Essays on Incentives and Choice in Education and Health Economics

by

Keshav Garud

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2023

Doctoral Committee:

Assistant Professor Sara B. Heller, Co-Chair
Associate Professor Sarah M. Miller, Co-Chair
Professor Emeritus Charles Brown
Professor Tanya Rosenblat

Keshav Garud

kgarud@umich.edu

ORCID iD 0009-0004-8666-981X

To all underrepresented researchers, whose time is now

ACKNOWLEDGEMENTS

I am fortunate to have had the opportunity to work with a dissertation committee that constructively challenged me as I formed these chapters. Professors Sara Heller and Sarah Miller, my co-chairs, are extremely dedicated faculty and scholars. They played a pivotal role in guiding me towards clarity and polish. I feel grateful that they believed in me throughout the various stages of ideation, piloting, grant writing, fielding, analysis, and writing. I particularly value the strong emphasis they placed on my writing and presentation skills, which are cornerstones in communicating research that relates to policy. Furthermore, Professor Sara Heller's course on experimental methods proved incredibly useful and was perhaps my favorite part of my doctoral coursework. Without the combined mentorship of my co-chairs, I would not have reached this stage. I am grateful for the support of Professor Tanya Rosenblat, from whom I feel I could continue to learn and explore exciting questions in behavioral and experimental economics. She provided me motivation and context to help me continue to enrich my project, even when I felt unsure of what the next step may look like. Finally, I extend profound gratitude towards Professor Emeritus Charles Brown, whose keen eye and attention to detail is admirable and is something I realized is crucial and very much a part of mastery in the field of Economics.

In addition to this remarkable set of mentors, I feel grateful for having found close friendship within the Economics doctoral program with Hayley Abourezk-Pinkstone, Max Huppertz, Shwetha Raghuraman, Clay Wagar, Cameron Fen, Gerardo Sanz-Maldonado, Stephanie Karol, Tereza Ranošova, John Olson, and Thomas Flanagan. Hayley played a tremendous influence on my doctoral experience, and we formed a close bond from co-authoring our paper and providing each other moral support during challenging times. I further feel grateful for the relationships I formed with colleagues during the final few months of my doctoral dissertation, particularly as we experienced a unique moment in the University's history as graduate students. I would also like to thank the administrative office in the Economics Department; individuals who played a role in my success include Laura Flak, Hiba Baghdadi, Lauren Pulay, and Julie Heintz.

I am further grateful for significant support from several individuals as I put together these chapters. In particular, I am grateful for the California Department of Education for providing us with Open Enrollment School lists and communicating with us about the Open Enrollment Act, which certainly was pivotal for the completion of Chapter III of this dissertation. I am grateful for many colleagues and seminar participants for helpful comments from both presentations and informal meetings.[1] Your guidance was invaluable. I am also thankful for institutional support as I completed my dissertation chapters.[2] I would like to extend additional thanks to Professors Hoyt Bleakley, Sarah Jacobson (Williams College), Eduardo Montero (University of Chicago), Min Young Yoon (Rider University), and Vicki Bogan (Cornell University) for helpful feedback along the way.

As I look back to my long journey as a student and trainee, it is hard to not feel sentimental when I think about individuals who invested and believed in me. For instance, my summer working with Professor Melissa Dell was particularly formative; she played a significant role as an individual who I continue to look up to, both in terms of her curiosity and in terms of her work ethic and perseverance in the field. I feel grateful for the role that Daniel Beltran, Ricardo Correa, and Juan-Miguel Londono-Yarce played in my development in the field and particularly in fostering my interest in producing policy-driven research and opening my view to multiple real-world applications of economics. I am grateful for my undergraduate instructors who played a central role in my initial discovery of economics: Professors Mallesh Pai, Jere Behrman, and Ufuk Akcigit were significant influences for me at this stage. I would be amiss to leave out particularly important influences from before this: my high school English instructor Sandra Wyngaard pushed us to craft sentences and claims with support and purpose and my Mathematics instructor Greg Somers, who helped me see math as if something that possesses its own artistic expression. I did not know it back then, but my real training as an Economist began with them.

On a more personal note, I look back to my childhood attending public school in New York City, which came with its own challenges and rewards, and I ultimately feel blessed to have experienced such a diverse and unique education. This early education taught me that above all else, empathy and the ability to see things from other people's perspective counts far more than just being "right." I am fortunate (or unfortunate) enough to come from a family of academics, who were my "informal committee" during my doctoral studies. Professors

[1]In particular, I am grateful for feedback from participants in the University of Michigan Behavioral and Experimental Economics Lab, the Labor Seminar, and the Causal Inference in Education Research Seminar (CIERS).

[2]The experiment in Chapter I of this dissertation is pre-registered with a pre-analysis plan on the American Economic Association's registry for randomized controlled trials (AEARCTR-0008691). The study received an exemption (HUM00190253) from the Health Sciences and Behavioral Sciences IRB at the University of Michigan. I am thankful for generous project funding from the University of Michigan Economics Department as well as a Rackham Research Grant.

Sumita Raghuram and Raghu Garud (or, more simply, mom and dad) have supported my intellectual curiosity, motivated me to see the larger picture, and emphasized the importance of paying attention to the details. I will look back fondly at the period during my doctoral studies that we spent in quarantine together during the COVID-19 pandemic. I am grateful for the unwavering support of my sister and brother-in-law, Professors Nandita Garud and Aaswath Raman. They always believed in my ability to succeed in a doctoral program and they both helped me find the courage to continue even after I experienced a few setbacks in the beginning of my program. It has been delightful to watch my nephew, Anirudh, grow. An informal joke is that you cannot be a Garud without being a doctor: well indeed, this prophecy has finally come true for myself!

I have learned that I am a highly emotional individual, but I have taught myself to channel and celebrate this part of myself. I furthermore believe that distinct parts of my identity are distributed in different close individuals. I am grateful and blessed to have incredibly supportive friends outside of the program. Every drop of my happiness and joy belongs to them (and they know who they are). I feel particularly fortunate for the support of my best friend, Robert Hsu, who has offered me laughter, postcards, and who listened to me when I needed to be heard the most. It is challenging to reckon with ending multiple decades of formal education and training as an Economist all with a walk to the podium, the flip of a tassel, and a diploma mailed home. However, this is just a new beginning, and I am just getting started. The best days are yet to come.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

This dissertation consists of three chapters focusing on incentives and choices in the context of the economics of education and health economics. The first two chapters explore potential explanations underlying a puzzle in health economics: why do individuals appear to remain inadequately insured for their health? Furthermore, Chapters I and II explore the possible role of salient information about health, delivered to individuals in various forms, on health plan choices. Chapter III focuses on the socioeconomic and racial sorting effects of a policy that expanded school choice options for students attending low performing public schools.

In Chapter I "Who Avoids Health Information? Experimental Evidence on Health Insurance Choice," I explore the phenomenon of avoidance of information related to mortality, illness, and salient health events. Because information avoidance cannot be classified using existing observational data, I design an experiment to separate a group of "avoiders" and "non-avoiders" based on their willingness-to-pay to avoid uncomfortable or tedious information related to health. This approach enables me to document several observable traits about this newly defined group of "avoiders." I simulate a potential future health shock for all respondents and I observe that insurance preferences of avoiders and non-avoiders appear to be different. I rule out the role of different information exposure in helping explain this modest, yet significant difference in insurance preferences.

In Chapter II "Private Health Insurance Patterns Following Spousal Health Shocks," I explore how individuals select private health plans after a spouse has experienced a hospital admission. I document a 7.4pp increase in private plan coverage, specifically for individuals who did not initially hold a private health plan when they entered the longitudinal Health and Retirement Study. Across all individuals, both the previously uninsured *and* those who held insurance, I observe an increased $82 annual premium associated with private plans, suggesting individuals switched towards more generous coverage as a result of the spousal hospital admission. While I cannot rule out other possible explanations, such as lowered transaction costs, these results are consistent with individuals responding to salient health information by selecting into higher levels of insurance coverage.

Finally, in Chapter III "School Choice and Student Mobility from Low-Performing Schools:

Evidence from the California Open Enrollment Act," Hayley Abourezk-Pinkstone and I answer the question of how minority and low-income students respond to a policy incentive that allows them to transition to better-performing public schools of their choice. Using variation in the policy roll-out across different schools in different years, we find that minority and low-income students respond to the policy by exiting low-performing schools and substituting towards higher-performing schools. The effect is likely driven by a combination of both a *nudge* to shift schools (from a letter mailed home) and the actual policy that allowed students to switch schools.

CHAPTER I

Who Avoids Health Information? Experimental Evidence on Health Insurance Choice

1.0 Abstract

Does information avoidance, a phenomenon where individuals actively choose to not view costless information that may be possibly distressing or cognitively taxing, help explain why individuals select less generous health insurance plans? Using an online survey experiment to identify individuals' willingness to pay to avoid health consequence information, I find avoidance is relatively common: 24% of individuals in my sample are willing to pay to avoid health information. Avoiders are older, have lower income, are less likely to perform calculations when selecting plans, and have lower initial knowledge about definitions of insurance terms than non-avoiders. Presented with a hypothetical health issue, avoiders require a 2-3 pp higher likelihood of an adverse health event before selecting a more generous health plan. After randomly assigning avoiders to view facts and images highlighting the illness and mortality consequences of experiencing health issues, I find no measurable treatment effect on respondent preferences for health plans. The paper sheds new light on the fact that avoiders of health information are a distinct group with identifiable traits and finds that, independent of exposure to treatment information, they systematically select less generous health plans.

1.1. Introduction

Across the United States, 14% of adults under 65 years had no health insurance in 2021 (Cohen et al., 2021). Furthermore, more than 40% of the adults who have a private health plan are enrolled in a high deductible plan, resulting in the risk of out-of-pocket medical care spending (Kullgren et al., 2019). These choices may be rational: the marginal cost of higher premiums often associated with plans offering more generous coverage may outweigh the marginal benefit from utilizing these plans for medical care. Alternatively, information frictions or behavioral biases may lead individuals to make suboptimal choices.

There is existing empirical evidence suggesting that information frictions play a non-trivial role and may lead to underinsurance relative to an optimal individual and social level. First, individuals who previously went without health insurance altogether or who enrolled in plans with less generous coverage and were nudged by informational interventions into selecting more generous plans end up experiencing lower rates of mortality as a result (Goldin et al., 2021; Abaluck et al., 2020). Additionally, experimental research finds that educating individuals about insurance terms and making expected financial outcomes clearer results in individuals shifting towards more optimal plans (Samek and Sydnor, 2020; Bhargava et al., 2017; Kairies-Schwarz et al., 2017; Loewenstein et al., 2013; Schram and Sonnemans, 2011). Second, although uninsured individuals frequently cite plan costs as a deterring force against insuring, close to half of those who are uninsured have access to highly subsidized or zero-premium plans through multiple government programs that they still do not sign up for (Cox and McDermott, 2021).

This existing work has found that individuals alter their plan choices after information costs are lowered. Yet, providing low-cost information may be insufficient because individuals may actively avoid this information. Following Golman et al. (2021), active avoidance takes place in a wide variety of settings when an individual knows that information exists and this would be low-cost or costless or to obtain, and yet an individual refuses this information. People may refuse to obtain certain types of information because they may have fear or they may believe this information is mentally discomforting, among other potential reasons.[1] Indeed, individuals may go so far as to be willing to pay to avoid information, making avoidance itself costly. While past studies document widespread avoidance of personal health risks (Kiefer et al., 2014; Barbour et al., 2012), this paper devises a method to test how avoiding salient health information about illness and mortality relates to insurance plan choices.

A key challenge to answering this question is that it is impossible, using existing data, to hold inattention and information acquisition costs fixed in a way that would allow a clean separation of avoiders and non-avoiders. Additionally, by definition, it is not possible to observe whether avoiders would change their insurance choices if they actually observed this information. To explore this hypothesized role of information avoidance in explaining

[1]Other reasons an individual may avoid health information include belief preservation, cognitive stability, mental discomfort, fear, autonomy, obligations to act, cognitive dissonance, and self-image concerns (Exley and Kessler, 2021; Hertwig and Engel, 2016; Sweeny et al., 2010). Distinguishing between specific reasons for avoidance is not the focus of this paper.

risky health insurance choices, I use the online participant platform Prolific to conduct an experiment with a representative sample of 1,638 unmarried adults aged 26 and above from all 50 states across the United States.

I start by measuring respondents' motivation, learning, and knowledge about health risks and insurance terms. After initially quizzing respondents on four key insurance terms, I teach them the meaning of the terms and then quiz them again. This sequence allows me to measure how quickly respondents learn, which I later explore as a baseline characteristic. The remainder of my experiment allows me to answer three questions: (1) how frequently do individuals avoid information about health consequences, (2) do avoiders have distinctive real-world characteristics and select significantly different plans, and (3) do avoiders respond differently to treatment with health information that they preferred to avoid in the first place?

To answer the first question, I measure avoidance by using an experimental approach to classify avoiders as a distinct group of individuals who are willing to pay to avoid health consequence information in favor of neutral information. Participants may select whether they prefer to view information about health information for a higher monetary incentive or whether they prefer to view neutral, non-health related information that would take the same amount of time, for a lower incentive. In other words, the marginal individual who chooses to avoid health information would be willing to forsake *at least* the additional bonus amount they would be provided by viewing health information in the experiment.[2] Some of my respondents face a low-stakes scenario where they have to be willing to forsake a "small" potential bonus to not view health information. Other respondents face larger stakes such that for an individual to be classified as an avoider, they must be willing to forsake a "larger" bonus.

I document that the prevalence of health information avoidance is non-trivial. In my sample, 24% are willing to pay to avoid health information. I find that for individuals surveyed using the "small" bonus spread, the identified group of avoiders select significantly less generous health plans in the lab setting. The "large" bonus spread condition provides compelling evidence as well: 17% of respondents under this condition are willing to forsake an additional bonus of $0.40, which amounts to almost 10% of pre-bonus survey compensation, in order to avoid health information (Figure A.1).

I explore the second question of who avoiders are and how they select plans by studying

[2]Strictly defined, "information avoiders" might actively dislike any form of information in any context (including neutral or positive contexts). However, this is a strong assumption about information avoiders; instead, it is more plausible that avoiders have a distaste for particular domains of information, of which information on consequences of health risks may be one such example. In the context of this paper, it is important to interpret avoiders in this context as those who specifically avoid health information.

avoiders' traits. Avoiders differ significantly from non-avoiders along numerous observable characteristics. Avoiders tend to be lower-income, older, uninsured, and with lower baseline levels of health insurance term knowledge compared to non-avoiders. Avoiders are significantly less likely to report performing calculations during the experimental choice between health insurance plans.

I can also measure avoiders' plan choices by exploring their insurance plan choices in a lab setting. I measure all individuals' plan preferences from a hypothetical health insurance choice scenario where individuals choose between three distinct plan options for varying probabilities they will encounter a health issue in the next year. The plan options individuals may select between are: (1) no insurance, (2) a "Bronze" (high deductible, low premium plan), and (3) a "Gold" (low deductible, high premium plan). By studying the switching patterns between plans for individuals, I form outcome measures that shed light on preferences for riskier plans. I find that avoiders choose riskier health plans. Avoiders wait until their health risks are around 3pp higher before switching to a low-deductible plan from either a high-deductible plan with a lower premium or from no insurance and around 2pp higher before switching to any form of health insurance. While previous studies have documented that risky health behaviors may result from individuals deciding not to undergo medical testing (e.g. Oster et al. (2013)), to my knowledge, my paper is the first to document a connection between information avoidance and risky plan choices.

