spend. Thus, the maximum expected discounted lifetime utility $V_t(\Omega_t)$ is:

$$V_t(\Omega_t) = \max[V^A_t(\Omega_t), V^m_t(\Omega_t, 0), V^m_t(\Omega_t, m^*_t)]$$

4 Empirical Implementation

4.1 Measurement Equation

To employ rich health information from the HIE and also allow for measurement errors in health, I use factor analysis to estimate a distribution of latent health. Specifically, I assume individual's health at the beginning of the HIE is distributed normally based on their characteristics:

$$log(H_1) = \lambda^H X_1 + \eta^H,$$

where $\eta_H \sim N(0, \sigma_{\eta^H})$. X_1 is a vector of individual characteristics that contains the log of education years and the log of age. It is important to allow initial health to depend on education levels and age because they may be important determinants of health.[21] It also allows the distribution of initial health to vary across individuals.

[21] Other individual characteristics, i.e., income, race and gender, may also matter. But when I include them in the measurement equation, their coefficients are not significant. The parameters for the measurement equations are reported in Table 14 of the Appendix.

The two measurements used are the number of conditions from screening examinations and the number of disorder concerns, both of which are measured by non-negative integers. The measurement system is as follows:

$$z_t^{H,j} = \lambda_{z,1}^j + \lambda_{z,2}^j \log(H_t) + \eta_t^j$$

where $z_t^{H,j}$ represents the jth measurement in period t. $\eta_t^j \sim N(0, \sigma_{\eta^j})$, which is the measurement error in the jth measurement. $\lambda_{z,2}^2$, the loading for the number of disorder concerns, is normalized to -1 for identification. This normalization also means that having a higher value of H_t denotes better health.

I allow that observed spending z_t^m, i.e., recorded data for spending, to be a noisy measure of latent spending, m_t. For example, administrative mistakes may occur so that the spending I observe may not be what consumers chose, e.g., incorrect inputting of amounts or dates. The measurement equation for spending is as follows:

$$z_t^m = \begin{cases} \exp(\lambda_{m,1} + \log(m_t) + \eta_t^m) & \text{if } \ m_t > 0 \\ 0 & \text{if } \ m_t = 0 \end{cases}$$

where $\eta_t^m \sim N(0, \sigma_{\eta^m})$, which measures the degree of measurement errors in observed spending. $\lambda_{m,1}$ is set to be $-\sigma_{\eta^m}^2/2$ to ensure that $E(z_t^m|m_t > 0) = E(m_t|m_t > 0)$. This measurement system assumes no measurement errors when medical spending is zero.

Consumers do not directly observe their experienced coinsurance rates directly. Instead, they see their reimbursements. Since reimbursements may also contain measurement errors, I set another measurement equation for reimbursements. Denote the true reimbursement in period t as rb_t, and the observed reimbursement in period t as z_t^r, i.e., the recorded data for the reimbursement in period t. The relation between the latent reimbursements and observed reimbursements is as follows:

$$log(z_t^r + 1) = \lambda_{r,1} + log(rb_t + 1) + \eta_t^r$$

where

$$rb_t = \begin{cases} (1 - s_t^p)m_t \quad \text{if } \ (m_t > 0 \ \ \& \ \ MDE_t - s_t^p m_t > 0) \ \text{ or } \ p\text{=free} \\ (1 - s_t^p)\frac{MDE_t}{s_t^p} + (m_t - \frac{MDE_t}{s_t^p})(1 - s_t^{free}) \quad \text{if } \ m_t > 0 \ \ \& \ \ MDE_t - s_t^p m_t <= 0 \\ 0 \quad \text{if } \ m_t = 0 \end{cases}$$

where $\eta_t^r \sim N(0, \sigma_{\eta^r})$. I fix $\lambda_{r,1} = -(\sigma_{\eta^r})^2/2$ to ensure that, when spending is positive, the mean of true reimbursements equals the mean of observed reimbursements, i.e., $E(rb_t|m_t > 0) = E(z_t^r|m_t > 0)$.[22] The latent reimbursement is set to be $log(rb_t + 1)$ since reimbursement could be zero.

4.2 Likelihood

Let the density of some random variable x be denoted as $f_x(.)$. The likelihood for the jth measurement of health is as follows:

[22]The detailed proof is as follows: $E(z_t^r|m_t > 0) = E\{\exp(\lambda_{r,1})\exp(\eta_t^r)(rb_t + 1) - 1|m_t > 0\}$. Since $E\exp(\eta_t^r) = \exp(\sigma_{\eta^r}^2/2)$ and $\lambda_{r,1} = -(\sigma_{\eta^r})^2/2$, $E(z_t^r|m_t > 0) = E[rb_t|m_t > 0]$.

$$L_t^{H,j} = \begin{cases} f_{\eta^j+\lambda_{z,2}^j\eta^H}[z_t^j - \lambda_{z,1}^j - (\lambda_{z,2}^j\lambda^H X_1)] & \text{if } t=1 \\ f_{\eta_j}[z_t^j - \lambda_{z,1}^j - \lambda_{z,2}^j \log(H_t)] & \text{if } t=T+1 \end{cases}$$

where $f_{\eta^j+\lambda_{z,2}^j\eta^H}(.)$ represents the density function of a normal distribution that equals the sum of the distributions of η^j and $\lambda_{z,2}^j\eta^H$. Period T is the end period of the experiment. Therefore, $z_{T+1}^{H,j}$ represents the jth health measurements right after leaving the experiment. The parameters for the measurement equations of health are estimated before the estimation of other parameters in the model.

After the estimates for the measurement equation of health are derived, I use MLE to estimate the parameters for the utility, the health production function, the measurement equations on spending and reimbursements, and the parameters characterizing learning. The observations are attrition choices, A_t, observed spending, z_t^m, reimbursement amounts, z_t^r, and health measurements at exit $z_{T+1}^{H,j}$ for those who have not attritted. Denote $L_t^C(A_t, z_t^m, z_t^r|\Omega_t^k)$ as the likelihood of observing $\{A_t, z_t^m, z_t^r\}$ for Type k consumers given Ω_t^k in period t. The likelihood of observing all history of $\{A_t, z_t^m, z_t^r\}$ and $z_{T+1}^{H,j}$ for Type k consumers is as follows:

$$\begin{aligned} L(\{A_t, z_t^m, z_t^r\}_{t=1}^T, \{z_{T+1}^j\}_{j=1}^2|Type=k) = \int_{H_1}\int_{\{\epsilon_{t+1}^H, s_t\}_{t=1}^T} & \underbrace{[\prod_{t=1}^T L_t^C(A_t, z_t^m, z_t^r|\Omega_t^k\}]}_{\text{likelihood of observing } \{A_t, z_t^m, z_t^r\}_{t=1}^T} \\ & \underbrace{[\prod_{j=1}^2 L_{T+1}^{H,j}(z_{T+1}^{H,j}|\Omega_T^k)]}_{\text{likelihood of observing } \{z_{T+1}^{H,j}\}_{j=1}^2} \\ & f_{H_1}(H_1)\prod_{t=1}^T[f_{\epsilon_{t+1}^H}(\epsilon_{t+1}^H)f_{s_t}(s_t)]d\{\epsilon_{t+1}^H, s_t\}_{t=1}^T dH_1 \end{aligned}$$

where

$$\begin{aligned} L_t^C(A_t, z_t^m, z_t^r|\Omega_t^k) =& Pr(A_t=1)\mathbb{1}_{(A_t=1)} + Pr(m_t=0)\mathbb{1}_{(z_t^m=0)} \\ &+ Pr(m_t>0)\mathbb{1}_{(z_t^m>0)}\underbrace{f_{\eta_t^m}[\log(z_t^m) - \lambda_{m,1} - \log(m_t)]}_{\text{likelihood for observed spending}} \\ & \underbrace{f_{\eta_t^r}\{log(z_t^r+1) - \lambda_{r,1} - log(rb_t+1)\}}_{\text{likelihood for observed reimbursements}}, \end{aligned}$$

As there are K types of consumers, the final likelihood of observing all the history of $\{A_t, z_t^m, z_t^r\}$ and health measurements at exit would be as follows:

$$L(\{A_t, z_t^m, z_t^r\}_{t=1}^T, \{z_{T+1}^j\}_{j=1}^2) = \sum_{k=1}^K Pr(Type=k)L(\{A_t, z_t^m, z_t^r\}_{t=1}^T, \{z_{T+1}^j\}_{j=1}^2|Type=k)$$

Since the above likelihood needs to be integrated over the distribution of health and signals for at least 39 periods for individuals without attrition, calculating it would be very difficult.