Third, I test whether treatment with health information impacts avoiders differently. I use stratified randomization to assign avoiders and non-avoiders to a treatment where they are taught information about health risks and treatments that plans commonly cover. To avoid deception, I assign individuals with a higher probability to their preferred information option. There are two treatment arms: (1) only facts or (2) facts and images. The treatment arm with facts exposes individuals to statistics and facts relating to health risks such as annual doctors' and emergency room visits, along with health consequence information about hospitalization visits, annual deaths due to remaining inadequately uninsured, and an anecdote about an individual passing away after falling ill without health insurance. The second treatment arm shows respondents the same facts, along with accompanying images that heighten the consequences experiencing a health issue: an individual being loaded into an ambulance, sick individuals being treated in hospital beds, a deceased individual, and burials and family members grieving. Plausibly, the "facts and images" treatment allows me to study the effect of increased salience due to heightened emotional responses to images not shown to the "facts-only" treatment.

Using the outcome measures from the health insurance choice scenario (which respondents complete post-treatment), I find a small, insignificant impact of health information on insurance choices made by all survey respondents and can rule out evidence that treatment results in meaningful shifts toward more generous health insurance plans. Further, avoiders do not respond more to health information than non-avoiders. I find no distinction in risk preferences for the "facts"-only and "facts and images" treatments. Why is there no observable treatment effect? A possible explanation is that individuals are not sensitive to the specific form of health consequence information provided in this intervention. To further probe the causal nature of avoidance on risky plan choices, future research may explore whether there are other forms of information interventions, perhaps those more specifically tailored to individual characteristics, that result in measurable shifts in plan preferences, particularly for avoiders.

This paper contributes to a wide literature in health economics whose aim is to explore how consumers choose insurance plans. This literature has already documented that individuals do a relatively poor job in choosing plans that are appropriate for their underlying health risks (Handel et al., 2020). Furthermore, it has been established that inertia, high switching costs, imperfect information, and inattention appear to result in inconsistent plan choices (Brot-Goldberg et al., 2021; Handel et al., 2020; Handel, 2013; Polyakova, 2016; Abaluck and Gruber, 2011, 2016). However, the existing papers do not establish whether some people actively resist information even when there are interventions to overcome inattention. This would result in persistent inefficiencies. This paper contributes by exploring the extent that underinsuring would persist even after lowering information acquisition costs and holding inattention fixed. I thereby shed light on the previously under-explored role of information avoidance.

This paper contributes to a literature in behavioral and experimental economics that explores information avoidance. Golman et al. (2017) suggest that there may be a strategic advantage to ignoring information. People often avoid information that makes them feel bad as there are immediate hedonic benefits to avoiding negative information, which also carries direct and negative utility. Golman et al. (2021) introduce the term 'motivated attention' to refer to the desire to savor good news and ignore bad news. Avoidance refers to the incentive individuals have to avoid information even when it is freely available ('ignorance is bliss'), which contrasts with the notion that individuals would prefer to have more information, or that individuals are uninformed because information is costly to obtain. Researchers have established multiple determinants of avoidance (e.g. Exley and Kessler (2021)). The role of avoidance has not been previously explored in the health insurance context. I contribute by exploring the consequences of avoidance by providing new evidence linking avoidance and financial decision-making. I also provide evidence that avoidance plays a role in the setting of health insurance decisions.

The paper proceeds as follows: Section 1.2 devises an experimental solution that enables measurement of avoidance and studies whether avoiders respond more to information about health consequences than non-avoiders. Section 1.3 describes characteristics about avoiders and establishes a link between avoidance and more risky plan choices. Section 1.4 presents results for the causal impact of treatment with health information on plan risk preferences for avoiders and non-avoiders. Section 1.5 concludes.

1.2. Experimental Design

1.2.1 Survey and Sample Details

I pre-registered my study in the American Economic Association's RCT Registry[3]. I fielded the survey between December 2021 and February 2022[4] to participants using an online participant recruiting platform called Prolific. I screened my sample to include respondents who were 26 years and older, did not have dependents, were not currently married, and were U.S. citizens and nationals[5]. I directed participants to complete a Qualtrics survey that on average took 17.2 minutes to complete. Upon opening the survey link, respondents read a description of the survey and consented to partake.

I dropped 90 individuals based on a pre-specified decision rule in my pre-analysis plan[6], giving me a sample size of 1,558 respondents across the United States. Table 1.1 presents sample characteristics.[7] The sample is highly educated and relatively young with a mean age of 35. This is younger than the mean age of an analogously-constructed sample (that passes the age and demographic criteria required for survey participation) from the American Community Survey 2020, primarily because online survey platforms tend to skew towards younger participants. Survey respondents are from all 50 states in the United States.

[3] The corresponding RCT ID is AEARCTR-0008691.

[4] The national open enrollment period for health insurance took place between November 1st, 2021 to January 15th, 2022.

[5] I impose the restrictions "unmarried" and "no dependents" so that I can present hypothetical plans to respondents in the experiment that individual plans rather than family plans.

[6] I dropped individuals who displayed inconsistent switching choices, following my pre-analysis plan (Appendix A.3 contains more details).

[7] I do not re-weight the survey sample to match the American Community Survey 2020. I make this decision so as to preserve accuracy for inference purposes. As with any online experiment, my results are internally consistent for population represented by internet survey-takers.

Table 1.1: Sample Characteristics of Survey Respondents

	Survey Sample	5-Year 2020 ACS
Number of respondents	1558	–
Male*	45.13%	51.36%
Female*	52.03%	48.63%
% Assigned to Treatment	52.82%	–
% Avoider	23.91%	–
% Uninsured (health)	17.27%	13.23%
Mean Age	34.69	51.51
% With High School Degree	75.87%	87.36%
% earning below $20,000	26.85%	38.33%
% White	74.78%	70.45%
% Democrat	55.97%	–

Note: [1] Column 1 presents sample characteristics. [2] *Male* and *Female* in Column 1 do not add up to 100% because a small proportion of survey respondents classify themselves as *Non-binary*. [3] Column 2 presents population characteristics for individuals who pass the filters required to qualify to be a survey respondent: (1) Above 26 years, (2) No dependents, (3) Not currently married, (4) U.S. National. The population values are drawn from the 5-year 2020 American Community Survey. I apply individual-level weights to calculate the population level characteristics in Column 2.

1.2.2 Experiment Design

A brief overview of the survey design follows, with details on each portion provided in subsequent subsections. I start by measuring baseline knowledge as well as measuring learning rates of the health insurance terms that I teach the respondents. I then identify individuals' willingness to pay to avoid health information and then assign respondents to treatment arms through stratified randomization. Individuals step through a series of materials related to health facts (treatment) or a logic game (control). After this randomization into treatment, all respondents face a prompt where a hypothetical health issue may occur in the next year with different probabilities. They may select between different plan options that range in their plan deductible and out of pocket costs. This allows me to measure individuals' preferences for risky plans. I form primary outcome measures of plan preferences based on individuals' responses. After the hypothetical health issue, individuals answer several background demographic questions, and then the survey concludes.[8]

1.2.2.1 Measuring Baseline Knowledge, Learning, and Motivation

I measure pre-treatment characteristics about individuals' perceptions of relatively common medical risks (the percent of adults who had a visit to a doctor in the past year and the number of annual ER visits). Individuals complete a definition quiz on four common insurance terms (premium, deductible, coinsurance, and out of pocket maximum), are then all shown the term meanings, and complete the same quiz after seeing the definitions. I incentivize respondents by informing them they will receive a bonus if they get all terms correct on either or both quizzes and inform them to attempt their best effort for the first quiz because this will maximize their chances of receiving a bonus. I use this sequence to measure baseline health and insurance knowledge and learning. I also ask individuals to state how much time they would be willing to spend to learn about plan details in real life if they were selecting an insurance plan. This question gives me a measure of an individual's motivation to research plans characteristics.

1.2.2.2 Identifying Information Avoiders

I devise an approach to plausibly identify avoiders of health information. To do so, I elicit individuals' willingness to pay to avoid salient health information. I ask individuals to select between two information options: viewing health information that discusses mortality and illness for bonus payment X or viewing non-health related information, which is later revealed to be a walk-through on how to play a logic game, for bonus payment Y (Figure 1.1).

[8]Appendices A.3.2 and A.3.3 discuss additional experiment details related to survey design decisions. Appendix A.6 shows screenshots of the actual survey instrument.

Figure 1.1: Avoid Prompt

There are two options to proceed. You will receive the full survey compensation for either option. We have no preference for which option you select. You are not guaranteed to get the option you choose. But, you should select the option you prefer the most because you will be assigned with a higher probability to the option you choose.

Option 1: We will walk you through information about issues such as hospitalizations, mortality, and what health insurance covers. **There will be facts about illness and death. You may also be shown graphic images. Either of these could make you feel uncomfortable or may be upsetting, unpleasant, tedious, or could even feel monotonous and tiring.** If assigned to Option 1, you will be given an additional **$0.15 bonus**.

Option 2: We will walk you through different information that is **not related to** health insurance, illness, or death. This will take the same amount of time to go through as Option 1. If assigned to Option 2, you will be given an additional **$0.10 bonus**.

Please select the option you most prefer:

Option 1 (Health insurance information)

Option 2 (Non-health information)

Note: The prompt classifies two groups of individuals: avoiders and non-avoiders of health information. All respondents in the survey view this prompt (or a variation thereof) where the bonus that an individual assigned to Option 1 receives varies between $0.15 to $0.50.

X is strictly greater than Y. Individuals are informed that they will receive the respective bonus if they are assigned to that particular option. Holding Y fixed at $0.10, I vary X in increments of $0.05 between $0.15 to $0.50. Figure A.1 shows the number of respondents I gathered for each value of X. Respondents select their preferred option with the knowledge that they will be assigned with a higher probability to this option. Respondents are assigned with a 55% chance to their selected option, although they are not informed of the exact probability.

1.2.2.3 Treatment Arms

I stratified treatment assignment on avoidance. To separate out the role that facts alone play from the role of heightened salience from an emotional response from images, I create two treatment arms for individuals assigned to view "health information" about general health risks and the consequences of experiencing these (facts about mortality and illness). The first treatment arm is a "low salience" group that only views facts (25% of all treatment individuals are assigned to this treatment). This treatment first provides context on general health information by reporting the percent of individuals in the United States who visited a medical provider in the past year and the number of people foregoing needed care because of inadequate health insurance. Individuals may understand their chances of experiencing certain health issues, such as pneumonia or a fracture, but may lack information about what happens after this occurs (i.e. the need for hospitalization or the possibility of mortality). So, interspersed with the general health context, the treatment emphasizes information about the consequences of experiencing particular health issues.

To highlight health consequences, the treatment provides information about the common reasons individuals are hospitalized in the United States, along with statistics on the number of annual visits to the emergency room and hospital stays. The treatment explains consequences of falling sick through facts of annual deaths due to individuals being underinsured as well as a real anecdote about an individual in the United States who lost health insurance and ended up passing away because she felt deterred from seeking out medical care when she experienced a health issue. The second treatment arm is a "high salience" group that is presented the exact same facts as treatment one along with images (75% of all treatment individuals are assigned to this treatment). The images emphasize the consequences of experiencing health issues through heightened emotions. The images are of hospitalized individuals, people boarding an ambulance, and individuals grieving at burial sites, emphasizing mortality. I also quiz all individuals in both treatment arms once they have finished going through information and they have the opportunity to receive bonuses for additional correct answers.

In addition to the two treatment arms, there is a control arm where I teach individuals how to play a logic game. The purpose of this arm is to ensure that control-group individuals are still actively involved in a task that is not related to health information.

Fifteen percent of my sample were assigned to Treatment 1 (facts only), 38% were assigned to Treatment 2 (facts and images), and 47% were assigned to the control group. Table A.1 and Table A.2 show that pre-treatment survey characteristics are balanced across treatment groups. Almost all covariates are balanced across treatment groups, with the exception that treatment appears to contain a larger proportion of males and Asians (Table A.1). For treatment balance carried out independently for avoiders and non-avoiders (Table A.2), a few covariates appear slightly unbalanced. This is as would be expected by chance. Furthermore, joint F-tests for all of the covariates in both cases are insignificant.

1.2.2.4 Soliciting Plan Choice Preferences

After respondents have finished going through either the treatment or control interventions, all respondents are redirected to a common page and informed they could experience a hypothetical health issue in the next year, costing $25,000 without insurance. I provide respondents with three different plan options to choose between. They may choose no insurance, a high-deductible and low premium plan, or a generous plan that has a low deductible and high premium. The high-deductible plan is called "PPO Bronze" and the generous plan is called "PPO Gold". Each hypothetical plan has four characteristics associated with it: monthly premium, deductible, coinsurance, and out of pocket maximum, which are terms that I previously taught all respondents prior to treatment. These characteristics correspond to those found on actual health insurance plans listed on the health exchange, but I change the plan names (Appendix A.3.4). As Figure 1.2 depicts from the survey, for different probabilities (ranging from 0% to 100%, in increments of 10%) of potentially experiencing the health issue, respondents are asked to select their preferred plan choice. To make these responses compatible with preferences, I also ask respondents to write a paragraph explaining their choices in the health simulation, and I reward "thoughtful and insightful" responses with an additional bonus.

Figure 1.2: Health Insurance Choice Scenario

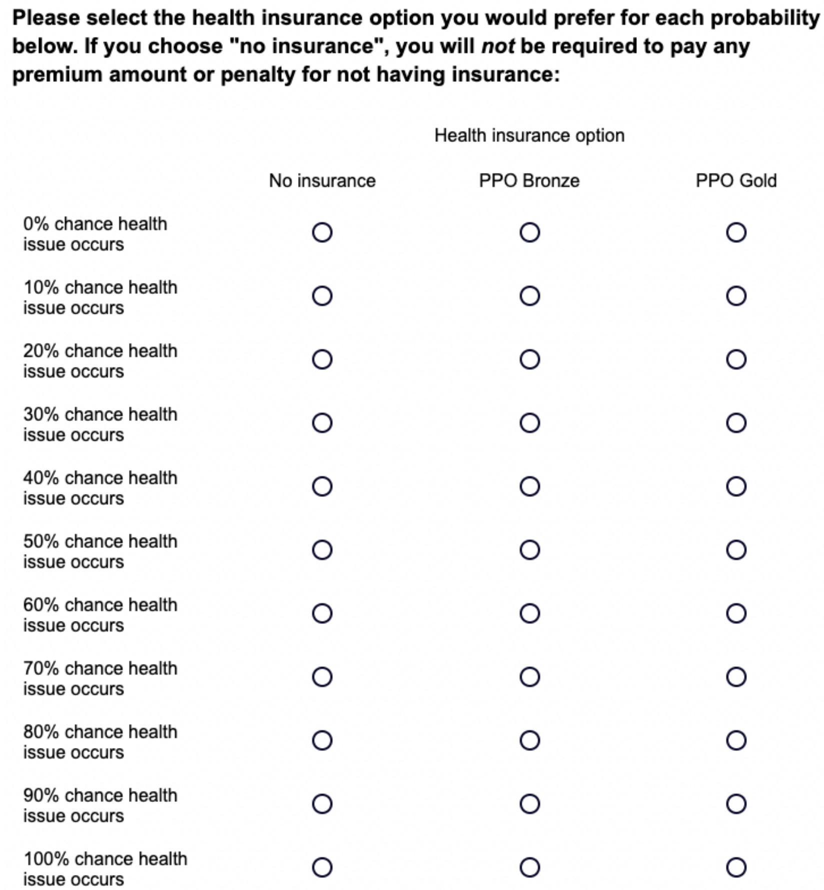
Please select the health insurance option you would prefer for each probability below. If you choose "no insurance", you will *not* be required to pay any premium amount or penalty for not having insurance:

	Health insurance option		
	No insurance	PPO Bronze	PPO Gold
0% chance health issue occurs	○	○	○
10% chance health issue occurs	○	○	○
20% chance health issue occurs	○	○	○
30% chance health issue occurs	○	○	○
40% chance health issue occurs	○	○	○
50% chance health issue occurs	○	○	○
60% chance health issue occurs	○	○	○
70% chance health issue occurs	○	○	○
80% chance health issue occurs	○	○	○
90% chance health issue occurs	○	○	○
100% chance health issue occurs	○	○	○

Note: All respondents may choose between three plan options: no insurance, PPO Bronze ("high deductible" plan), and PPO Gold ("generous" plan). Individuals may select any option for each health issue probability. I pre-specified that I would drop individuals who displayed non-monotonically increasing switch decisions. To form my primary outcome variables, I study the health risk probabilities when individuals switch to any insurance plan and when individual switch to the generous plan.