I use the simulated maximum likelihood method to approximate the likelihood. Specifically, I use the initial distribution of health estimated from the first stage and draw N times H_1 as well as $N \times T$ times health shocks and signals. Then I approximate the likelihood using the following equation:

$$\frac{1}{N}\sum_{i=1}^{N}\sum_{k=1}^{K} Pr(Type = k)L(\{A_t, m_t, z_t^r\}_{t=1}^T, \{z_{T+1}^j\}_{j=1}^2 | \Omega_{i,1}^k, \{\epsilon_{i,l+1}^H, s_{i,l}\}_{l=1}^T)$$

where $\Omega_{i,1}^k$ contains the ith draw of initial health. $\{\epsilon_{i,l+1}^H, s_{i,l}\}_{l=1}^T$ are the ith draws for health shocks and signals from period 1 to the end of the experiment.

4.3 Identification

4.3.1 Identification of the Parameters that Characterize Learning

In this subsection, I briefly explain how the learning model can be identified. The identification of the parameters that characterize learning relies on the assumption that spending and reimbursements are noisily observed by the econometrician such that signals (or experienced coinsurance rates) are also noisily observed. Although the observed signals may be a noisy measurement, they still provide valuable information regarding the true signals. For example, the means of signals equal one minus the proportion of the average observed reimbursements to the average observed spending, which identifies the means of signals.[23] Furthermore, observing a high signal in the data indicates the consumer is more likely to receive a high signal, which also helps to identify the standard deviations of signals. Intuitively, although the variations in observed signals are also influenced by the distributions of the measurement errors in medical spending and reimbursements, only the distributions of latent signals affect consumers' healthcare decisions and, therefore, are identified. For example, a lower standard deviation of signals indicates that this risk-aversed consumer faces less uncertainty about their out-of-pocket payments and, therefore, is more willing to spend.

Then how consumers respond to their observed noisy signals helps to identify the means and the standard deviations of the initial beliefs. Specifically, the means and the standard deviations respond differently to noisy signals, which allows them to be separately identified. For example, in the first period, a consumer who is observed to receive a high (noisy) signal will be more likely to update the mean of their belief upwards and thereby tend to spend less. The standard deviation of their belief will also shrink because they receive additional information and become more certain about their belief. Another identical consumer in the first period who is observed to receive a low (noisy) signal and spend the same amount of medical spending will be more likely to update the mean of their belief downwards and tend to spend more. The standard deviation of their belief will shrink the same amount as

[23] Take the free plan as an example. The following proof is conditional on the plan being the free plan. Since $s_t^{free} \sim N(\theta^{free}, (\sigma_s^{free})^2)$, we have $E(1 - s_t^{free}|m_t > 0) = E(1 - s_t^{free}|m_t \ \& \ m_t > 0)$. Then $E(rb_t|m_t > 0) = E[(1 - s_t^{free})m_t|m_t > 0] = E(m_t E(1 - s_t^{free}|m_t \ \& \ m_t > 0)|m_t > 0) = E(m_t E(1 - s_t^{free}|m_t > 0)|m_t > 0) = E(1 - s_t^{free}|m_t > 0)E(m_t|m_t > 0)$. Because $E(z_t^r|m_t > 0) = E[rb_t|m_t > 0]$, $E(z_t^m|m_t > 0) = E[m_t|m_t > 0]$, we have the true mean of the free plan $E(s_t^{free}|m_t > 0) = 1 - \frac{E(z_t^r|z_t^m>0)}{E(z_t^m|z_t^m>0)}$.

the previous consumer because both consumers receive only one signal and are equivalently certain about their beliefs. Without health shocks, these two consumers only differ in their means of beliefs and, therefore, how they spend differently helps to identify the mean of the initial beliefs. More generally, the means of initial beliefs are identified by how consumers change their spending based on different observed signals. Then the standard deviations of the initial beliefs can be identified by how sensitive consumers change their spending. If consumers significantly change their spending after observing some signals, it may indicate that consumers have a high degree of uncertainty about their effective coinsurance rates, i.e., the standard deviations of the beliefs are relatively big, and thus put a high weight on their signals.

4.3.2 Identification of the Preference Parameters

After the parameters characterizing learning are identified, the remaining parameters in the utility functions can be identified by the variation in medical spending. r^o and α_1^o are the two utility parameters when the experiment ends that can be identified by consumers' amounts of medical spending in the last period, when r, the risk aversion coefficient for health during the experiment, does not enter the value function. Then r can be identified by the variation in medical spending before the last period. c^f, the fixed cost, is identified by the probability of zero spending. All the parameters in the values of attrition are identified by both the variations of attrition choices and medical spending. For example, attrition values may affect consumers' future expected discounted values and, thus, affect their decisions on medical spending.

4.3.3 Identification of the Measurement Equations and the Health Production Function

According to Cunha and Heckman (2008), the measurement equations with only one latent factor can be identified under the following two conditions: there are at least two measurements, and one has its factor loading normalized and thereby set to be one. I have two health measurements for one latent health, and the loading of the number of disorder concerns is set to be -1. Thus the measurement equations of health are identified. The measurement equations for spending and reimbursements are identified because the loadings on the two measurements are set to be one, and I also fix the constant term to ensure that the means of the observed measurements equal the means of the latent ones. Therefore, there is only one parameter, the standard deviation of the corresponding measurement error, to be estimated in each measurement equation. How true medical spending and reimbursements predicted by the model deviate from the observed measurements helps to identify the standard deviations of the measurement errors. Then, the health production function can be identified because all variables in the function, i.e., health at enrollment and exit and medical inputs (spending), are observed with measurement errors and the magnitudes of measurement errors can be identified.

5 Estimation Results

Table 5 presents the parameter estimates of signals. θ^{free}, the estimated true effective coinsurance rates for the free plan, is 2.64%. This suggests that almost all medical spending was reimbursed if a consumer was assigned to the free plan. θ^{25} and θ^{95}, the estimated true effective coinsurance rates for the 25% coinsurance rate plan and the 95% coinsurance rate, are 28.24% and 88.49%, respectively. The estimated true effective coinsurance rates are very close but not precisely equal to the coinsurance rates.

Table 5 also reports the parameter estimates of the beliefs. I end up with two types of consumers.[24] Type 2 consumers have their initial beliefs about the effective coinsurance rates that are very close to the true rates. However, for Type 1 consumers, their initial beliefs about the means significantly deviate from the true rates. In the free plan, Type 1 consumers initially believe that the mean is 34.90%, while the true rate is 2.64%. They seem to believe that, although the plan is called a "free plan", a significant amount of services are not covered. For both the 25% and the 95% coinsurance rate plans, the initial beliefs of Type 1 consumers about means lie below the true rates. Among consumers enrolled in the free plan, the probability of being Type 2 consumers is 62.78%, which suggests that most consumers can understand the free plan. For the other two plans, the probabilities of Type 1 consumers center around 45%, suggesting that nearly half of consumers can not understand plans with non-zero coinsurance rates.

The weights of signals when consumers update for the first time are reported in Table 5. In all three plans, the weights are below 2%, which indicates that consumers do not trust their signals and learn very slowly. Figure 1 displays the population means of beliefs over periods. Periods over 39 are excluded as most consumers, i.e., those who were assigned to the 3-year plans, had exited the experiment. For Type 2 consumers, the means of their beliefs do not change much because they mostly observe signals that are close to their beliefs. For Type 1 consumers, their beliefs slightly change during the experiment. Although they put small weights on signals, constantly observing signals that deviate from their beliefs still gives them opportunities to learn. However, because the weights on the signals are too small and they only receive signals every 2-3 periods, their beliefs show limited improvement during the experiment.

Table 6 presents the estimates for the remaining parameters. The coefficient of the Participation Incentive α^{PI} is fixed at 0.01, meaning the utility will be measured in dollars if divided by 0.01.[25] For the health production function, the estimate of the constant term γ_1^H is negative, capturing the expectation that consumers' health depreciates over time. Specifically, if a consumer with a median level of health had no medical spending and received no health shock, their health in the next period would depreciate roughly by 1.66%. The coefficient of health γ_2^H is very close to 1, indicating that consumers' health does not vary much without big health shocks or large amounts of medical spending. I also find that the effect of medical spending on health is positive ($\gamma_3^H > 0$), consistent with the expectation that utilization of medical services will improve health (Cronin, 2019; Liu et al., 2023). The parameters for the measurement equations of health are reported in Table 14 of the appendix.

[24]The model fits spending patterns well. Therefore, I did not include more types.

[25]I also estimated a specification without α^{PI} being fixed. The point estimate was 0.008, but it was not precisely estimated (the standard error was 0.008), and the log-likelihood was not substantially different.

Table 5: Estimates for Beliefs and Signals

	Est.	Std. Err.
The free plan		
Mean of initial beliefs (Type 1), $\mu_1^{1,free}$	0.349	0.004
Mean of initial beliefs (Type 2), $\mu_1^{2,free}$	0.027	0.003
Prob. of Type 1, pr^{free}	0.372	0.058
Std. dev. of initial beliefs, σ_1^{free}	0.015	0.001
True mean, θ^{free}	0.026	0.002
Std. dev. of signals, σ_s^{free}	0.154	0.002
Weight on signals in the first updating	0.009	Calculated
The 25% coin. rate plan		
Mean of initial beliefs (Type 1), $\mu_1^{1,25}$	0.049	0.009
Mean of initial beliefs (Type 2), $\mu_1^{2,25}$	0.325	0.020
Prob. of Type 1, pr^{25}	0.436	0.136
Std. dev. of initial beliefs, σ_1^{25}	0.007	0.005
True mean, θ^{25}	0.282	0.008
Std. dev. of signals, σ_s^{25}	0.124	0.009
Weight on signals in the first updating	0.004	Calculated
The 95% coin. rate plan		
Mean of initial beliefs (Type 1), $\mu_1^{1,95}$	0.176	0.013
Mean of initial beliefs (Type 2), $\mu_1^{2,95}$	0.771	0.015
Prob. of Type 1, pr^{95}	0.463	0.105
Std. dev. of initial beliefs, σ_1^{95}	0.003	0.001
True mean, θ^{95}	0.885	0.015
Std. dev. of signals, σ_s^{95}	0.027	0.001
Weight on signals in the first updating	0.015	Calculated

Note: The weights on signals in the first updating are calculated using the following formula: $\frac{(\sigma_1^p)^2}{(\sigma_s^p)^2+(\sigma_1^p)^2}$. The remaining parameters for the utility, the value of attrition, the measurement equations on spending and reimbursements, and the health production function are in Table 6.

Figure 1: Beliefs Over Time

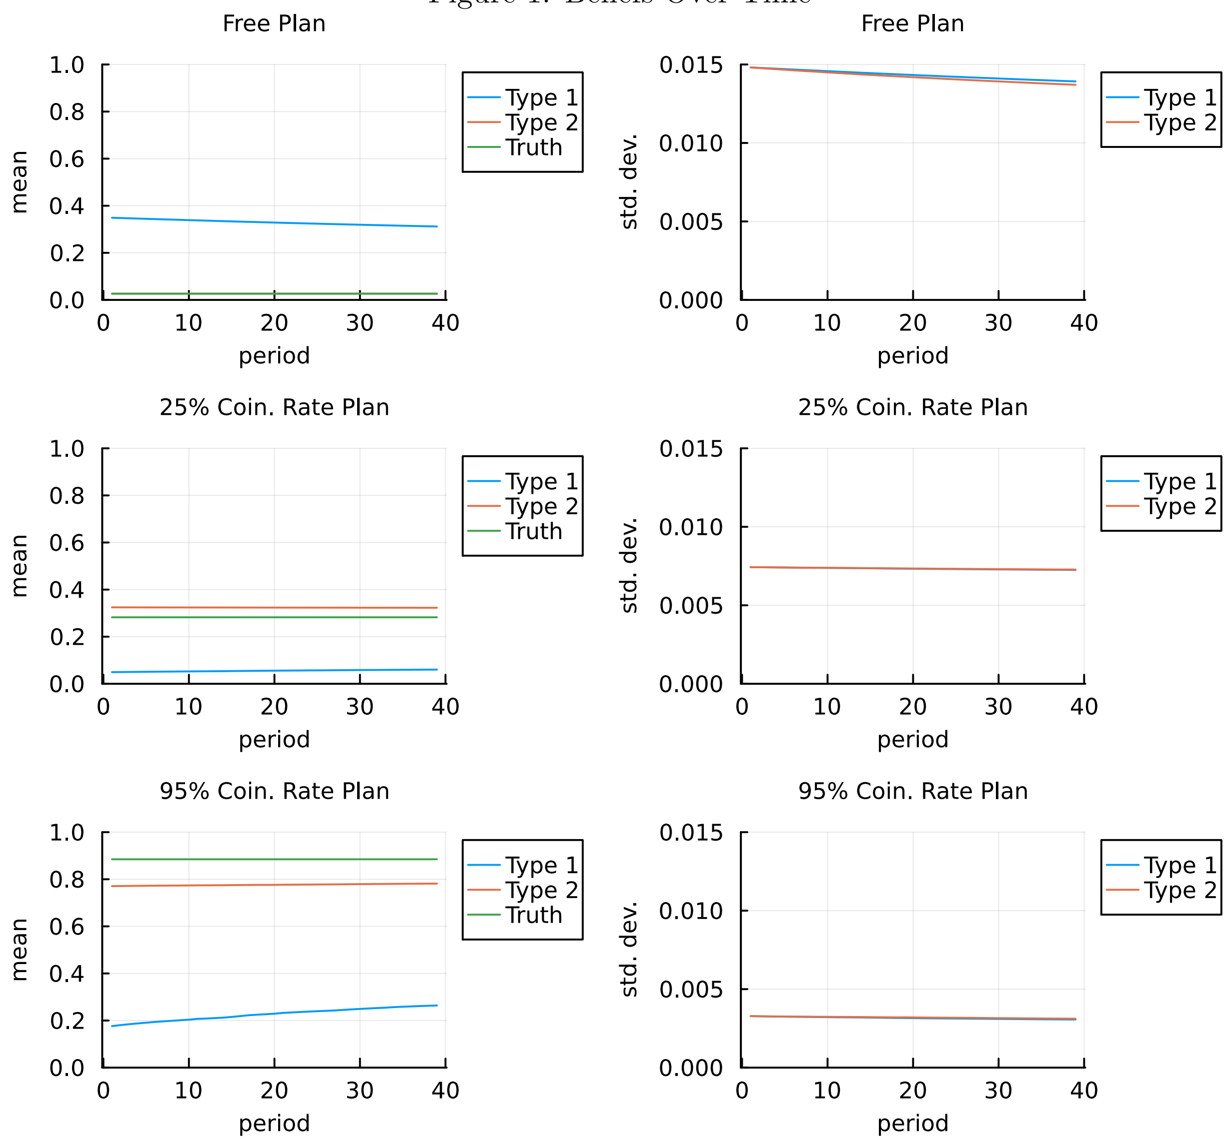

Note: The figure presents the average of the means and the standard deviations for beliefs over periods from the simulations using the derived estimates. The truth refers to the true effective coinsurance rate. Each person is simulated 500 times. The means of Type 2 Consumers in the free plan are not displayed because they overlap with the truth.

Table 6: The Remaining Estimates

	Est.	Std. Err.
Utility and value when the experiment ends		
Risk aversion coefficient of health during the experiment, r	1.994	0.003
Risk aversion coefficient of out-of-pocket payments, r^c	0.047	0.000
Coefficient of Participation Incentive, α^{PI}	0.01	Fixed
Fixed cost of positive spending, c_f	9.033	0.110
Risk aversion coefficient of health when the experiment ends, r^o	1.619	0.016
Coefficient of utility of health when the experiment ends, α_1^o	119.758	1.229
Discount factor, β	0.996	Fixed
Health production function		
Constant term, γ_1^H	-0.024	0.000
Coefficient of health, γ_2^H	0.992	0.000
Coefficient of $log(spending + 1)$, γ_3^H	0.0095	0.000
Coefficient of $log(age_1)$, γ_4^H	-0.0002	0.000
Std. dev. of the error term, σ^H	0.116	0.000
Value of attrition		
Constant term, α_1^A	-502.416	42.264
Coefficient of health, α_H^A	340.819	14.879
Coefficient of age, α_{age}^A	-9.312	0.892
Coefficient of quadratic age, $\alpha_{age^2}^A$	0.050	0.015
Coefficient of high school diploma, α_{hs}^A	-137.955	27.409
Measurement errors		
Std. dev. of measurement errors on spending, σ_{η^m}	1.142	0.023
Std. dev. of measurement errors on reimbursement, σ_{η^r}	1.173	0.022
Log-likelihood	-12482.480	

Note: The estimates for the measurement equations are reported in Table 14 of the appendix.

Table 7: Actual and Simulated Probabilities of Zero Spending

Plan Type	Data	Simulations
free	0.501	0.553
25% coin. rate	0.620	0.608
95% coin. rate	0.780	0.673

5.1 Model Fit

To assess the fit of the model, healthcare spending, the regressions of log(spending+1) on the number of observed signals, attrition, health, and reimbursements are compared between the data and the simulations. Each consumer is simulated 500 times.

5.2 Spending

The fit of spending is extremely important. The parameters that characterize learning are identified from spending dynamics. Spending is also one of the main interests of the counterfactuals. Therefore, fitting the spending well is not only necessary for a good fit of learning patterns but also important for the accuracy of counterfactuals. Given the prevalence of corner solutions, i.e., zero spending, I separate spending into extensive (probabilities of zero spending) and intensive (expenditures when choosing to spend) margins, both of which would be informative about beliefs.

5.2.1 Pooled Spending

I start by assessing the fit of pooled spending by plan. Table 7 presents the data and the simulations for the probabilities of zero spending. In both the data and the simulations, the probability of zero spending increases with coinsurance rates, i.e., consumers in the free care plan spend more frequently than consumers in the 25% coinsurance rate plan, and the latter spend more frequently than consumers in the 95% coinsurance rate plan. Although the probabilities are very close for the other two plans, the probability of zero spending for the 95% coinsurance rate plan is 10.7% lower in the simulations. On average, the overall probability is pretty close between the simulations and the data. Regarding spending when consumers choose to spend, the distributions are extremely close between the data and the simulations, as shown in Figure 2.