1.2.3 Outcome Variables

The health choice task above allows me to construct two primary outcomes and three secondary outcomes.[9] The primary outcomes focus on understanding how high health risks need to be before an individual chooses a more generous plan. The primary outcomes are: (1) the health risk probability when an individual switches to any health plan from no insurance and (2) the health risk probability when an individual switches to the "generous" (low deductible, high premium) health plan from the "high-deductible" (low premium) plan or

[9]Appendix A.3.5 presents the decision rules I followed from my pre-analysis plan along with additional details about dropping inconsistent choice patterns. The full pre-analysis plan is accessible through the following link: **https://doi.org/10.1257/rct.8691**

from no insurance. My secondary outcomes are counts of the number of times an individual selected the generous plan, high-deductible plan, or no insurance plan.[10]

1.3. Observable Characteristics of Avoiders

1.3.1 Who Are Avoiders?

Avoidance behavior is not typically observable in standard datasets commonly used by health economists. I provide descriptive evidence that characterizes avoiders and compares them to non-avoiders, taking advantage of the elicitation approach that allows me to identify them. Table 1.2 summarizes differences between avoiders and non-avoiders. Column (3) provides relevant demographic differences for avoiders. There is suggestive evidence that avoiders are more likely to be uninsured for health care ($P < .10$). Avoiders are 5pp less likely to earn above $20,000 and are 6pp less likely to hold a college degree. Avoiders are 5pp less likely to know anyone who has been hospitalized in the past year. Avoiders are, on average, almost two years older than non-avoiders in the sample. One possible explanation is that older individuals are more averse to potentially experiencing negative emotions from information shocks and prefer to avoid these due to prior health experiences.

There are also notable differences in knowledge and choice characteristics for avoiders. For instance, avoiders are 10 pp less likely to perform calculations during the hypothetical plan choice task than non-avoiders. Second, avoiders have an initial average score on the insurance term quiz that is statistically significantly lower by 0.23 points, out of 4 possible points ($p < 0.01$) . After teaching insurance term definitions to avoiders and non-avoiders, the average quiz score for both groups increases and they both have almost identical scores that are not statistically different. The finding that avoiders initially have worse knowledge than non-avoiders about insurance terms suggests that the measurable distaste for health consequence information ("avoid" in this sample) correlates to measurable, pre-existing gaps in *other* forms of insurance knowledge such as financial competency, which have been documented to result in sub-optimal insurance choices (Samek and Sydnor, 2020). Moreover, the finding that avoiders catch-up to non-avoiders suggests that it could be possible to, with the right informational nudge, help these individuals fill in missing information due to prior avoidance tendencies. However, despite the suggestive nature of this result, the final row of Table 1.2 suggests that treated avoiders appear to perform worse on a quiz about health knowledge that they were just taught. In Appendix A.1.2, I provide graphs with additional details on different characteristics of avoiders.

[10] I include further discussion of outcome variables in Appendix A.3.4.

Table 1.2: Characteristics of Avoiders

	(1) Avoiders	(2) Non-avoiders	(3) Difference
N	371	1,187	
Demographic Characteristics			
% Uninsured (health)	19.67	16.51	3.92^{*} (2.27)
% Earning Above $20,000	68.55	74.56	-5.28^{**} (2.62)
% w/ at least college degree	18.86	23.67	-5.77^{**} (2.55)
% Knowing any acq. hospitalized	68.75	73.80	-5.43^{*} (2.69)
Average Age	36.04 yrs.	34.27 yrs.	1.88^{***} (0.57)
Knowledge and Choice Characteristics			
% Performed Calculations	28.03	37.99	-10.16^{***} (2.87)
Pre-treatment health risk beliefs (out of 100)	47.92	49.71	-1.92^{*} (1.05)
Insurance term quiz 1 score (out of 4)	2.64	2.85	-0.23^{***} (0.08)
Insurance term quiz 2 score (out of 4)	3.71	3.71	-0.00 (0.04)
Post-treatment health knowledge (out of 100)	80.91	85.05	-4.19^{**} (1.79)

Note: Column (3) reports the difference between avoiders and non-avoiders after controlling for the bonus spread. *, **, and *** denote significance at the 10, 5, and 1 percent level.

1.3.2 How Do Avoiders Select Plans in a Plan Choice Simulation?

I estimate the following equation to study whether avoiders choose to enroll in different plans than non-avoiders:

$$Outcome_i = \beta_1 Avoid_i + \beta_2 X_i + \epsilon_i \tag{1.3.1}$$

Equation 1.3.1 documents whether avoiders make different health insurance plan choices than non-avoiders. The coefficient β_1 measures the extent to which avoidance is correlated with plan choices. In some of the model specifications, I include a vector of controls X_i, which includes gender dummies, income bin dummies, age, health insurance status, race dummies, education dummies, number of acquaintances hospitalized in the past year, indicators for whether an individual is Democrat, whether they performed calculations in the health insurance choice simulation task, self-rated mental and physical health, procrastination scores, time willingness to spend learning about health plans, insurance term quiz scores (before and after being taught the terms), initial beliefs about health risks, survey characteristics (duration in minutes, pre-treatment time, and clicks), a dummy for assignment to treatment, and the bonus spread between viewing health information and non-health related information that different experimental arms were shown.

Across all respondents, the majority switch to the generous plan when the probability that hypothetical health event from the experiment is between 30%-70% and a majority of respondents switch to any insurance plan when the probability is 30% or under. Yet, switching patterns differ across groups, as is apparent from Figure 1.3. For every potential health risk, a smaller proportion of avoiders (red bars) switch to either any insurance plan or choosing the generous plan. Table A.3 presents the outcome means for avoiders and non-avoiders.

Figure 1.3: Avoiders versus Non-Avoiders: Distribution of health risk probabilities when respondents switch in the health insurance simulation

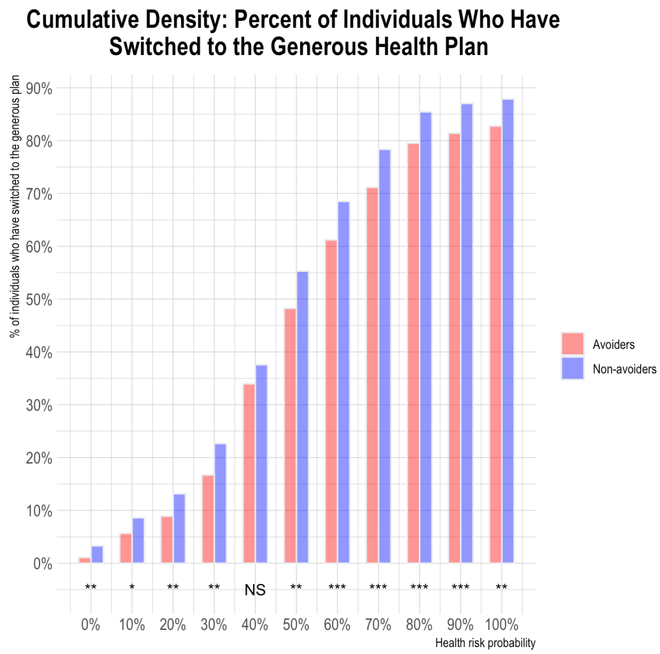

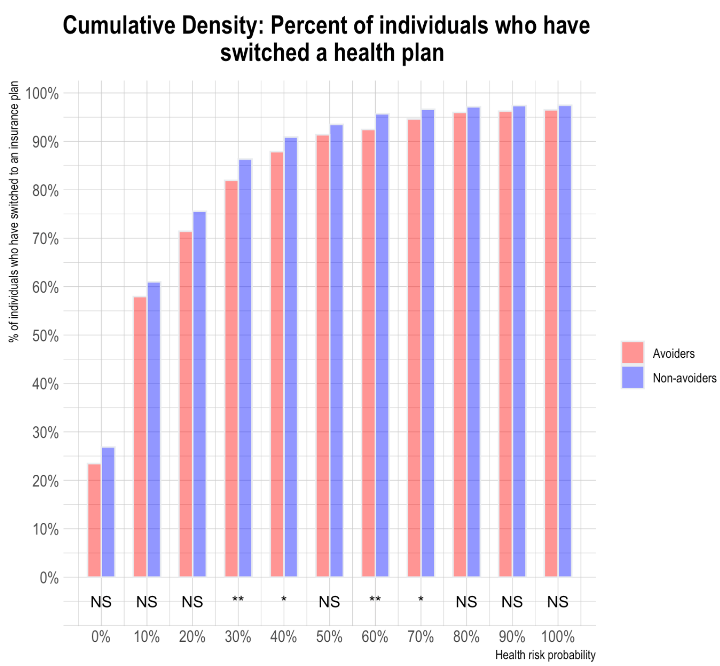

Note: I report statistical significance from running separate regressions comparing the % of individuals in each relevant comparison group (treatment versus control, avoiders versus non-avoiders) that have switched to the generous plan by the time their health risk probability has reached X%. The figures report coefficients from the following regression: $\mathbb{1}(Switch\, before\, X\%)_i = \beta \text{Avoid}_i + \epsilon_i$

Table 1.3 shows that on average, avoiders wait until their health risks are higher before switching to more generous plan choices. Prior to including controls X_i, I find that avoiders wait until their health risks are higher by 2.0 pp before switching to any insurance and higher by 3.0 pp before switching to the most generous plan option (column 1).[11] As observable in column (2), I control for differences in informational content (assignment to treatment), which by design is correlated with avoidance. After controlling for information, avoiders still appear to select riskier plans than non-avoiders: avoiders wait until health risks are 2.1 pp and 3.1 pp higher before switching to insurance and to the generous plan, respectively. This suggests that there may be an underlying tendency of avoiders to choose riskier plans. Moving from column (2) to (3), it is apparent that controlling for covariates found in the full vector X_i results reduces the size of the magnitude on avoidance somewhat, but it still appears that avoiders select different plans even after including these controls. Avoiders wait until their health risk is higher by 2.7 pp before switching to the generous plan.[12]

[11]Appendix A.1.1 provides estimates on the explanatory relevance of avoidance on the health risk probability when individuals in specific population groups switch to the most generous health plan. The findings suggest that within specific population groups, avoidance plays an even larger role in explaining risky plan choices. For instance, non-democrat avoiders switch to the generous plan when their health risks are 6.5 pp higher than non-democrat non-avoiders.

[12]Even after further limiting the sample in columns 1 and 2 to match the sample in column 3, I find that the magnitude shrinks due to the inclusion of covariates.

Table 1.3: Estimates of insurance choice outcomes on independent variable of interest: *Avoid*

Outcome:	(1) No controls	(2) Control for treatment	(3) Control for treatment + full vector of controls	(4) Outcome Mean
[1] Switch to generous plan health risk %	3.05**	3.14**	2.66**	48.36
	(1.41)	(1.41)	(1.44)	
[2] Switch to plan health risk %	2.03*	2.11**	1.19	16.24
	(1.04)	(1.04)	(1.03)	
[3] Count no insurance choices	0.29**	0.30**	0.07	1.88
	(0.14)	(0.14)	(0.13)	
[4] Count high-deductible plan choices	0.29*	0.28*	0.36**	3.75
	(0.16)	(0.16)	(0.17)	
[5] Count generous plan choices	−0.57***	−0.57***	−0.43**	5.37
	(0.17)	(0.17)	(0.17)	
Max. N respondents	1,558	1,558	1,492	1,558

Note: This table reports coefficients on the independent variable of interest (Avoid) under three model specifications: (1) no included controls, (2) controlling for the bonus spread and assignment to treatment, and (3) controlling for the full vector X_i of individual demographic characteristics, which also includes the bonus spread and treatment. The first two rows present results for my main outcomes (health probability when an individual switches plan options) and and the last three rows describe the number of times individuals selected each plan choice, out of 11 possible options in the health simulation. The number of respondents in each column is reported as a range where the lower number in each range corresponds to the number of non-missing observations for row [1] and the higher number corresponds to the number of non-missing observations for rows [3] through [5]. Column 4 presents the mean value of each outcome variable in the dataset (for both avoiders and non-avoiders combined).
*, **, and *** denote significance at the 10, 5, and 1 percent level.

1.3.2.1 Interpreting Magnitudes

I interpret the magnitudes from Table 1.3 by benchmarking these against the mean and standard deviation of each outcome variable. I present the benchmarked magnitudes in Table 1.4. The point estimates suggesting that avoiders wait until their health risks are 3pp and 2pp before switching to the generous plan and to any plan, respectively, correspond to a 1/7 to 1/9 of a standard deviation increase in health risk probability relative to non-avoiders. These coefficient sizes, along with comparisons of the estimate sizes relative to *mean* choice patterns observed in the simulation, suggest that the sizes of the magnitude for preferring less generous plans is moderate for avoiders.

A different way to benchmark the effect sizes is to compare the observed difference in simulation choice patterns between avoiders and non-avoiders to the observed difference across the "actually" uninsured and insured. Uninsured individuals in the sample wait to switch to any insurance when their health risks in the health simulation are 13.8pp higher than insured individuals and to the generous health insurance plan when their health risks are 7.7pp higher than insured individuals. For instance, the point estimate (Table 1.3) on the health risk when avoiders switch away from no insurance to a plan is 2pp higher, which corresponds to almost 15% of the magnitude of when uninsured individuals wait to switch to an insurance plan in the health simulation. The point estimate on when avoiders switch to the generous plan is a 3pp higher health risk, corresponding to almost 40% of uninsured individuals' switch behavior.

Table 1.4: Relative size of coefficient on *Avoid*

	(1) Mean	(2) σ	(3) Estimate size Relative to Mean	(4) Estimate size Relative to σ
[1] Switch to generous plan health risk %	48.3629	21.7281	6.30%	14.02%
[2] Switch to plan health risk %	16.2442	17.2120	12.49%	11.78%
[3] Count no insurance choices	1.8793	2.2883	15.20%	12.48%
[4] Count high-deductible plan choices	3.7528	2.6702	7.60%	10.69%
[5] Count generous plan choices	5.3677	2.8922	10.64%	19.75%

Note: This table reports the size of the point estimates for the independent variable of interest, Avoid, without including any covariates (i.e., results from column (1) in table 1.3). These are expressed as a percent of the mean and standard deviation of each outcome measure.
*, **, and *** denote significance at the 10, 5, and 1 percent level.

1.4. Treatment Effect from Salient Health Information

I randomly assign respondents to view salient health information in the form of facts about mortality and illnesses (Treatment one) and the same statistics in addition to images about hospitalizations along with the same facts from group one (Treatment two). I measure the intent-to-treat estimate from assigning individuals to either of these two treatments (i.e., pooled treatment effect) using the equation below:

$$Y_i = \alpha + \theta_1 Treatment_i + \theta_2 Avoid_i + \epsilon_i \tag{1.4.1}$$

where the coefficient of interest is θ_1. I also include an indicator for whether an individual is an avoider ($Avoid_i$) because avoiders have a different treatment assignment probability from non-avoiders. I also test whether the treatment effect differs significantly for avoiders compared to non-avoiders, which I measure in an interaction term, θ_3, between avoidance and treatment assignment. The equation below tests this:

$$Y_i = \alpha + \theta_1 \mathbb{1}(Avoid_i) + \theta_2 Treatment_i + \theta_3 \mathbb{1}(Avoid_i) * Treatment_i + \epsilon_i \tag{1.4.2}$$

This equation tests the hypothesis that avoiders respond more to treatment material than non-avoiders (θ_3 is negative and significant). I hypothesize that avoiders would respond more to treatment by switching to more generous plan options because, by definition, they are more likely to have "avoided" this information in the real world. Random assignment to treatment, therefore, may have a greater effect due to the information being more novel and more shocking to a group of individuals who previously had a proclivity to avoid it.