5.2.2 Spending Dynamics

Learning patterns are the main interest of this paper, but because beliefs are not observed, the fit of learning patterns cannot be directly assessed. Instead, I compare the fit of spending dynamics to assess whether learning patterns fit well. Specifically, I compare the regression coefficients of log(spending+1) on the number of observed signals to assess the fit. The regressions have been discussed previously as preliminary evidence of learning. If the regressions using the simulations produce coefficients that are very similar to the ones using the data, this suggests the model is capturing learning patterns.

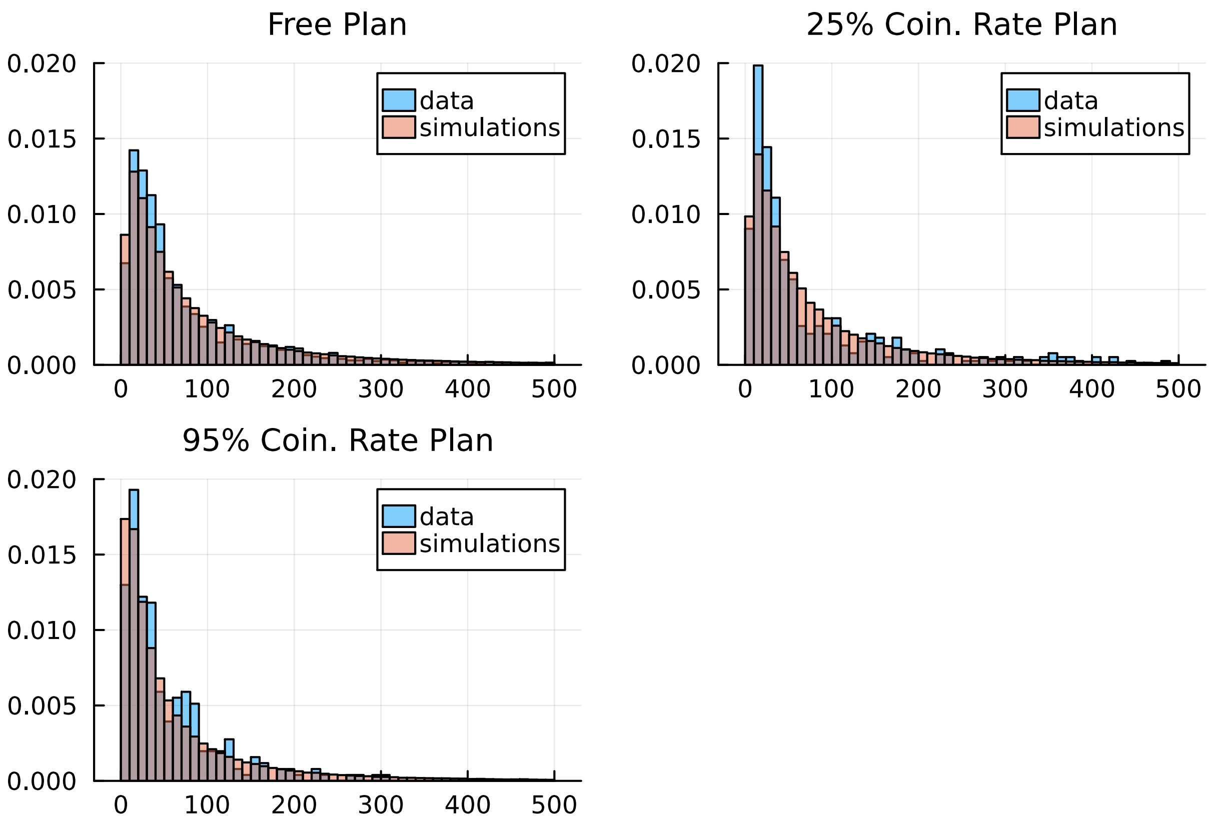

Figure 2: Histograms of Spending

Note: The above histograms are conditional on spending being positive and below 500. The histograms of spending from the data and the simulations are both presented to compare the model fit.

Table 8: Regression of log(spending+1) on the Number of Signals Observed

	Free		25% Coin. Rate		95% Coin. Rate	
	Data	Simulations	Data	Simulations	Data	Simulations
Num. of signals	0.090**	0.074	0.087**	0.062	0.101**	0.044
	(0.004)		(0.012)		(0.012)	
Num. of person-period obs.	4250		1060		1195	

Standard errors in parentheses.

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$

The coefficients on the other controls are presented in Table 12 in the appendix.

The regression results are displayed in Table 8. In the data, the coefficients on the number of observed signals are positive and significant for all three plans. This suggests that, as more signals are observed, consumers who were enrolled in these three plans were inclined to spend more. The regressions using the simulations produce coefficients that are also positive for all three plans. Although the three coefficients are smaller than the ones derived from the data, the difference is relatively small for the free plan and the 25% coinsurance rate plan. For the 95% coinsurance rate plan, the simulation produces a coefficient that is less than half of that derived from the data. This may be in part due to the higher probability of zero spending under the 95% coinsurance rate plan in the simulations. Overall, the coefficients have the same signs between the simulations and the data, and the magnitude of the coefficients does not differ too much, especially in the free plan and the 25% coinsurance rate plan, which accounts for 74% of the population. Therefore, the regression results suggest that the learning patterns fit well with the model.

5.3 Health

It is essential to fit health well because it helps to isolate spending dynamics resulting from learning. Specifically, because the parameters characterizing learning differ by plans, not only should the distributions of health be captured well for the whole sample, but they should also fit each plan separately. If the health transitions within each plan are not captured, then the spending dynamics resulting from learning may not be isolated well, which may bias the estimates that characterize learning. As shown in Figure 3, overall, the average health at enrollment and exit is very close between the simulations and the data. Table 13 in the appendix displays the average health at enrollment and exit by plan. At enrollment, the means and the standard deviations of simulated health are very close to the ones in the data, even by plan. However, the simulated health by plans differs slightly from the data at exit. While in the data, consumers from the free plan have the worst average health at exit, the simulations predict that consumers in the 95% coinsurance rate plan end the experiment with the worst health.

5.4 Attrition

While spending contains significant information about beliefs, attrition may also help us understand learning. For example, as time goes on and people get more signals, they may be more certain about their effective coinsurance rates and, thus, are less likely to attrit.

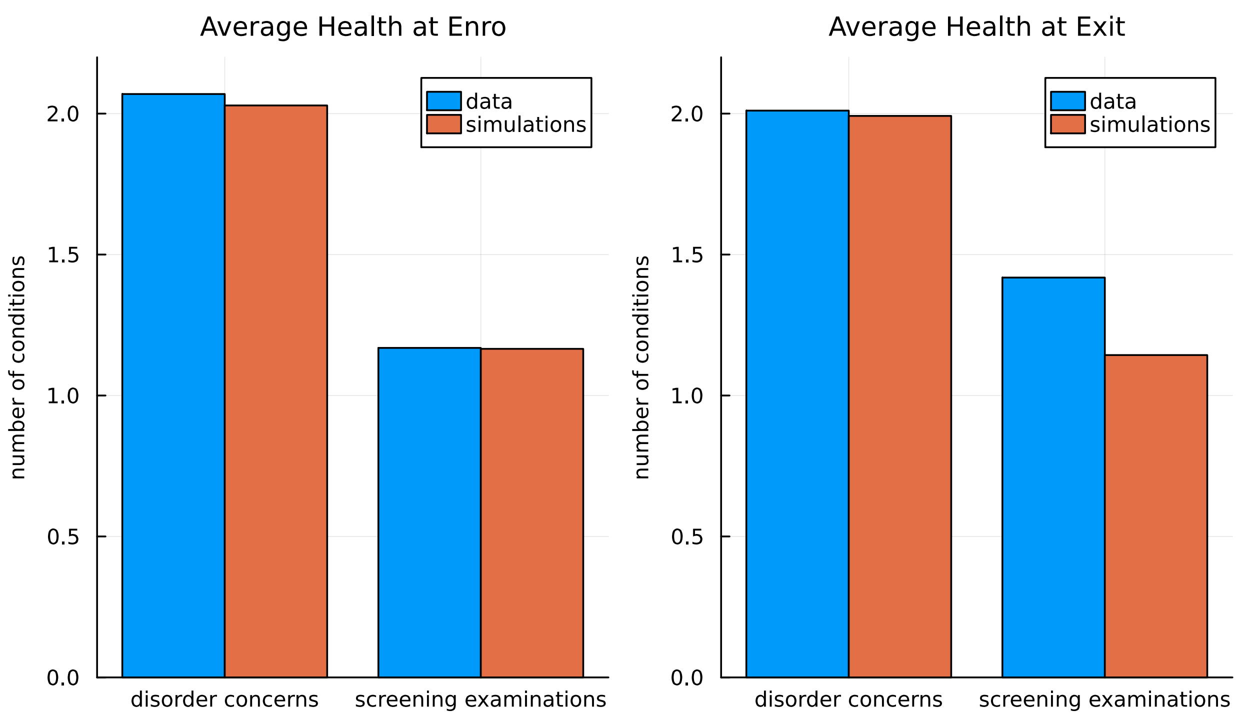

Figure 3: Average Health at Enrollment and Exit

Note: The figure displays the average disorder concerns and the average number of conditions from screening examinations in the data and the simulations.

However, as attrition is rare across all plans, it may be less informative about learning compared to spending dynamics. To assess the fit of attrition, I estimate the survival functions for both the data and the simulations and compare their predicted survival rates over time. As shown in Figure 5 of the appendix, the simulated survival rates capture the trend that attrition rates decline over the years. While the predicted attrition in the 95% coinsurance rate plan is underestimated, the model still predicts that this plan has the highest attrition rate, which is consistent with the data. Additionally, the attrition patterns in the other two plans are quite similar between the data and the simulations. Because only 24% of consumers are in the 95% coinsurance rate plan, the attrition behaviour is well captured by the model for the majority of the population.

5.5 Reimbursement Amount

Reimbursements are what people actually observe. Whether fitting reimbursements well matters for identifying the distributions of signals. From Figure 6 in the appendix, the distributions of reimbursements have a good fit with the data.

6 Counterfactual

In this section, I construct two counterfactuals to investigate the importance of biased beliefs and learning. Specifically, I compare spending, health, attrition and discounted experienced utility at enrollment in different scenarios.[26] The base case uses the derived estimates. In order to better understand how consumers are affected by different plans, the outcomes in the base case are simulated by plan, assuming that each plan has the same set of consumers, i.e., nearly all consumers in my sample. Because consumers in the free plan do not have positive MDEs, I follow the experimental design that randomly assigns consumers a MDE that is 5%, 10% or 15% of their income with a maximum of $1000. Three individuals are excluded because they had no income in the previous year but were assigned to the free plan such that their assigned MDEs would be zero. I also set the Super Participation Incentive to be zero for everyone because the payment was only available for the 95% coinsurance rate plan. Individuals whose data are only included partially, e.g., those who were terminated during the experiment and thus the periods that are one year before the terminations are not included, are allowed to stay in the experiment until the end of the experiment. Each individual is simulated 500 times.

6.1 Counterfactual 1: Consumers with Full Information

The first counterfactual assumes all consumers have full information (*informed consumers*), i.e., unbiased beliefs and no uncertainty about their effective coinsurance rates. Comparing the base case with the counterfactual under full information helps quantify the importance of biased beliefs and learning about effective coinsurance rates. Informed consumers and Type

[26] Experienced utility refers to the actual utility that consumers experience when their effective coinsurance rates are realized. The ex-ante utility may not be appropriate to measure consumers' welfare because their beliefs may be biased, and their ex-ante values may significantly deviate from what they actually experience.

2 consumers are expected to behave similarly and, thus, end the experiment with similar levels of health. However, the true effective coinsurance rates are quite different from the beliefs of Type 1 consumers. As a result, informed consumers should spend a very different amount on medical expenditures. Specifically, in the free plan, informed consumers should spend more and be less likely to attrit than Type 1 consumers. This is because the true effective coinsurance rate is much lower than the initial mean of Type 1 consumers. Informed consumers should also end the experiment with better health due to their higher spending on health. In the other two plans, informed consumers should spend less and be more likely to attrit compared to Type 1 consumers because the true effective coinsurance rates lie above the initial priors of Type 1 consumers. They should end the experiment with worse health due to their lower medical spending. In all three plans, consumers' discounted experienced utility at enrollment should be higher in the full information case because consumers make their decisions under the correct information sets.

Table 9 shows the average spending in both the base case and the counterfactual with full information. The differences in average spending between informed consumers and Type 2 consumers are small. This is because the true effective coinsurance rates are very close to Type 2 consumers' beliefs. However, informed consumers spend very differently from Type 1 consumers. They, on average, spend \$41.37 (192.96%) more in the free plan, \$34.22 (53.06%) less in the 25% coinsurance rate plan, and \$30.82 (69.45%) less in the 95% coinsurance rate plan on a four-week basis. Because Type 1 consumers account for almost half of the population, informed consumers, on average, still spend somewhat differently from consumers in the base case, i.e., \$15.25 (32.06%) more in the free plan, \$13.08 (30.17%) less in the 25% coinsurance rate plan, and \$15.29 (53.00%) less in the 95% coinsurance rate plan.

Understanding how health is affected by biased beliefs and learning is also essential. Table 10 shows the average health at exit in the two cases.[27] Since the average spending is very close between informed consumers and Type 2 consumers, they end the experiment with similar levels of health. Compared with Type 1 consumers, informed consumers, on average, have 25.00% better health in the free plan, 14.29% worse health in the 25% coinsurance rate plan, and 22.22% worse health in the 95% coinsurance rate plan. Table 15 in the appendix compares the average health at exit measured by the two health measurements. Compared to Type 1 consumers, informed consumers, on average, would have 0.36 fewer conditions from screening examinations and 0.26 fewer disorder concerns in the free plan, 0.26 more conditions from screening examinations and 0.17 more disorder concerns in the 25% coinsurance rate plan, and 0.43 more conditions from screening examinations and 0.28 more disorder concerns in the 95% coinsurance rate plan. Therefore, consumers' health could be significantly affected by biased beliefs about plan characteristics.

Because attrition is rare, it does not change much when consumers are fully informed. Table 16 in the appendix presents the attrition rates by plan. The attrition rates are very close between informed consumers and Type 2 consumers. Although the attrition rates are small in all three plans, informed consumers do attrit slightly differently compared with Type 1 consumers: 1.81% less informed consumers choose to attrit in the free plan, 1.06% more choose to attrit in the 25% coinsurance rate plan, and 3.41% more choose to attrit in

[27] Health for those who left the experiment earlier are excluded. This exclusion does not affect the results because attrition rates are rare and do not differ too much by case.

Table 9: Average Spending in the Base Simulations and the Counterfactuals

	Base Case			Full Info.	More Frequent Signals		
Plan Type	All	Type 1	Type 2	All	All	Type 1	Type 2
free	47.56	21.44	62.74	62.81	48.29	23.40	62.77
25% coin. rate	43.35	64.49	26.91	30.27	42.55	62.52	27.04
95% coin. rate	28.85	44.38	15.15	13.56	24.97	36.42	14.87

Note: This table presents the average amounts of spending in the base case and the counterfactuals. The full information cases assume that consumers have unbiased beliefs and no uncertainty about their beliefs. Since types only differ in their beliefs, consumers under full information do not differ by type. The counterfactual with more frequent signals assumes that consumers receive signals even when they consume no medical service.

Table 10: Health at Exit in the Base Simulations and the Counterfactuals

	Average Latent Health						
	Base Case			Full Info.	More Frequent Signals		
Plan Type	All	Type 1	Type 2	All	All	Type 1	Type 2
free	0.46	0.40	0.50	0.50	0.46	0.41	0.50
25% coin. rate	0.44	0.49	0.41	0.42	0.44	0.49	0.41
95% coin. rate	0.40	0.45	0.36	0.35	0.39	0.42	0.36

Note: This table presents the average latent health at exit in the base case and the counterfactuals.

the 95% coinsurance rate plan. Consumers who attrit due to their biased beliefs may suffer more because their healthcare decisions are made under non-optimal plans. However, as we do not have their plan characteristics after attrition, we cannot discuss this in more detail.

It is also important to evaluate how much consumer welfare has improved with access to full information. Table 11 presents the average discounted experienced utility at enrollment. In the per-period utility, the Participation Incentive PI_t is measured in dollars. Therefore, the experienced utility can be measured in dollar units if divided by α^{PI}. Because the effective coinsurance rates are very close to Type 2 consumers' beliefs, the average discounted experienced utility at enrollment is very close between informed consumers and Type 2 consumers. However, due to much more accurate information sets, informed consumers do make better decisions and receive much higher utility than Type 1 consumers in all three plans. The difference in discounted experienced utility at enrollment increases with the coinsurance rates: informed consumers, on average, receive \$2465.6 more utility in the free plan, \$6820.4 more utility in the 25% coinsurance rate plan and \$40341.8 more utility in the 95% coinsurance rate plan. The differences in experienced utility show that the value of information is extremely important, especially for those consumers with significantly biased beliefs.

Table 11: Average Discounted Experienced Utility at Enrollment in the Base Simulations and the Counterfactuals

	Base Case			Full Info.	More Frequent Signals		
Plan Type	All	Type 1	Type 2	All	All	Type 1	Type 2
free	-37631.9	-39171.7	-36719.9	-36706.1	-37549.4	-38950.8	-36719.2
25% coin. rate	-43490.8	-47321.3	-40553.7	-40500.9	-43185.2	-46620.2	-40551.4
95% coin. rate	-62341.3	-83973.3	-43765.2	-43631.5	-55110.6	-68340.9	-43749.3

Note: Experienced utility refers to the actual utility that consumers experience when their effective coinsurance rates are realized. The utility has been divided by 0.01, the coefficient on the PI_t in the per-period utility, so that the values are measured in dollars.