Columns (1) and (3) in Table 1.5 present results from equation 1.4.1, which regresses the primary outcome variables on an indicator for pooled treatment (treatment arms one and two). Columns (2) and (4) presents results from equation 1.4.2, which regresses the primary outcomes on separate indicators for treatment assignment and avoidance preference, as well as an interaction term. I find that neither avoiders' or non-avoiders' plan choices respond to the treatment. The point estimates on treatment (columns 1 and 3) are slightly positive but close to zero. These are not statistically significant at conventional levels. The 95% confidence interval for the treatment estimate (Column 1) is (-1.58, 3.06). The lower bound rules out the fact that even if there were a treatment effect, this results in a meaningful shift towards selecting the more generous plan option. Given the fact that individuals may select different health plans for health risks in increments of 10%, the point estimate is precisely estimated, as a meaningful treatment effect would be that individuals switch to a generous plan when their health risks are at least 10pp lower. Figure 1.4 presents visual evidence suggesting that treatment and control individuals have a similar distribution in plan choices across different health risk probabilities. Additionally, the point estimates on the interaction term are small and statistically insignificant. I conclude that there is no effect of assigning either avoiders or non-avoiders to view the particular health information I selected for individuals. In Appendix A.4, I also present results for Treatment 1 (facts only) and Treatment 2 (facts and images) separately. Similar to the pooled treatment effect, there is no measurable effect of the treatment arms separated.

Figure 1.4: Treatment versus Control: Distribution of health risk probabilities when respondents switch in the health simulation

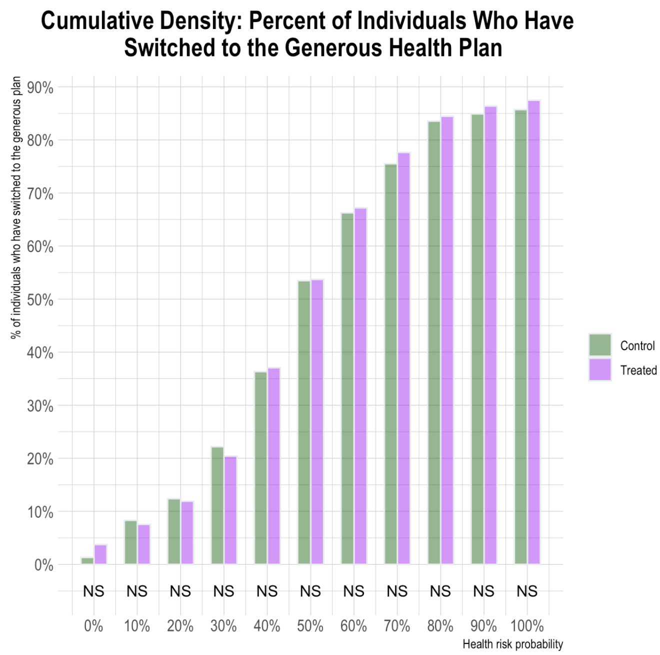

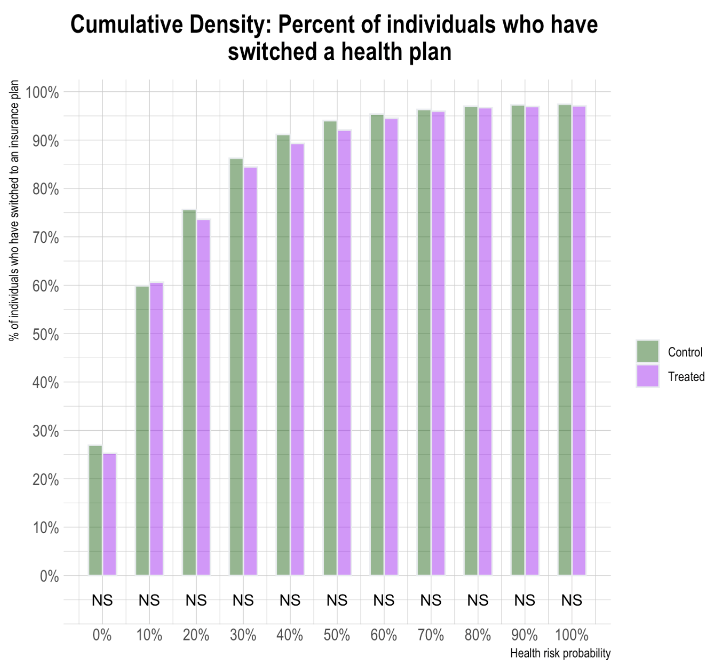

Note: I report statistical significance from running separate regressions comparing the % of individuals in each relevant comparison group (treatment versus control, avoiders versus non-avoiders) that have switched to the generous plan by the time their health risk probability has reached X%. The figures report coefficients from the following regression: $\mathbb{1}(Switch\, before\, X\%)_i = \beta$ Treatment$_i + \epsilon_i$

It is potentially surprising that respondents do not respond to salient health images

(Treatment two). This finding stands in contrast to recent research that finds that providing respondents with graphic images in certain contexts results in significant behavioral changes (Ahn and Mermin-Bunnell, 2022). Future research ought to explore alternative forms of information to study whether avoiders, in particular, can be incentivized to select more generous plans with the provision health information.

Table 1.5: Pooled Treatment Effect on Health Risk Switch Point for Health Insurance

	Switch to generous plan		Switch to plan	
	(1)	(2)	(3)	(4)
Treatment	.95	.58	.91	.57
	(1.19)	(1.35)	(.89)	(1.02)
Avoid	3.14**	2.35	2.11**	1.42
	(1.41)	(1.41)	(1.04)	(1.44)
Treatment*Avoid		1.64		1.45
		(2.83)		(2.09)

Note: Columns 1 and 3 report treatment effects without including an interaction between *Treatment* and *Avoid*. Columns 1 and 2 correspond to treatment effects for primary outcome 1 (the health risk probability when an individual switches to the generous insurance plan), and 3 and 4 are for outcome 2 (the health risk probability when an individual switches to an insurance plan).
*, **, and *** denote significance at the 10, 5, and 1 percent level.

Treatment may have no effect on people's choices because information on the consequences of health issues is insufficient in motivating individuals to switch their health plan choices. As the conceptual model points out (Appendix A.5), a likely alternative may be that individuals primarily rely on information about their individual health risks, such that general information about averages is not highly relevant to them.[13] Or, it is possible that more impactful information would provide respondents specific scenarios on how they can utilize their plans for common medical issues, as opposed to this information, which emphasized health and mortality consequences of "more extreme" issues coinciding with inadequate insurance. Aside from the lack of a treatment effect for all individuals, perhaps the reason there was no effect for the group of avoiders, specifically, is that they expected significantly more upsetting health information provided in treatment.

[13]Avoiders and non-avoiders have at least an 80% accuracy on a post-quiz about the health information provided to treated respondents (Table A.5). If anything, this suggests that individuals internalized health consequence information, yet *still*, this did not impact choices. Perhaps, avoiders (as well as non-avoiders) would react more to information about the individual risks of actually experiencing a health issue in the health simulation.

1.5. Conclusion

In the United States in 2020, 31.2 million adults under 65 years of age were uninsured for medical care. Uninsured adults are more likely than insured adults to have poor or fair health (CDC, 2020). Individuals who potentially have the most to gain in terms of life quality as well as returns to longevity are also those who are least likely to hold a health plan.

This paper sheds light on the role of information avoidance, a previously unexplored facet underlying preferences for health plan coverage. My experiment, fielded to adult respondents across the United States, allows me to plausibly identify individuals who have a propensity to avoid salient health information, which existing research previously could not measure by relying on observational data alone. I document that a significant portion of individuals (24%) are willing to pay to avoid information about health. A larger proportion of respondents (34%) are willing to pay to avoid health information when they would have to forsake $0.05 in potential bonus ("low bonus"). Even when individuals would have to forsake the highest possible bonus amount of $0.40, 17% of respondents are still willing to pay to avoid health information. Under the "low bonus spread" condition, the group of individuals classified as avoiders displays significantly different plan preferences.

I offer new descriptive evidence that avoiders appear to select riskier health plans. After controlling for treatment assignment, avoiders wait until their health risks are 2 pp higher before switching to an insurance plan and 3 pp higher before switching to a "generous plan" that has a relatively low deductible and high premium. This paper also contributes by documenting characteristics about avoiders: they are more often older, earn less, know fewer acquaintances hospitalized in the past year, are less likely to have post-college education, are less likely to report performing calculations, and have lower initial insurance term knowledge. Even after controlling for demographic and pre-treatment attributes, I find that avoidance still predicts individuals selecting less generous plan choices.

When I randomly assign individuals to view a form of salient health information containing information about overall health risks, mortality, plan treatment benefits, as well as salient health images, I observe no treatment effect on either avoiders' or non-avoiders' plan choices. Potentially, the lack of a treatment effect specifically for the group of avoiders is slightly more important to delve into because of their riskier plan choices. It may be of particular policy importance for subsequent research to focus on exploring effective incentives for the group of avoiders to fully internalize information, which in turn could address the observed positive relationship between avoidance and risk. One possibility is that avoiders may simply inherently prefer less generous plans. Or, an explanation could be that an unobserved variable correlated with both avoidance and plan choice explains the insurance choice patterns observed across avoiders and non-avoiders. This study already documents and accounts for numerous behavioral and observable attributes about avoiders. It may be useful to also consider the role of alternative cognitive or behavioral explanations that may relate to information avoidance, such as procrastination or time-varying preferences based on readily available information in individuals' memories.

A different explanation could be that avoiders may be sensitive to alternative health information not provided in this study. The treatment I explore in this study is but one form of potentially many types of necessary policy-relevant information interventions that may help in nudging individuals, and particularly avoiders, towards more generous plans. It is possible that avoiders struggle with attention (as evidenced by their lower quiz score about information provided in the health information). Future research designed with the aim of eliciting a treatment response for avoiders, particularly, may focus on a longer treatment duration or a more interactive information intervention, such as repeated interventions to re-enforce health facts and images. Avoiders may specifically lack information about their own health risks. Research may explore the role of targeted health risk and consequence information that is tailored to an individuals' specific risks. Future experimental studies may consider field interventions to measure individuals' actual insurance choices in response to a treatment deployed through repeated, personalized information in the form of a longitudinal study.

This paper provides primary evidence that avoidance is likely to play a role in explaining why individuals may select no health insurance and high deductible plans instead of more generous plans. Furthermore, the paper suggests that nudges from informational interventions may not be enough in this context to change individuals' choices. Instead, we need further research to shed light on what would encourage avoiders to attend to information that would change their insurance choices. With these considerations in mind, this paper hopefully clarifies an additional piece of the complex and often interconnected puzzle relating to why individuals may underinsure for medical care and offers a few possible avenues for future exploration.

CHAPTER II

Private Health Insurance Patterns Following Spousal Health Shocks

2.0 Abstract

Using an event study, this paper explores how an individual's health insurance coverage changes after their spouse is hospitalized. I examine longitudinal data spanning 1992 to 2018 from the Health and Retirement Study, which includes insurance status and hospital admissions for individuals and their spouses. For the sample of individuals who did not hold a private health plan in the first survey wave, those whose spouses are hospitalized are 7.4pp more likely to hold a private health plan in the survey wave following the spousal health shock. Furthermore, for all individuals who either newly signed up for private insurance or who already held private insurance, the average overall total annual private insurance plan premium increases by $82 as a result of the spousal hospitalization event, representing a 7.3% increase over their average private plan annual premium. Increased plan coverage persists for multiple years following the spousal hospitalization event. I further identify large initial and persistent take-up in private plan coverage for white respondents, individuals with higher initial asset holdings, and individuals whose spouses held insurance at the time of their own hospitalization. In contrast, I observe a relatively small, short-term increased probability in private plan coverage for those with lower asset holdings, no detectable increase for minorities, and a delayed response for individuals whose spouse did not hold a plan at the time of their hospital admission. These findings suggest that there are more complicated and widespread impacts of hospital admissions than than the existing literature, which focuses on additional financial burden on hospitalized individuals, may suggest. These findings may motivate future research to explore behavioral selection into health insurance in response to particular forms of salient information.

2.1. Introduction

It is well-established that individuals with health insurance are less likely to forgo needed medical care and are more likely to report excellent or very good health (Institute of Medicine, 2001). Yet, 27.5 million individuals under age 65 were uninsured in 2021 (Tolbert et al., 2022). It is consequently of significant research and policy interest to analyze events that lead to greater plan coverage among the population, as this could shed light on effective informational interventions and specific populations that are responsive. It is particularly challenging to study this question without relying on comprehensive surveys of populations who "recently took up" private insurance.[1] However, longitudinal data may help researchers deduce certain patterns. For instance, an understudied phenomenon has been whether individuals' exposure to health shocks that are not their own impacts plan coverage. This is a compelling hypothesis because health shocks are common and individuals are regularly exposed, second-hand: in 2018, there were 130 million emergency department visits in the United States and 12.4% of these visits resulted in a hospital admission (Cairns et al., 2021).

A silver lining of these health shocks might be that they have spillover impacts on other people who witness them, resulting in shifts in insurance coverage for the initially uninsured or those who hold health plans associated with high deductibles. I build on Dobkin et al. (2018), who documents that hospital admissions result in unpaid medical bills, bankruptcies, and negative economic impacts on earnings and access to credit.

I use longitudinal data from the Health and Retirement Study between 1992 and 2018 to study health plan coverage leading up to and following a spousal hospital admission. I use a similar event study approach to that carried out by Dobkin et al. (2018), with the central difference that I explore the impact of an individuals' *spouse's* hospital admission, a plausibly exogenous shock, on an individuals own private plan coverage. Observing spousal health shocks may directly influence individuals to sign-on for health insurance as they learn about health risks and and update their own priors from observing a salient health event. There are, however, alternative reasons why a spouse's hospital admission may play a role in greater health insurance take-up, such as lowered sign-on costs or financial incentives to smooth spending for couples and families.

[1] Existing literature has documented important aspects about the role that health shocks have on direct outcomes for individuals who were hospitalized: for instance, people who previously faced health shocks adjust their marginal utility from consumption (Wagstaff, 2005). Research has focused on evaluating policies with an understanding of how to incentivize individuals to take out plans that suit them, for instance by exploring the role of paternalistic defaults (Brot-Goldberg et al., 2021). Further research has focused on providing individuals with information to educate individuals about plan characteristics (Bhargava et al., 2017). This paper contributes by identifying that individuals are likely to shift towards more generous health plans after spousal health shocks and discusses potential explanations for why this appears to happen.

I specifically focus on selection into private plans in this study for two reasons. First, the majority of those who hold insurance and who fall under the age of 65 primarily rely on private health insurance, which as of 2021 covers 68.4% of the United States population, including children under 18 and some individuals above the age of 65 (Congressional Research Service, 2023). Second, the majority of individuals may only sign on to public plans if they meet certain criteria. For instance, individuals are permitted to sign on as a Medicare beneficiary once they reach the age of 65, with exceptions for individuals with disabilities and certain conditions. Furthermore, once eligibility criteria are met, there is already high participation in public programs; over 90% of individuals aged 65 and older participate in Medicare (Geifer and King, 2021). Further, Card et al. (2008) suggests that nearly universal coverage under Medicare after the age of 65 lowers many health and medical utilization outcomes between minorities and non-minorities due to increased coverage, but persistent gaps remain between certain groups for hospitalizations, with one potential explanation that Medicare is valuable *in addition* to other (i.e. private) insurance plans in lowering prohibitive costs. This suggests that private health insurance take-up is important to explore, even independent of public plan adoption.

This paper contributes by providing primary evidence that exposure to a health shock, from a spousal hospital admission, results in increased private plan coverage for respondents. This paper also analyzes the potential explanations underlying this increase in health coverage by describing the extent to which increased coverage may be due to an individual's own behavioral shift due to increased salience or due to information costs becoming lower. This paper also sheds additional light on long-term insurance patterns for individuals who have a "propensity" to remain uninsured. Particularly, I examine health plan take-up and changes in total annual premium payments for a group of "initially uninsured" individuals (i.e. those who entered the HRS without a *private* health plan). I separately explore these outcomes on a larger sample consisting of all individuals.[2]

From the main event study approach, I find that in the wave corresponding to the first observable time that a person's spouse was admitted to the hospital (primary spousal hospital admission), private health coverage increases by 7.4pp among initially uninsured respondents. I also find that exposure to a spousal hospitalization event results in an average increase in the total annual premium by $82, for the combined group of individuals who already held

[2]I make no distinction between the elderly (those 65 years and older) and near-elderly in my study, although in my empirical analysis, I control for age. It is evident that the impact I would observe in this study would be larger for the near-elderly population, or those directly below the age of 65 years, since Medicare offers near universal coverage for the elderly and less than 1% of individuals above 65 years are uninsured (Tolbert et al., 2022). I provide additional empirical evidence by censoring the panel above and below the age of 65 years and present the results in Appendix Figure B.2.

insurance (and were already paying a premium) and those who newly selected into a private plan (who were not previously paying a premium). This provides suggestive evidence that respondents are insured under plans with more generous coverage, as a result of the spousal hospitalization event. These effects are persistent over time: private plan uptake among the treated, initially uninsured population remains elevated for three survey waves following the spousal hospitalization and increased annual premiums remain elevated for four survey waves for all individuals.

I decompose the increased take-up in private insurance between individuals who are insured under at least one plan where the spouse is a policy-holder and those who are insured entirely by plans for which they are policy-holders. I find that 60% of increase in the take-up of private plans is driven by individuals privately insuring as the main policy-holder, instead of from a spouse's employer or marketplace plan. This finding is consistent with survey respondents responding to novel information provided in the form of the increased salience of health risks. The decline in magnitude of my point estimates over time is also consistent with people responding immediately to a health shock because of heightened initial salience.

However, there are other possible explanations that may underlie the increase in private plan coverage I observe for treated respondents. One such alternative is that the hospitalization event may have lowered transaction costs associated with signing up for private insurance (for instance, information about the plan sign-up process becomes easier to access). To explore this empirically, I assume an underlying negative connection between the transaction costs an individual faces in signing up for private health insurance plans and whether their spouse *already held* a form of private or government-sourced health insurance at the time of their hospitalization event. Using my event study approach, I divide the sample of treated individuals based on this spousal characteristic to find separate treatment effects for each group of respondents. I find that respondents whose spouses held an insurance plan during their hospitalization event experience a 9.6pp increase, in the survey wave corresponding to the spousal hospitalization, in private plan take-up as a result of the health shock. In contrast, there is no immediate significant impact of the spousal hospitalization for individuals whose spouses were uninsured at the time of their hospitalization. Instead, this group of individuals experiences a 10.5pp increase ($P < 0.05$) in private plan possession after five waves. This finding suggests that those who faced higher transaction costs took longer to respond to the spousal hospitalization than those who did not. Part of the long-run significant effect for individuals whose spouse did not hold insurance at their hospitalization event may be mechanical; individuals in this group may have anyways signed on for private plans as they age. In addition to transaction costs, I provide evidence that there could be persistent barriers in obtaining private plans for less wealthy (measured by liquid asset holdings) individuals and minority individuals in plan take-up: I observe a range of no effects to relatively small effects that taper off relatively quickly.

The remainder of the paper proceeds as follows. Section 2.2 provides a descriptive overview about the longitudinal HRS dataset and characterizes individuals who experienced spousal hospitalization events. Section 2.3 presents evidence using an event study that individuals respond to spousal health shocks by taking out additional insurance coverage. Section 2.4 examines heterogeneity in the treatment effect of spousal hospitalizations across different groups of individuals. Section 2.5 concludes.

2.2. Data and Descriptive Summary

I use longitudinal data from the Health and Retirement Study (HRS), which is nationally representative of the elderly and near-elderly population. There is linked data with information about respondents and their spouses, which allows me to explore how individuals' insurance choices evolve after a spousal hospitalization event. A spousal hospitalization refers to whether a survey respondent's spouse reported any overnight stay in a hospital in the reference period. In the HRS, this period is all time since the last interview, or the past two years up to the current interview wave w. I utilize 14 biannual survey waves that span the years 1992 and 2018. I limit the sample to include individuals whom I can observe in at least one interview wave without reporting a spousal hospitalization event. To do so, I drop individuals whose spouse was hospitalized in the first wave they appeared in the HRS. This restriction ensures that the index spousal hospitalization event occurs at least three years after any previous hospitalizations, in expectation.

Table 2.1 presents sample characteristics for individuals who experienced at least one spousal hospitalization event (within the relevant study window) and for those that never observed one. There are 35,248 unique individuals in the full sample. In the overall sample, individuals whose spouses were hospitalized tend to be older and hold significantly more assets. Individuals exposed to spousal hospitalization events are also more likely to hold insurance themselves and less likely to be uninsured in their first survey wave. There are 11,144 individuals who were uninsured in their initial HRS interview wave. The descriptive summary of the sample of respondents who entered uninsured into the survey is in Appendix Table B.1.

Table 2.1: Average Sample Characteristics

Characteristic	(1) Spouse Ever Hospitalized	(2) Spouse Never Hospitalized	(3) Difference
Continuous Variables			
Age	66.31	64.75	1.55***
	(10.68)	(11.53)	(0.05)
Assets (Checking + Savings)	\$31,154	\$22,217	\$8,936***
	(140,876)	(161,051)	(650)
Total Annual Private Plan Premium	\$1,221	\$1,049	172.50***
	(2,663)	(2,617)	(12.29)
Binary Variables			
% Above 65	53%	43%	10.83***
			(0.21)
% White (not Hispanic)	77%	62%	14.98***
			(0.21)
Self-hospitalized	24%	25%	−0.36**
			(0.18)
Individual Holds Private Plan	66%	58%	8.00***
			(0.21)
Individual Covered by Spouse	16.87%	10.60%	6.3***
			(0.15)
Individual Uninsured First Wave	19.90%	32.47%	−12.57***
			(0.19)
Number of Unique Respondents	11,398	23,850	35,248

Notes: This table presents averages across individual survey wave pairs, allowing for the same individual to be repeated up to 14 waves, using the universe of data between 1992 to 2018. The proportions represent proportions of outcomes across all of the survey waves represented by all individuals. Note that the sample reflects (a) dropping individual-wave pairs with missing weights and (b) dropping individuals whose spouses were hospitalized in the first wave that an individual appeared in the survey. Column 3 presents estimates and standard errors for a t-test of similarity between Column 1 and Column 2. I include standard errors for continuous (non-binary) variables in parentheses below their respective means. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.10$. *Data Source:* Health and Retirement Study.

It is clear from Table 2.1 that individuals whose spouses were ever hospitalized descriptively differ on observable (and likely on unobservable) characteristics. The absence of randomly-assigned exposure to spousal health shocks necessitates the use of quasi-experimental methods to explore the causal impact of a spousal hospitalization event on health insurance outcomes.

2.2.1 Outcomes

I focus on two main outcomes of private health insurance take-up. The first outcome is whether an individual holds at least one private health insurance plan. For this outcome measure, as well as the second outcome measure below, I consider a respondent to be insured if they are insured through a private plan held either by themselves or insured through a private plan their spouse holds that also covers the respondent.[3] The percent of private plan holders across all waves is lower, on average, for individuals who entered the HRS without a private plan (henceforth, for simplicity, I refer to this group as the "initially uninsured") as opposed to all individuals. On average, 19% of the "initially uninsured" holds at least one private plan across all survey waves, while 62% of the overall sample holds at least one private plan.

While the measure of whether an individual holds at least one private plan provides insight into plan uptake at the extensive margin, it is also useful to deduce whether the quality of plan selection changes in response to a spousal health shock. More generous health plans are often associated with higher premiums. The second main outcome variable I examine is a measure of the total annual premium across a respondent's private plans (including zeros, which represent an individuals who do not hold a plan).[4] In all survey waves, if a respondent indicates that they pay all or some of the cost of the premium (or if there is no initial question regarding how much of the premium the respondent pays), the respondent is asked how much they pay to contribute to plan p. Consequently if individual i and this individual's spouse, i_s are both covered under plan p (i.e. through either of their employers), both person i and i_s may pay different portions of the premium associated with plan p. However, my second constructed outcome measure only reflects how much individual i contributes towards private

[3]The HRS provides an exact count for the number of private health plans an individual holds in a given survey wave; I utilize this to form my outcome variable to be a binary indicator as to whether a respondent holds *at least one* private health plan.

[4]The HRS provides information on the monthly premium associated with the top three private plans that an individual is a beneficiary under. To impute the second outcome ("total annual premium"), I sum up the premiums up across each respondents' three highest premium private plans and multiply the sum by 12. I apply an additional transformation to set "total annual premium" equal to zero for individuals who did not hold any private plans in the survey wave (as opposed to treating these outcomes as missing).

plans.[5] The average annual premium for the "initially uninsured" sample is \$297 and the average annual premium for the full sample is \$1,126.

I test main outcomes (a) private plan take-up and (b) total annual premium on two different samples: (1) the sample of "initially uninsured" respondents (individuals who did not hold any private plan their first HRS survey wave) and (2) the "full sample" of all respondents. However, given the relatively high mechanical level of private plan ownership among the "full sample", the first main outcome (private plan take-up) is particularly meaningful to interpret for the sample of "initially uninsured" respondents. During the remainder of the paper, I note the specific sample I test a particular outcome on.

2.2.2 Definition of "Treatment"

I benchmark treatment relative to the first survey wave that an individual experiences a spousal hospitalization in the past two years at baseline in the relevant survey window. I use this baseline measure of a "primary" (first observed) spousal hospital exposure to form time-varying leads and lags, for "ever-treated" individuals.

Following the recent literature in TWFE (Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021), it is important to note that treatment units in an unrestricted, standard event study design may act as effective "control" units for either prior to or after they have been assigned to treatment. Thus, in this paper, I go only so far as to make the distinction between "ever-treated" and "never-treated" units. I interpret treatment effects with the context that the effective comparison group in each period relative to treatment may contain some "ever" and "never" treated individuals. However, I show in Appendix Sections B.3.2 and B.3.3 that results are robust to stricter comparisons (a) using only a sample of "ever-treated" units and (b) using a specifically defined control group of "never-treated" units in a newer event study approach that is robust to potential treatment effect heterogeneity.

2.2.3 Covariates and Additional HRS Variables

I gather an individual's age to include as a control variable, which varies contemporaneously across successive time periods in the constructed panel. I gather self-hospitalization status to explore the impacts of a spousal hospitalization event on an individuals' own medical utilization in response to treatment. For subgroup analyses, I use data on race/ethnicity to

[5] For someone who is Medicare-eligible (i.e. above 65 years, for instance), it is possible to sign on for Medicare Part D and to still hold a private health plan either through a current/former employer, union, or from the marketplace. If this is the case, the employer or union will notify an individual if their drug coverage is creditable. Importantly, individuals who sign up for Medicare Part D are advised that they *may* lose their employer or union health drug coverage once they sign up for Medicare drug coverage (Source: Medicare Government Online).

divide the sample between Black, Hispanic, and white respondents. To provide a measure of liquid assets, I use an HRS variable that sums a respondent and spouse's total value in their checking, savings, or money market accounts.[6] This proxies for an individual's ability to pay private insurance premiums. I also use a variable on the years a respondent has been married to their current spouse in each wave to determine marriage length around the time of spousal hospitalization. Lastly, I use a variable detailing whether respondent's spousal plan status, which allows me to determine whether a respondent's spouse held either a government or a private plan at the time of their hospitalization.

2.3. Event Study of Spousal Hospitalization Events

I present an event study to examine how insurance decisions change based on time since a spousal hospitalization. In my preferred specification, I denote $Event_i$ to represent a variable recording the time period (in "absolute" survey wave terms, which translates to biennial calendar years) t when individual i's spouse is hospitalized. I analyze the following specification[7]:

$$y_{it} = \sum_{w=-3}^{-2} \mu_w(\mathbb{1}[t = Event_i - w])_{it} + \sum_{w=0}^{5} \mu_w(\mathbb{1}[t = Event_i + w])_{it} + \beta X_{it} + \gamma_i + \tau_t + \varepsilon_{it} \quad (2.3.1)$$

where μ_w are coefficients on indicators for survey waves relative to the wave w that an individuals' spouse was hospitalized in the observable survey window. I define a relevant survey window to include up to five waves following the spousal hospitalization event, which corresponds to, on average, 10 years post-spousal hospitalization. X_{it} is a vector of individual-level controls that includes a respondent's age, an indicator for whether individual i is above 65 years (the age when the majority of individuals become eligible for Medicare), and

[6] For analyses, I transform the HRS variable using the natural logarithm. To avoid dropping individuals who have zero assets I add $1 to every individual's asset holdings.

[7] I include never-treated units in my main specification as my, for whom μ_w equals zero in all treatment periods. However, in Section 2.3.2.1, I present evidence that there are similar results even after only including ever-treated individuals. I do not trim my sample beyond the fifth wave following the spousal hospitalization or prior to three waves leading up to the spousal hospitalization event. However, I find that even after trimming in this manner, the results are consistent with my results discussed below in Section 2.3.4.

an interaction of this indicator with respondent i's age.[8] γ_i represents individual fixed effects[9] and τ_t represents biennial calendar fixed effects. The coefficients of interest are the μ_w's, which estimate the outcome variables at a given wave w relative to the omitted indicator, μ_{-1}. The omitted indicator, μ_{-1} reflects (on average) an interview conducted one year prior to the spousal hospitalization. I cluster standard errors at the individual level. Equation 2.3.1 is fitted using weighted least squares using person-level analysis weights. I provide a detailed explanation of the sample construction and empirical setup in Appendix B.1.

2.3.1 Panel Structure and Balance

By construction, the HRS is a longitudinal survey with staggered entry and attrition in different survey waves. For the purpose of preserving power, my baseline empirical specification does not rely on a balanced sample for inference. Yet, one concern due to not restricting to a balanced panel could be potential bias arising from survey attrition, particularly following the time of treatment. The inclusion of individual fixed effects in all of my specifications is one approach to address this concern. However, I also run my empirical specification by restricting the sample of "initially uninsured" respondents to a balanced panel of "ever-treated" individuals whom I observe in survey waves -2 and 2, which allows me to observe respondents that appear in survey waves around the time of treatment, and which may help address the potential concern that compositional changes contribute to the results. I present the results from this approach in Appendix Table B.3 and find similar results to my baseline specification.

2.3.2 Identifying Assumptions

A feature of the event study design is the presence of multiple cohorts that may be exposed to unexpected spousal hospital admissions in different waves of the HRS. I follow Sun and Abraham (2021) by walking through three identifying assumptions that are relevant for the validity of the event study.

[8]In Appendix B.2, I present results from running the empirical specification separately on individuals above and below the age of 65. In this specification, the same individual i could appear in both an "under 65" panel and an "over 65" panel (in different survey waves), depending on their age in the current survey wave. The same central restrictions still apply to individuals in both panels: for instance, if an individual i's spouse was hospitalized in individual i's first HRS survey wave, they are excluded from *both* the "under 65" and "over 65" panels.

[9]Note that Dobkin et al. (2018) only includes these as a form of robustness; however, inclusion of these affects the point estimates substantially.