6.2 Counterfactual 2: More Frequent Signals

The speed of learning depends on many factors. Although consumers learn very slowly during the experiment, there might be ways to increase their learning speed. One approach is to send signals to consumers even when they consume no medical service.[28] In the data, consumers spend every 2-3 periods and, therefore, they do not observe signals in every period. Sending additional signals to consumers helps them update their beliefs more frequently, which increases their speed of learning.

Figure 4 presents the average differences between the means of beliefs and the true effective coinsurance rates for Type 1 consumers in the base case and the case receiving signals in every period. The average differences for Type 2 consumers are not displayed because their beliefs are fairly close to the true rates. Although consumers' beliefs do not change much in the 25% coinsurance rate plan, consumers receiving frequent signals tend to update more quickly in the free plan and the 95% coinsurance rate plan. Type 1 consumers in the 95% coinsurance rate plan receive the highest gain in information due to their relatively higher rate of learning and more frequent learning opportunities. After 39 periods, the average difference between their means of beliefs and the true rate is -62.39% in the base case, but it increases to -44.89% with more frequent signals. Although their beliefs still differ significantly from the true effective coinsurance rate, Type 1 consumers in the 95% coinsurance rate plan still receive a moderate improvement in their information, which substantially improves their healthcare decision-making.

Table 9 presents the average spending for consumers receiving more frequent signals. In all three plans, consumers with more frequent signals, regardless of which type they are, make better spending decisions than consumers in the base case, i.e., the average spending is closer to those made by informed consumers. The change is the most evident for Type 1 consumers in the 95% coinsurance rate plan. By receiving more frequent signals, they,

[28] Alternatively, one could reduce the standard deviations of signals to increase the speed of learning. However, this reduction will also increase consumers' medical spending in both the base case and the full information case because consumers would have less uncertainty about their out-of-pocket payments, which reduces their utility loss of spending. The welfare loss from biased beliefs may be amplified because consumers' medical spending may deviate further from those under full information. Sending more frequent signals not only gives consumers more opportunities to learn but also fixes the amount of medical spending under full information. Therefore, learning by receiving more frequent signals reduces the welfare loss from biased beliefs.

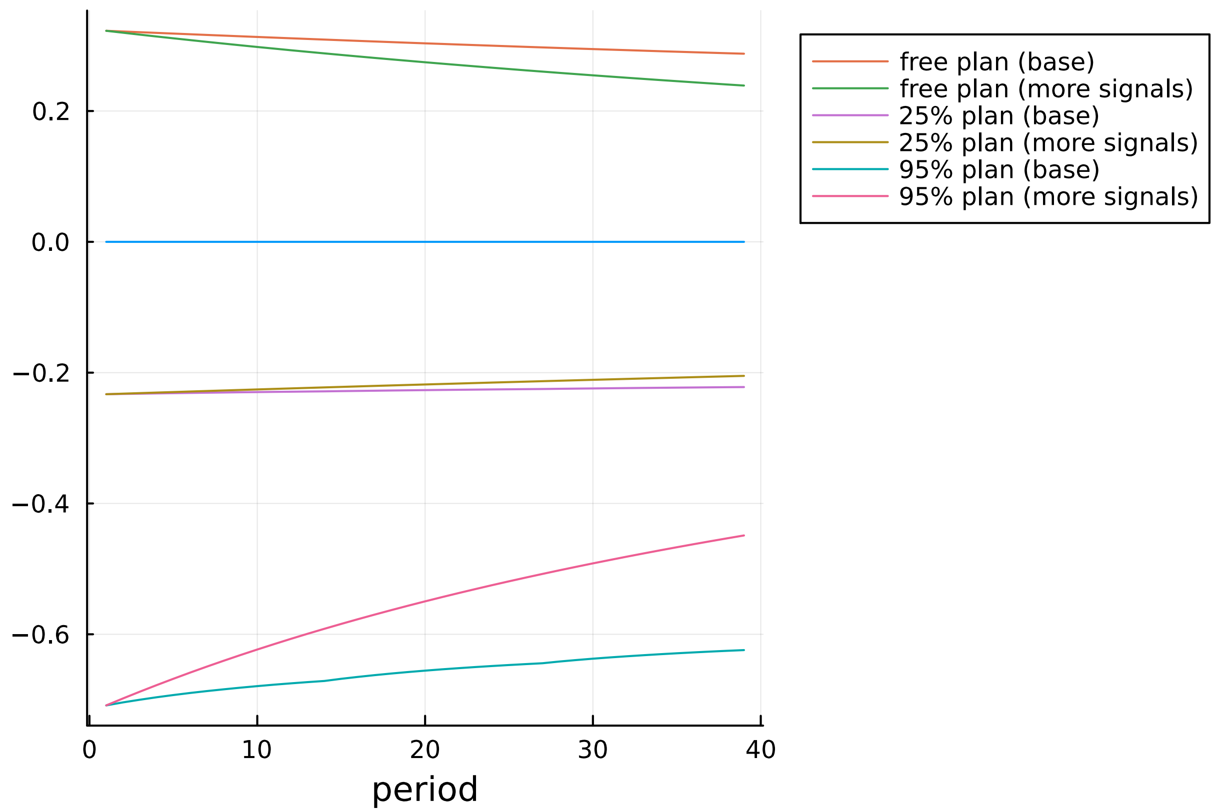

Figure 4: Differences Between the Means of Beliefs and the True Effective Coinsurance Rates for Type 1 Consumers

Note: The figure displays the average differences between the means of beliefs and the true effective coinsurance rates for Type 1 consumers in the base cases and the cases receiving signals in every period.

on average, reduce their spending by \$7.96 compared to those in the base case. Table 10 and Table 16 present the average health at exit and the attrition rates for consumers receiving more frequent signals. In all three plans, receiving more frequent signals would lead consumers to have their health and attrition rates closer to those in the full information case. Table 11 reports the average discounted experienced utility at enrollment. In all three plans, the experienced utility is higher if consumers receive more frequent signals. Type 1 consumers in the 95% coinsurance rate benefit the most: their utility, on average, increases by \$15632.4 from receiving more frequent signals. This increase indicates that consumers' welfare could be significantly improved even by a moderate gain in information.

7 Conclusion

In this paper, I measure the importance of consumers' biased beliefs and learning about the characteristics of their health insurance plans. To do so, I develop a dynamic model of consumer decision-making in healthcare, allowing for the possibility that consumers may have biased beliefs about their effective coinsurance rates and may learn about the rates via their experienced coinsurance rates. The previous literature on models of healthcare decisions typically assumes that consumers are well-informed about their plan characteristics. However, ample evidence has found that this assumption may be unrealistic. This paper extends existing models by relaxing the assumption of full information.

The results show that nearly 40% of the consumers have their initial beliefs significantly deviate from the true rates. This is consistent with previous evidence that consumers may have biased beliefs about their plan characteristics. Furthermore, these consumers learn very slowly, which is novel in the literature. A counterfactual with well-informed consumers shows that biased beliefs can significantly distort consumers' healthcare decisions and health outcomes. Another counterfactual in which consumers receive signals in every period also shows that, even with a moderate improvement in information, consumers' healthcare decision-making and welfare could be substantially improved. Therefore, how to improve consumers' understanding of their plan characteristics is important and needs to be addressed.

The paper could be further extended in several ways. First, future work could use more recent datasets. The HIE was implemented 40 years ago. Nowadays, consumers may have more tools to help them understand their plan characteristics. Therefore, the results may not apply to consumers in the current context. Second, while this paper provides some insights into how biased beliefs may affect plan choices (or attrition), I cannot investigate the problem further due to the data limitations. However, how biased beliefs influence consumers' plan options may be important because their welfare could be exacerbated further by choosing non-optimal plans. Future research could provide a more comprehensive understanding of the effects of biased beliefs, particularly regarding plan options. Third, future studies may investigate other sources of information that may help consumers learn about their plan characteristics. For example, this paper focuses on consumers who are single. However, numerous health insurance plans are offered at a family level, which may increase the value of learning and help consumers learn more quickly. Hodor (2021) has explored this question in a reduced-form framework. Future research could investigate the issue by using a structural model to provide more insights.

Appendix A Participation Incentive Payments

The Participation Incentive (PI) was designed to compensate for people who were financially worse off in the HIE compared to their existing health insurance plans. It was calculated by the maximum loss when switching from the existing plan to the plan in the HIE. If the premium exceeded the maximum difference, the PI was equal to the premium payment.