2.3.2.1 Assumption 1: Parallel Trends Assumption

The claim that trends would have counterfactually evolved in parallel in the absence of treatment is more plausible if we observe them evolving similarly prior to treatment. One way for researchers to assess this in event study designs is to examine whether there are insignificant pre-trends prior to the treatment period. A theoretical threat to the parallel trends assumption in this context might be that respondents whose spouses were hospitalized fundamentally differed from those whose spouses were never hospitalized, in a way that differentially impacted take-up of private insurance in the absence of treatment. For instance, individuals whose spouses were hospitalized at some point may have had more background exposure to the health and insurance context to begin with due to spousal doctors trips, even absent a more serious spousal hospitalization event. To address this potential issue, I follow Dobkin et al. (2018) and Sun and Abraham (2021) by examining the effects of limiting the sample to only respondents that ever experienced a spousal hospitalization event in the observable window. From this robustness check, presented in Figure B.4, I find that the results are highly similar to my main results.

2.3.2.2 Assumption 2: No Behavior Based on Anticipation

As Sun and Abraham (2021) explain about the setting underlying hospitalization events, given that these are significant, and often unusual events that occur (in my interpretation) as "health shocks", it is plausible to assume that individuals could not have anticipated that their spouse would be hospitalized. Furthermore, if anything, we may expect that individuals behaviorally pre-empt their *own* future hospitalization events by taking up insurance, but this seems less plausible in my setting where I am exploring a spouse's insurance take-up. As Borusyak et al. (2023) explains, some form of the no anticipation effects assumption is necessary, "*as there would be no reference periods for treated units otherwise" (p. 9).*

2.3.2.3 Assumption 3: Treatment Effect Homogeneity

As Sun and Abraham (2021) articulate: "*in settings with variation in treatment timing across units, the coefficient on a given lead or lag can be contaminated by effects from other periods, and apparent pretrends can arise solely from treatment effects heterogeneity" (Abstract).* Since respondents in my sample can observe primary spousal hospitalizations in different survey waves, there is a possibility for biased estimates due to treatment effect heterogeneity. For example, as Sun and Abraham (2021) point out about the HRS, individuals who are hospitalized in later waves are, as a matter-of-fact, more likely to be older – and in the case of this study, it is plausible that respondents whose spouses are hospitalized in later waves are *also* likely to be older, considering assortative marriage by age. There are also several other variables that may generate treatment effect heterogeneity in this context, lowering the cost of signing on to health plans for later cohorts, including evolving policies (i.e. the Affordable Care Act) and incentives (i.e. government subsidized plans) or a gradual transition during the survey period towards online marketplaces that might facilitate an easier plan sign-up process. For the purpose of this paper, I assume that treatment effects are relatively homogenous (a) for different treatment cohorts and (b) across different treatment event lags within each treatment cohort.

If this assumption were to be substantively incorrect, the point estimates corresponding to event leads or lags in my results could be biased by picking up terms consisting of treatment effects from other periods. Further, the pretrends presented in my results could be different, if this assumption is violated. However, I re-run the event study using an alternative estimation technique, the interaction-weighted estimator, proposed by Sun and Abraham (2021). I discuss my implementation of this technique and the results for one of my main outcome variables (private health plan take-up) in Appendix Section B.3.3. The results from this robustness are highly similar to the results from the main event study approach. This demonstrates that even if my assumption of treatment effect homogeneity were violated, I would still observe that private health plan take-up increases as a result of spousal hospitalization events.

2.3.3 Descriptive Evidence

I motivate my main findings from the event study approach with descriptive evidence of a shift in the percent of individuals who hold a private plan following a spousal hospital admission. Figure 2.1 presents evidence of a clearly observable jump from 20% to 30% in the probability that the *initially uninsured* hold a private plan, following a first-time spousal hospitalization event.

Figure 2.1: Probability an Initially Uninsured Individual Holds a Private Insurance Plan Before and After a Spousal Hospitalization Event

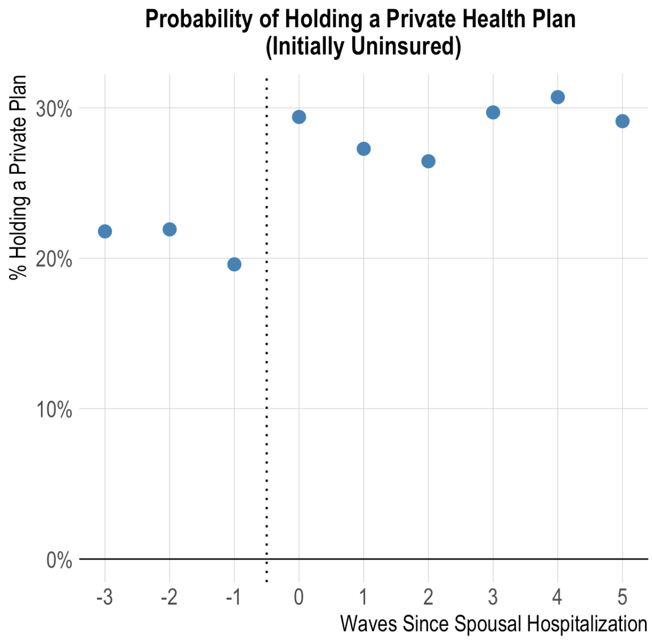

Notes: This figure plots the percent of initially uninsured individuals who hold at least one private health plan. All individuals displayed here entered the HRS in a survey wave where they did not hold a private health insurance plan. The plotted probability for each survey wave represents an average percent of individuals in the sample in that specific wave that hold at least 1 private plan (measured in percent, not percentage points). In other words, the composition of individuals changes across waves depending on the wave they entered the HRS for the first time and the wave they exited the HRS, so this does not plot the same exact individuals across time. *Data source*: Health and Retirement Study.

What underlies the patterns in Figure 2.1? The event study approach allows me to determine how much of this uptake of private plans is specifically a result of the spousal hospitalization. Then, analyses that explore heterogeneous treatment effects allow me to comment on the degree to which potential explanations for why the spousal hospitalization event may incite increases in private insurance uptake. The remainder of the empirical exploration presents evidence of the plausible treatment effect of spousal hospitalization events on private plan take-up.

2.3.4 Main Results

Figure 2.2 illustrates the results corresponding to equation 2.3.1. For "initially uninsured" individuals, private health insurance plan ownership increases by 7.4pp in the survey wave corresponding to the spousal hospitalization event and this elevated probability persists for several survey waves. There is a statistically insignificant impact of the spousal hospitalization event on holding private plans when I include *all* individuals in the sample, as opposed to limiting only to individuals who were uninsured in the first wave. By definition, we would expect that the average rate of holding insurance for individuals who already *held* health insurance would either stay constant or decrease over time.

Figure 2.2: Private Insurance Plan Take-up

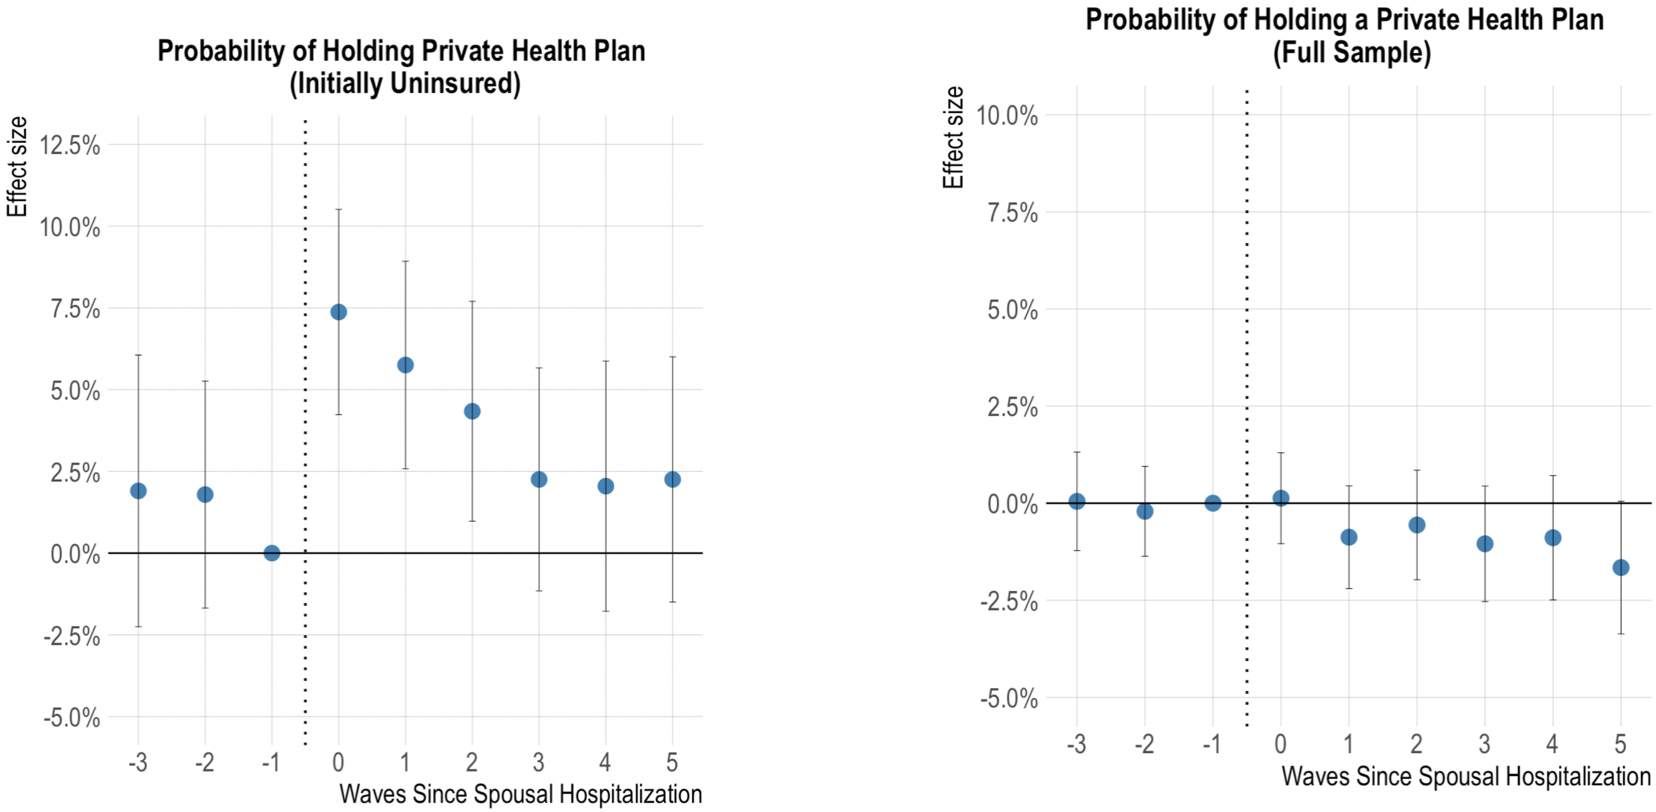

(a) “Initially Uninsured” Sample

(b) Full Sample

Notes: This graph presents estimates using equation 2.3.1. The relevant sample compositions (i.e. the number of treatment and control respondents) are listed in Table B.2. Panel 2.2a presents results for the “initially uninsured” sample and Panel 2.2b presents results for the full sample. On average, the proportion of the sample of “initially uninsured” respondents who hold a private health plan during any survey wave is 19%. The index spousal hospitalization takes place when the “waves since spousal hospitalization” is 0. I include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. *Data source:* Health and Retirement Study.

Figure 2.3 presents results on the total annual plan premium, for equation 2.3.1. The total premium across all private plans increases by \$82 (full sample) in the wave corresponding to the spousal hospitalization event, once again persisting over time. For the sample of "initially uninsured" respondents, there is an increase in the total premium by \$135 in the wave corresponding to treatment. These results suggest that there are two response mechanisms: (1) people are taking up insurance more frequently, and (2) people are taking out even more generous plans following the spousal hospitalization.

Figure 2.3: Private Insurance Total Premium

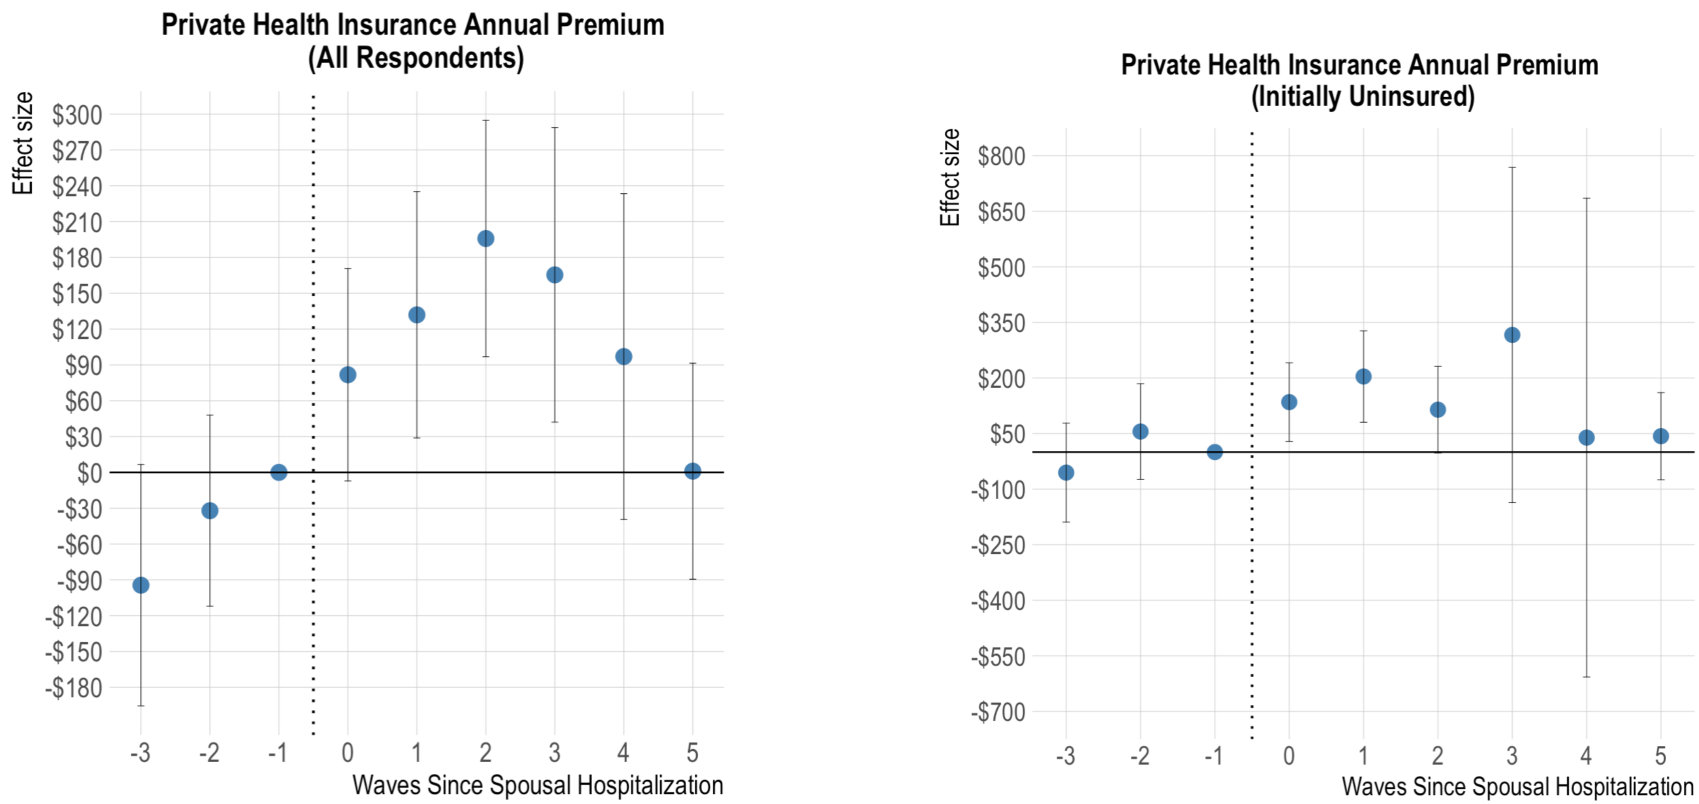

(a) Full Sample

(b) “Initially Uninsured” Sample

Notes: This graph shows estimates using equation 2.3.1. The relevant sample compositions (i.e. the number of treatment and control respondents) are listed in Table B.2. Panel 2.3a presents results for the full sample and Panel 2.3b presents results for the “initially uninsured” sample. The mean total annual premium of those who hold private plans is $1,126. The index spousal hospitalization takes place when the “waves since spousal hospitalization” is 0. I include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. *Data source:* Health and Retirement Study.

I decompose the effect of increased private plan take-up (for the "initially uninsured" sample) between plans through employers and private plans directly from an insurance company, and I present the results in Figure 2.4. There are positive effects of the spousal hospitalization event on the probability individuals insure through both plan sources. The effect size is initially relatively moderate for private plan coverage directly from an insurance company through marketplace (1.9pp increase in uptake in the initial wave, $P < 0.10$) and lasts only up to the second survey wave following the spousal hospitalization. The effect size is initially statistically insignificant for private plans taken out from employers. Yet, two waves after the spousal shock, there is an increased 2.4pp ($P < 0.10$) take-up of plans directly from the insurance company.

Figure 2.4 also presents a decomposition of plan source based on whether at least one of a respondent's private plans was covered by their spouse or whether all private plans are from a respondent's own sources. In the survey wave corresponding to the spousal health shock, uptake in private plans for treated individuals who are covered under spouses increases by 2.9pp and uptake in private plans for treated individuals *not* under spousal coverage increases by 4.4pp. Applying unrounded versions of these shares, I calculate that 60% of the initial uptake in private plans as a result of the spousal health shock is *independent* of potential expanded spousal coverage.

Figure 2.4: Private Insurance Plan Take-up, by Source

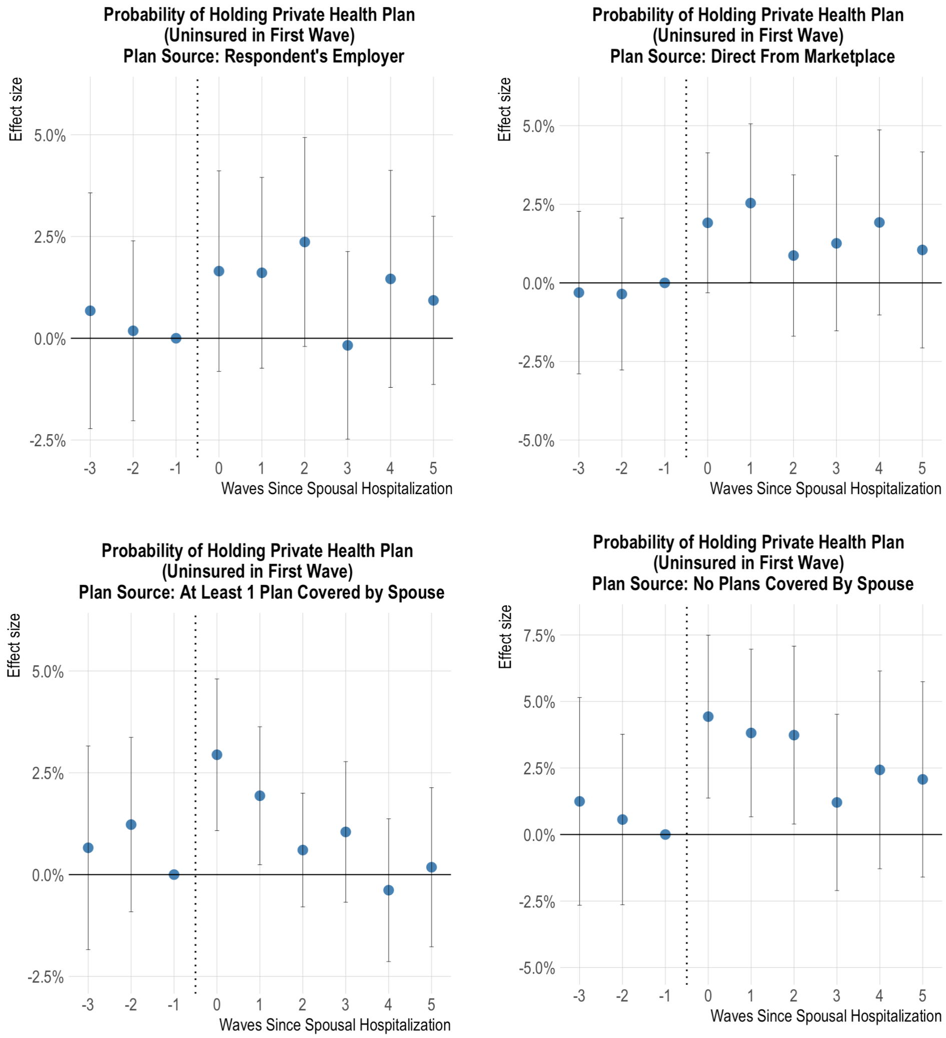

Notes: This graph presents estimates using equation 2.3.1 for the respondents who were initially uninsured in their first survey wave. The outcomes are private plans from based on four sources: (1) the respondent's employer, (2) directly from the health marketplace, (3) covered by the respondent's spouse, and (4) private plan not covered by their spouse (which includes (1) and (2), by definition). The sample is only the "initially uninsured" and the specific composition (i.e. the number of treatment and control respondents) is listed in Table B.2. On average, the proportion of the sample of "initially uninsured" respondents who hold a private health plan during any survey wave is 19%. The index spousal hospitalization takes place when the "waves since spousal hospitalization" is 0. I include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. *Data source:* Health and Retirement Study.

2.3.5 Discussion

I show above that the increase in private health plan take-up (for the initially uninsured) is highest in the survey wave corresponding to the spousal hospitalization event (Figure 2.2). Similarly, the effect size for respondents holding higher premium plans is highest for the first 4 waves after the spousal hospitalization, after which the effects become statistically indistinguishable from zero (Figure 2.3). Why does the effect taper off over time? A plausible explanation for the long-run dissolution in private insurance take-up in the long-run is a saliency effect. Under this explanation, the spousal hospitalization event has its largest impact in the first few years because it is more recent and also more memorable, incentivizing respondents or their spouses to take out more generous plan coverage. As time elapses, the saliency of the hospitalization diminishes, along with increased private plan uptake.

Another potential mechanism underlying long-term trends may be due to a direct decline in liquid assets. This may result in a direct substitution away from private plan take-up in later waves, because health coverage in the first few waves may be more expensive than people previously anticipate. There may also be an indirect effect. One possible explanation follows from Liu et al. (2022), who find that holding medical insurance is positively related to risky financial investments, even while health shocks have an insignificant effect on financial asset allocation. It is apparent that the initial spousal health shock increases private insurance rates, which may be positively correlated with risky financial investments and a consequent long-run deterioration in more liquid assets in checking and savings accounts. Yet, it is also possible that treated individuals begin to form an attachment towards insurance over time. For instance, treated individuals may substitute away from private plans towards government plans.

2.4. Heterogeneity Analyses

I carry out a series of analyses to study heterogeneous responses in health insurance take-up following a spousal hospitalization event. I re-run equation 2.3.1 on specific groupings of HRS respondents and study the pattern of indicators for time relative to the wave prior to the spousal hospitalization event for each group. This inference enables me to deduce a treatment effect of spousal hospitalizations for individuals with different spousal characteristics at the time of their spouses' hospitalization, specific racial/ethnic groups, and individuals with differing initial asset holdings. I discuss the findings for each of these sub-groups in this section.

2.4.1 Spousal Characteristics and Health Insurance Take-up

How do potentially prohibitive transaction costs factor into whether individuals sign onto private plans, particularly on the extensive margin? While there are likely several forms of costs and non-pecuniary barriers associated with signing on to private plans, to consider this question empirically, I assume that individuals whose spouses already held a health plan at the time of their hospital admission already had easier access to information that may make signing up for private plans easier, thereby facing lower transaction costs on average than individuals whose spouses did not hold a plan. I explore heterogeneous outcomes based on this spousal characteristic, in the wave of the spouse's hospital admission.

I explore this hypothesized role of transaction costs by running equation 2.3.1 for separately defined treatment groups: (a) survey respondents whose spouse held *any* form of health insurance in the wave in which they were hospitalized (insured spouse), and (b) survey respondents whose spouse did not hold any health plan (uninsured spouse). As is observable in Figure 2.5a, respondents whose spouses held a plan in the wave they were hospitalized display an immediate change in insurance coverage following the health shock: their probability of private plan take-up increases by 9.6pp ($P < .01$). This elevated probability of being privately insured continues up to three waves post-spousal hospitalization. In contrast, individuals whose spouses *were not* insured at the time of their hospitalization display a more delayed response to the spousal hospitalization. There is no impact for this group in the first few years following the spousal hospitalization event. However, five waves post-spousal hospitalization, private plan membership increases by 10.5pp ($P < .05$). This finding is consistent with the fact that individuals who face higher transaction costs take longer to take on more insurance, and could further be consistent with an alternative explanation that this group would have anyways mechanically signed on for insurance as they age. Access to relevant information may be potentially even more of a differentiating factor at the intensive margin: from Figure 2.5b, it is apparent that individuals whose spouses were insured are seen to take out plans with higher premiums ($97 higher in the year their spouse was hospitalized). Statistically higher premiums persist for this group up to four survey waves following the spousal hospitalization event. In contrast, respondents whose spouses were not insured experience a $441 ($P < 0.01$) decline in their annual premium across private plans in the year of the spousal hospitalization event. A possible explanation is that this group may have chosen to substitute assets away from more generous health insurance plans towards out-of-pocket costs generated from their uninsured spouse's hospitalization stay.

Other spousal characteristics could also shed additional light on the relevance of the health shock for individuals who experience the event. For instance, it is plausible that that length that an individual has been married to their spouse at the time of their hospital admission may correspond with how intensely potential losses seem, which may relate to how salient the health shock is. I explore the role of marriage length by running equation 2.3.1 on different subsamples, formed by defining "ever-treated" groups as: (a) those who were married to their spouse for a relatively short time (20 years or less), (b) a relatively medium length (20 to 40 years), and (c) a relatively long length (more than 40 years). Figure 2.6 provides evidence that the length of time an individual has been married to their spouse at the time of their hospitalization may also play a role for the "initially uninsured" sample. There appears to be an inverse relationship in the length of marriage at the time of the spousal hospitalization and the treatment effect. Respondents in the "short" marriage category demonstrate an increased 12.7pp uptake in private plans, those in the "medium" category demonstrate an increased 9.9pp uptake, and those in the "long" category demonstrate 5.6pp increased uptake in the survey wave corresponding to the spousal health shock. While potentially counterintuitive, this finding could be due to multiple reasons, including increased substitution towards Medicare and public plans for older couples or the possibility that previous health shocks (outside of the observable survey window) may have already had a stronger informational impact on couples married for longer, particularly those married longer than 40 years.

Figure 2.5: Private Health Insurance Take-up, by Spousal Insurance Status

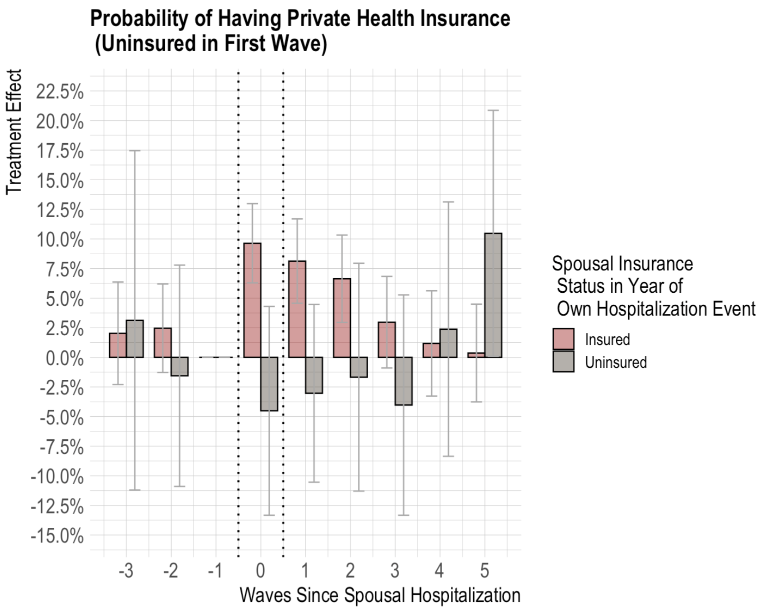

(a) Private plan take-up (Sample: "Initially Uninsured")

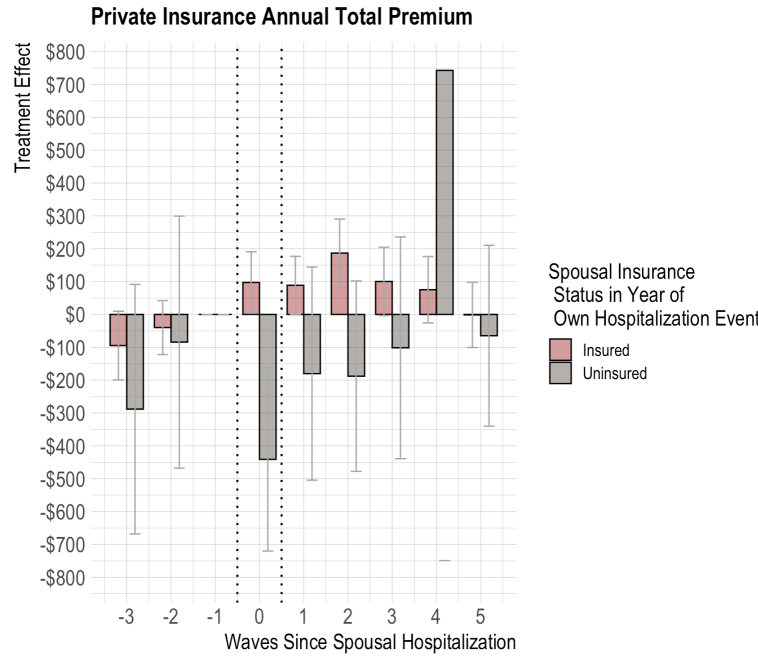

(b) Total premium (Sample: Full Sample)

Notes: I separate *ever treated* individuals (i.e. those who experienced spousal hospitalization events in the the observable survey window) between those whose spouses *held* either a private or government health plan in the year they were hospitalized and those whose spouses did not hold any form of health insurance. Sample compositions (i.e. the number of respondents in each separate dataset) are listed in Table B.2. In separate regressions for equation 2.3.1, each colored bar for each time period compares each of these distinct treatment groups to the (same) group of comparison individuals (i.e. those who did not experience a spousal hospitalization event during their time in the observable survey window). The graphs include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. *Additional note*: The confidence interval is omitted for one of the colored bars in the right panel due to a large standard error. This would have resulted in statistically insignificant results for that particular lag. *Data source:* Health and Retirement Study.