Specifically, the calculation of the PI ignored a family's true medical expenditures. If one had an existing plan with a deductible of \$100, above which they had to pay 20% coinsurance, and they were assigned to the plan with 25% coinsurance rate and \$1000 MDE in the HIE, then, the maximum loss occurred when the medical bills under the HIE were exactly \$4000. The maximum difference was thus $4000 \times 25\% - 100 - 20\% \times (4000 - 100) = 120$. This family or individual was paid \$120 per year.

The total PI could not exceed the MDE in the HIE unless its insurance premium exceeded the MDE. For example, if one family paid an insurance premium of \$900, then under the insurances above, its PI would be \$1000. But if the family paid a premium of \$1200, then the PI became \$1200.

To encourage every family to complete the experiment, a portion of the PI was withheld until the last term. In the last year, families received the full PI and the amount withheld was paid as a bonus for completing the physical screening examination and medical health questionnaire at exit. The percentage of PI withheld depended on the site and assigned enrollment term. Details of the percentages can be found in the codebooks for data.

Appendix B Full Results for Regressions on Number of Signals

Table 12: Full Regression of Number of Signals

	actual data			simulated data		
	free	25coin	95coin	free	25coin	95coin
	log_m	log_m	log_m	log_m	log_m	log_m
n_signal	0.090**	0.087**	0.101**	0.074**	0.062**	0.044**
	(0.004)	(0.012)	(0.012)	(0.000)	(0.000)	(0.000)
log_m_before	0.154**	0.126**	0.088**	0.155**	0.158**	0.166**
	(0.017)	(0.040)	(0.024)	(0.000)	(0.002)	(0.003)
period	-0.055**	-0.044**	-0.029**	-0.052 **	-0.041 **	-0.032**
	(0.002)	(0.006)	(0.003)	(0.000)	(0.000)	(0.001)
log_edu	0.410**	1.185 **	-0.585*	-0.391**	-0.196**	-1.08**
	(0.103)	(0.358)	(0.245)	(0.006)	(0.015)	(0.013)
log_age	0.484 **	0.386**	0.400 **	0.291**	0.442**	0.378**
	(0.085)	(0.190)	(0.113)	(0.004)	(0.008)	(0.006)
log_diff_MDE	0	-0.144**	-0.087**	0	-0.156 **	-0.193 **
	(.)	(0.050)	(0.029)	(.)	(0.002)	(0.001)
_cons	-1.469**	-2.619**	1.347 +	1.291 **	0.957**	3.51**
	(0.471)	(1.300)	(0.798)	(0.022)	(0.054)	(0.040)
N	4250	1060	1195	2101584	543064	662554

standard errors in parentheses

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$

Appendix C Additional Model Fit

Table 13: Average Health

Average Health at Enrollment	Disorder concerns		Screening Examinations	
	Data	Simulations	Data	Simulations
free	2.22 (1.97)	2.08 (1.91)	1.39 (1.15)	1.28 (1.13)
25% coin. rate	1.58 (1.47)	1.89 (1.88)	0.67 (0.82)	1.03 (1.06)
95% coin. rate	2.18 (1.81)	1.95 (1.89)	0.82 (0.91)	1.12 (1.14)
Overall	2.07 (1.85)	2.03 (1.91)	1.17 (1.09)	1.17 (1.08)
Average Health at Exit				
free	2.23 (2.19)	1.98 (1.85)	1.48 (1.25)	1.12 (0.91)
25% coin. rate	1.47 (1.84)	1.96 (1.84)	1.46 (1.33)	1.11 (0.93)
95% coin. rate	1.62 (1.04)	2.06 (1.86)	1.03 (1.14)	1.28 (0.97)
Overall	2.01 (2.02)	1.99 (1.84)	1.41 (1.24)	1.14 (0.93)

Note: The table presents observed and simulated average health by plans. Standard deviations are presented in parentheses.

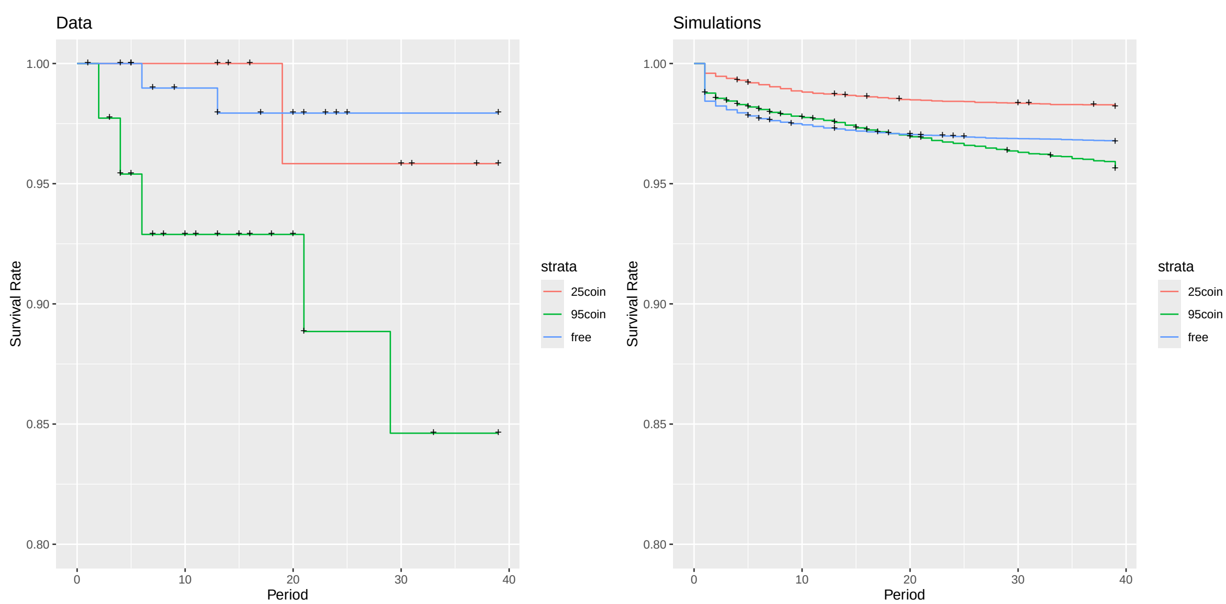

Figure 5: Survival Rate Over Periods

Note: The survival function is estimated using the Kaplan–Meier approach, a non-parametric method. The plus signs represent censored cases in the data, e.g., people who do not have more data after those periods and have not yet attrited.

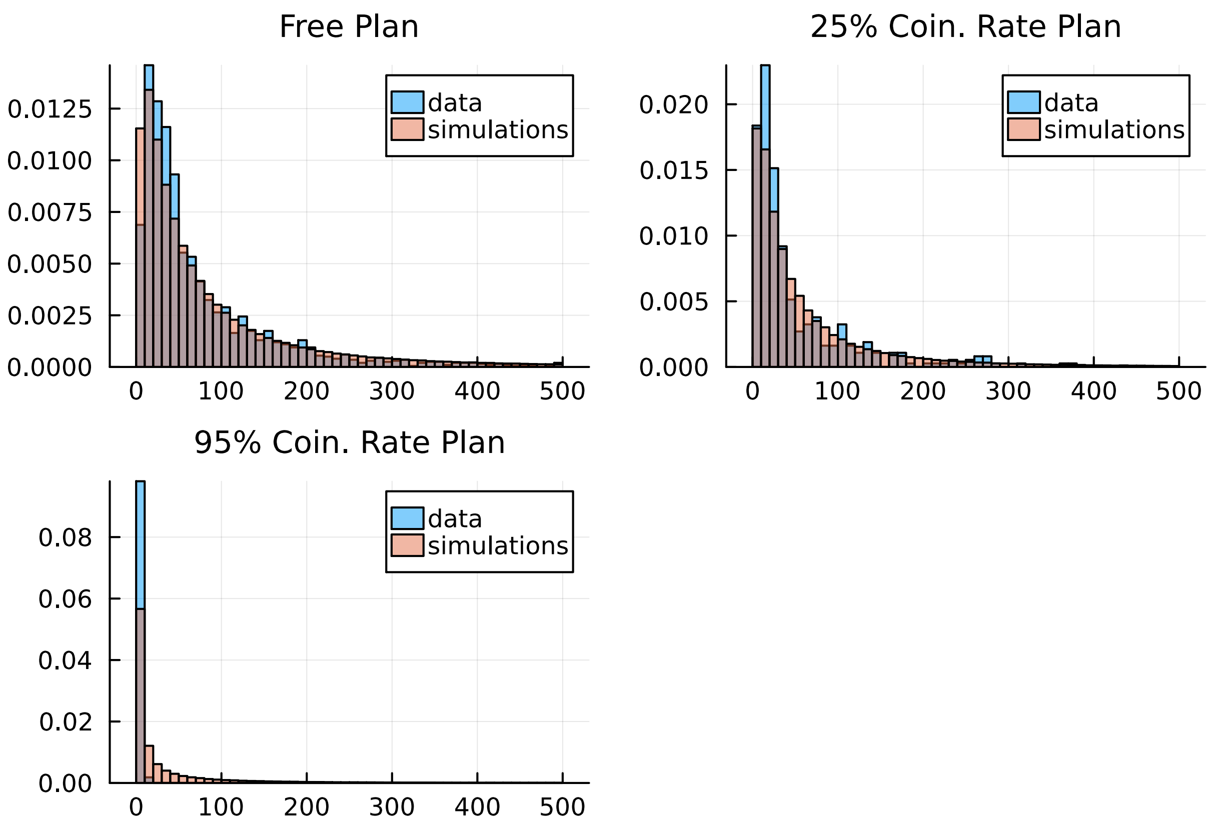

Figure 6: Histograms of Reimbursement by Plans

Note: The above histograms are conditional on reimbursements being positive. The histograms of reimbursements from the data and the simulations are both presented to compare the model fit.

Appendix D Estimates for Measurement Equations

Table 14: Estimates of Measurement Equations

health_exam		
$\log(H)$	-1.45996***	-1.45919***
	(0.45972)	(0.45246)
_cons	-0.16830	-0.20993
	(1.57418)	(1.37172)
disorder_concerns		
$\log(H)$	-1.00000	-1.00000
	(.)	(.)
_cons	1.07864	1.05104
	(1.03977)	(0.92513)
$\log(H_{exit})$		
_cons	-1.08356	-1.11255
	(1.05157)	(0.93175)
$\log(H_1)$		
male	-0.00188	
	(0.13840)	
log(income+1)	0.00226	
	(0.04953)	
white	0.02959	
	(0.41161)	
log(edu_years)	0.84449**	0.84851***
	(0.32973)	(0.31951)
log(age)	-0.85418***	-0.85189***
	(0.26049)	(0.25626)
var(e.$\log(H_1)$)	0.19262*	0.19258*
	(0.10589)	(0.10581)
var(e.$\log(H_{exit})$)	0.54952**	0.55011**
	(0.27696)	(0.27556)
var(e.health_exam)	0.36574	0.36567
	(0.31206)	(0.30454)
var(e.disorder_concerns)	3.19023***	3.19113***
	(0.32085)	(0.31910)
N	171	171

Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Appendix E Additional Tables for the Counterfactuals

Table 15: Health at Exit in the Base Simulations and the Counterfactuals

Average Number of Conditions from Examinations							
	Base Case			Full Info.	More Frequent Signals		
Plan Type	All	Type 1	Type 2	All	All	Type 1	Type 2
free	1.10	1.33	0.96	0.97	1.08	1.29	0.96
25% coin. rate	1.15	0.98	1.28	1.24	1.16	1.00	1.28
95% coin. rate	1.34	1.12	1.52	1.55	1.39	1.24	1.53
Average Number of Disorder Concerns							
	Base Case			Full Info.	More Frequent Signals		
Plan Type	All	Type 1	Type 2	All	All	Type 1	Type 2
free	1.95	2.10	1.87	1.86	1.94	2.08	1.87
25% coin. rate	1.98	1.87	2.06	2.04	1.98	1.89	2.05
95% coin. rate	2.13	1.99	2.25	2.27	2.17	2.07	2.26

Note: There are 8 conditions included in the number of conditions from examinations and 21 conditions included in the number of concerns. The detailed list is in Section 2.1.

Table 16: Attrition Rates in the Base Simulations and the Counterfactuals

	Base Case			Full Info.	More Frequent Signals		
Plan Type	All	Type 1	Type 2	All	All	Type 1	Type 2
free	3.11%	4.23%	2.44%	2.42%	3.08%	4.15%	2.45%
25% coin. rate	2.97%	2.28%	3.51%	3.34%	2.96%	2.25%	3.50%
95% coin. rate	4.03%	2.55%	5.30%	5.96%	4.10%	2.62%	5.36%

References

J. Abaluck and J. Gruber. Choice inconsistencies among the elderly: Evidence from plan choice in the medicare part d program. volume 101, page 1180–1210, June 2011.

J. Abaluck and J. Gruber. Evolving choice inconsistencies in choice of prescription drug insurance. *American Economic Review*, 106(8):2145–2184, Aug. 2016. ISSN 0002-8282.

D. A. Ackerberg. Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination. *International Economic Review*, 44(3):1007–1040, 2003.

D. M. Anderson, A. Hoagland, and E. Zhu. Medical bill shock and imperfect moral hazard. *Journal of Public Economics*, 236:105152, Aug. 2024. ISSN 0047-2727.

P. Arcidiacono, E. Aucejo, A. Maurel, and T. Ransom. College attrition and the dynamics of information revelation. *Journal of Political Economy*, 2024.

M. Boyer, P. D. Donder, C. Fluet, M.-L. Leroux, and P.-C. Michaud. Long-Term Care Insurance: Information Frictions and Selection. *American Economic Journal: Economic Policy*, 12(3):134–169, 2020.

A. T. Ching, T. Erdem, and M. P. Keane. Learning models: An assessment of progress, challenges, and new developments. *Marketing Science*, 32:913–938, 2013.

G. S. Crawford and M. Shum. Uncertainty and learning in pharmaceutical demand. *Econometrica*, 73(4):1137–1173, July 2005.

C. Cronin. Insurance-Induced Moral Hazard: A Dynamic Model of Within-Year Medical Care Decision Making Under Uncertainty. *International Economic Review*, 60(1):187–218, 2019.

F. Cunha and J. J. Heckman. Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation. *Journal of Human Resources*, 43(4):738–782, 2008.

F. Cunha, J. J. Heckman, and S. M. Schennach. Estimating the Technology of Cognitive and Noncognitive Skill Formation. *Econometrica*, 78(3):883–931, 2010.

M. Darden. Smoking, Expectations, and Health: A Dynamic Stochastic Model of Lifetime Smoking Behavior. *Journal of Political Economy*, 125(5):1465–1522, 2017.

L. Einav, A. Finkelstein, S. Ryan, P. Schrimpf, and M. Cullen. Selection on Moral Hazard in Health Insurance. *American Economic Review*, 103(1):178–219, 2013.

T. Erdem and M. P. Keane. Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science*, 15(1):1–20, 1996.

D. Gilleskie. A Dynamic Stochastic Model of Medical Care Use and Work Absence. *Econometrica*, 66(1):1–45, 1998.

D. Gilleskie. Work absences and doctor visits during an illness episode: The differential role of preferences, production, and policies among men and women. *Journal of Econometrics*, 156(1):148–163, 2010.

M. Grossman. On the Concept of Health Capital and the Demand for Health. *Journal of Political Economy*, 80(2):223–255, 1972.

B. Handel, J. Kolstad, T. Minten, and J. Spinnewijh. The socioeconomic distribution of choice quality: Evidence from health insurance in the netherlands. *American Economic Review: Insights*, 6(3):395–412, 2024.

B. R. Handel and J. T. Kolstad. Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare. *American Economic Review*, 105(8):2449–2500, 2015.

K. Ho, J. Hogan, and F. S. Morton. The impact of consumer inattention on insurer pricing in the medicare part d program. *RAND Journal of Economics*, 48(4):877–905, 2017.

M. Hodor. Family Health Spillovers: Evidence From the RAND Health Insurance Experiment. *Journal of Health Economics*, 79, 2021.

R. Hosseini, K. A. Kopecky, and K. Zhao. The evolution of health over the life cycle. *Review of Economic Dynamics*, 45:237–263, 2022.

C. Kaufmann, T. Müller, A. Hefti, and S. Boes. Does personalized information improve health plan choiceswhen individuals are distracted? *Journal of Economic Behavior and Organization*, pages 197–214, 2018.

J. Ketcham, C. Lucarelli, E. Miravete, and C. Robeuck. Sinking, Swimming, or Learning to Swim in Medicare Part D. *American Economic Review*, 102(6):2639–2673, 2012.

J. D. Ketcham, C. Lucarelli, and C. A. Powers. Paying attention or paying too much in medicare part d. *American Economic Review*, 105(1):204–233, 2015.

J. R. Kling, S. Mullainathan, E. Shafir, L. C. Vermeulen, and M. V. Wrobel. Comparison friction: Experimental evidence from medicare drug plans. *Quanterly Journal of Economics*, pages 199–235, 2012.

T. Liu, N. Mehta, and S. Richards-Shubik. Multidimensional Health Capital and the Production of Health. *Working Paper*, 2023.

G. Loewenstein, J. Y. Friedman, B. McGill, S. Ahmad, S. Linck, S. Sinkula, J. Beshears, J. J. Choi, J. Kolstad, D. Laibson, B. C. Madrian, J. A. List, , and K. G. Volpp. Consumers' Misunderstanding of Health Insurance. *Journal of Health Economics*, 32(5):850–862, 2013.

N. Mehta, N. Jian, K. Srinivasan, and B. Sun. A Dynamic Model of Health Insurance Choices and Healthcare Consumption Decisions . *Marketing Science*, 36(3):338–360, 2017.

M. N. White. Self-reported health status and latent health dynamics. *Journal of Health Economics*, 88, 2023.