Figure 2.6: Private Health Insurance Take-up, Marriage Length

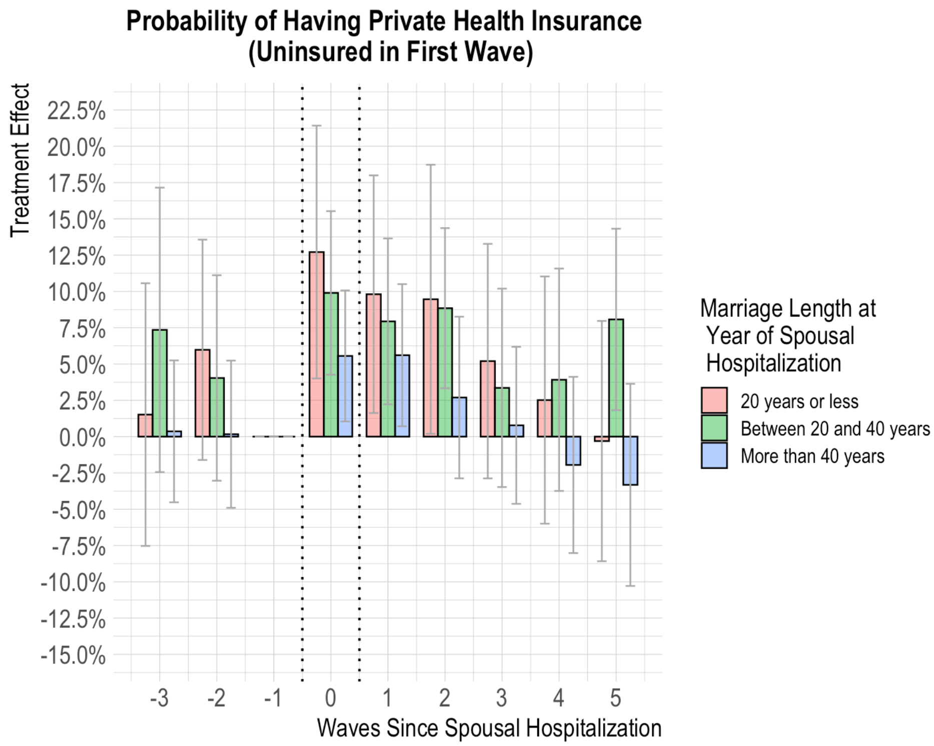

Notes: I separate *treated* individuals (i.e. those who experienced spousal hospitalization events in the the observable survey window) between those who had been married for a relatively short (20 years or less), medium (20 to 40 years) or long (greater than 40 years) time to their spouse at the time of that spouses' hospitalization event. The sample used in this analysis is the "initially uninsured" population, and specific sample compositions (i.e. the number of respondents in each separate dataset) are listed in Table B.2. In separate regressions for equation 2.3.1, each colored bar for each time period compares each of these distinct treatment groups to the (same) group of comparison individuals (i.e. those who did not experience a spousal hospitalization event during their time in the observable survey window. The graphs include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. *Data source:* Health and Retirement Study.

2.4.2 Racial/Ethnic Differences in Health Insurance Take-up

Figure 2.7 presents results from running equation 2.3.1 for three racial/ethnic groups that I determine using participant race and ethnic data in the HRS: non-Hispanic Black, non-Hispanic white, and Hispanic. For the group of "initially uninsured" individuals, take-up of at least one private plan for treated white respondents in the spousal hospitalization wave increases by 10.3pp. This effect falls in magnitude, but remains positive and significant for up to three survey waves following the spousal hospitalization. In comparison, there is no statistically significant positive impact of the spousal health shock on initially uninsured Hispanic or Black respondents in any survey wave following treatment. When interpreting these findings, it is important to take into account the relatively low rate of private health coverage in the survey wave directly preceding the spousal health shock. On average, only 10.1% and 13% of "initially uninsured" Hispanics and Blacks, respectively, hold a private plan in the wave prior to treatment. In contrast, 25.7% of "initially uninsured" white respondents hold a private plan in the wave prior to treatment.

A similar conclusion may be drawn for the total annual premium (for the full sample of respondents): on average, all white respondents who were exposed to a spousal hospitalization event, regardless of previous insurance status, pay $151 more in total premium across private plans than white respondents who are not treated (in the survey wave following treatment). In contrast, there appears to be a *decline* in the total private plan premium for Black respondents one wave after treatment ($230 decline, $P < 0.05$) and for Hispanic respondents five waves after treatment ($583 decline, $P < 0.05$). This is consistent with the literature, which suggests that the African American and Hispanic population has a greater prevalence of trigger events and socioeconomic characteristics associated with greater insurance loss and slower insurance gain (Sohn, 2017). These findings may suggest minorities adopt shorter-term solutions to respond to health shocks. Further, this finding may expand the policy scope to consider persistent barriers that minorities may face in responding to salient health events, which may be less of a hindrance to their white counterparts.

Figure 2.7: Private Health Insurance Take-up, by Racial/Ethnic Group

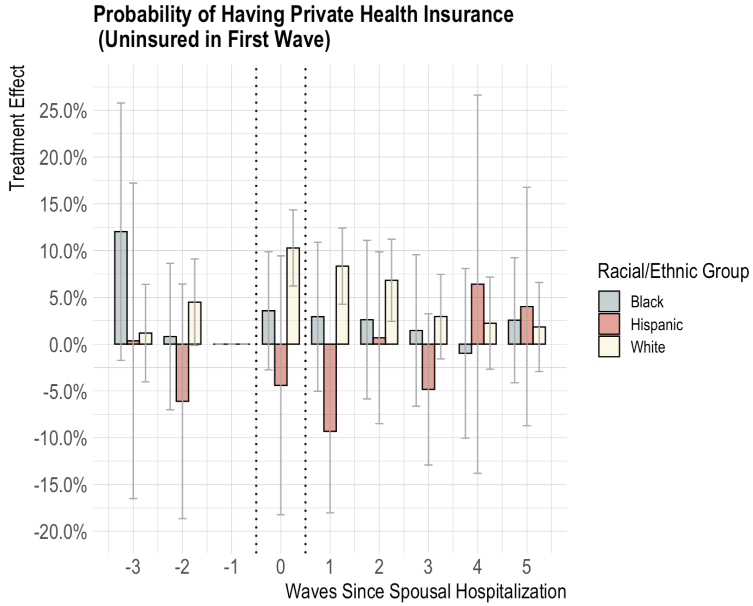

(a) Private plan take-up (Sample: "Initially Uninsured")

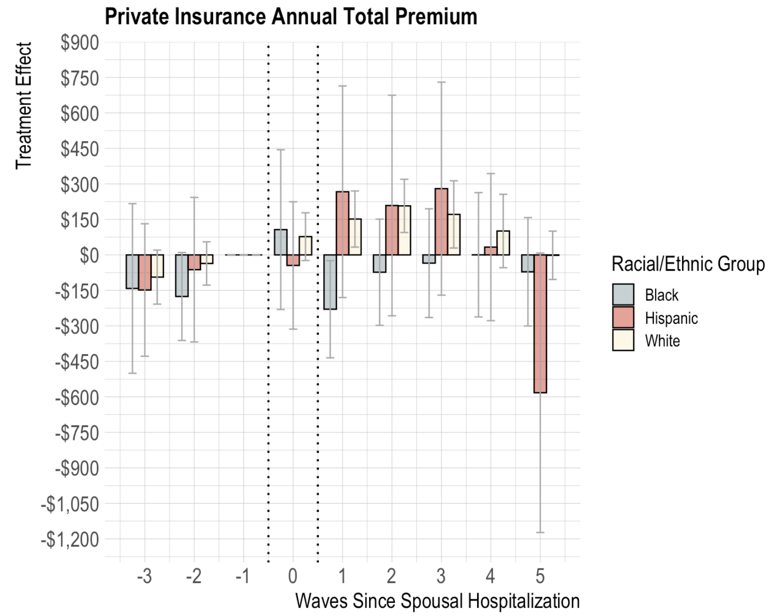

(b) Total premium (Sample: Full Sample)

Notes: I separate Black, Hispanic, and white respondents into distinct panels and run equation 2.3.1 separately for each of these groups. Sample compositions (i.e. the number of respondents in each separate dataset) are listed in Table B.2. Consequently, each colored bar for each time period depicts what happens to insurance coverage for *treated* individuals in Racial/Ethnic group X, relative to *comparison* individuals in Racial/Ethnic group X. The graphs include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. *Data source:* Health and Retirement Study.

2.4.3 Asset Holdings and Health Insurance Take-up

Figure 2.8 presents evidence of the treatment effect on respondent who have relatively "low" and "high" levels of one measure of liquid assets.[10] Figure 2.8 demonstrates that treated, "initially uninsured" individuals with relatively higher liquid assets experience 13.9pp increased uptake in private plans immediately following the spousal hospitalization, while treated respondents with relatively lower liquid assets experience a 3.9pp increase in private plan membership immediately following the spousal hospitalization.

[10]The HRS includes a measure for the amount of money held by *either* the respondent or their spouse in checking, savings, or money market accounts, which I refer to as "liquid assets" in this context because these are resources that respondents can plausibly easily withdraw through their financial institutions.

Figure 2.8: Private Health Insurance Take-up, by Asset Holdings

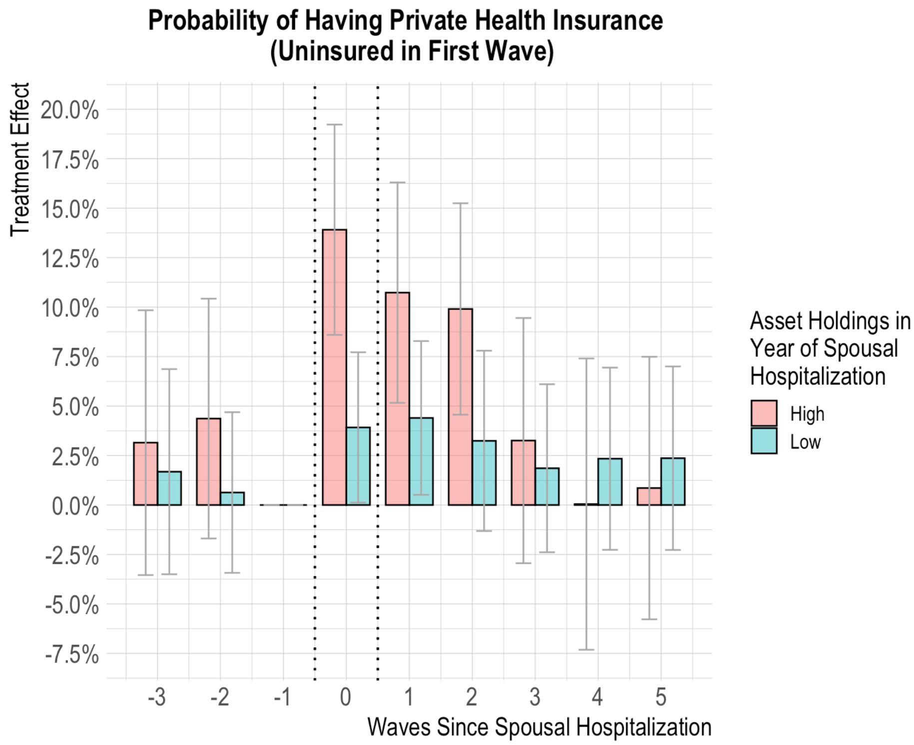

Notes: I separate *treated* individuals (i.e. those who experienced spousal hospitalization events in the the observable survey window) between those who held more than \$2,980 in assets (High asset group) and those who held less than \$2,980 in assets (Low asset group) at the time of the spousal hospitalization event. The sample used here is the "initially uninsured" group of respondents. Sample compositions (i.e. the number of respondents in each separate dataset) are listed in Table B.2. In separate regressions for equation 2.3.1, each colored bar for each time period compares each of these distinct treatment groups to the (same) group of comparison individuals (i.e. those who did not experience a spousal hospitalization event during their time in the observable survey window. I include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. *Data source:* Health and Retirement Study.

The fact that there is an observably relative large increase in private plan coverage for individuals with high assets, but a relatively smaller observable impact for individuals with low assets, suggests that there could be persistent barriers that interfere with up-take for certain disadvantaged groups. This finding is similar to the evidence above, which also suggests heterogeneous take-up by racial and ethnic group. An related explanation for the relatively low take-up for disadvantaged groups is that there are other, first-order issues preventing them from accessing health plans, and information in the form of salient health events is not sufficient to overcome these hurdles.

2.5. Conclusion

This paper examines private health insurance patterns following a spousal health shock in the form of a spousal hospitalization admission. There is significantly increased uptake (7.4pp) in private plans for survey respondents who entered the HRS without a private plan in their first survey wave. This is a combined effect of being covered by a spousal plan and an individual responding by taking out plans independently from their spouse. My results also suggest that overall, as a result of spousal hospitalizations, individuals pay significantly higher premiums in the following few survey waves.

The pattern – an initial upwards spike in treated individuals' private health insurance coverage followed by a decreasing effect size over time – is consistent with individuals responding to a salient health event that highlights health risks by selecting into more generous plan coverage. Over time, the initial shock of the health event diminishes, which is consistent with the effect size diminishing in the long-run (although alternatively, it could be that "never-treated" individuals sign-on for private insurance over successive survey waves, mechanically lowering the effect size).

Furthermore, there are alternative explanations beyond individuals responding to salient health information. One such explanation may relate to averse selection into more generous coverage. For instance, a portion of the effect of extended coverage for respondents can also be explained by being newly covered under hospitalized spouses plans ("spillover" effect). A second potential explanation underlying the observed patterns is that the transaction costs of gaining access to a private plan decrease following a spousal hospitalization event decrease: visits to the hospital may lower the cost of accessing information about plan options and navigating the sign-on process. I empirically explore this possibility further. I find that individuals whose spouses held insurance at the time of their own hospitalization event are initially 9.6pp more likely to sign-on to private coverage. A potential explanation underlying this pattern is that these respondents are more exposed to knowledge about the plan sign-on process and so information costs are lower this group. However, respondents whose spouses were *uninsured* are not initially more likely to sign-on to private coverage. If anything, the result may suggest that there may be persistent transaction and information costs for respondents, even immediately following a spousal hospitalization event. It is notable, though, that respondents whose spouses did *not* hold insurance at the time of their hospitalization event eventually do react in the long-run with greater private plan take-up. When I examine the sample by racial group, I find that white respondents are significantly more likely to hold private plans and higher premium plans as a result of spousal hospital admissions. Yet, spousal hospitalizations have no positive effect on plan take-up for Hispanic and Black respondents, and if anything, appears to lower total private premium payments at different post treatment waves for Hispanic and Black respondents.

How do individuals end up using the expanded coverage as a result of the spousal hospitalization event? I explore whether individuals are more likely to experience their own hospital admissions prior to and after the spousal hospital admission. However, I find no statistically significant effect of treatment on an individuals' *own* hospitalization visits, which lends credence to the conclusion that increased insurance coverage is "prudently" consumed; individuals with increased coverage plausibly seek out alternative treatment settings instead of ER or hospital stays. This explanation would be consistent with Miller (2012), who suggests that expanded health insurance may have a substantial impact on the efficiency of health services. It is possible that there could be potential efficiency-inducing impacts of spousal hospitalization events, due to greater levels of insurance coverage among their spouses.

Intrinsically, one of the limitations of relying on longitudinal data from the HRS is that I cannot rule out multiple possible explanations underlying patterns of higher private insurance take-up after a spousal health shock. It is plausible that individuals are motivated by the salient health event to insure themselves *or* provide additional coverage for their spouses. In spite of the need to further explore explanations underlying responses, which future research may seek to uncover using survey methods, my results highlight that individuals are more likely to be covered under more generous private health plans after a spousal health shock.

This paper sheds additional light on the complicated nature of examining the impacts of hospital admissions. Whereas existing research primarily focuses on negative financial impacts of hospital admissions their own earnings, income, and consumer borrowing (Dobkin et al., 2018), this examination suggests that any financial-impact analysis of serious health shocks should also take into account the general equilibrium effects of health shocks, which may result in second-order efficiency gains, particularly among those who observe hospital events and were initially uninsured. Furthermore, one potential explanation underlying the patterns that I identify in this paper – individuals responding to salient health information from a spousal health event by taking up additional private coverage – may have policy or social bearing on the type and format of information presented to potential consumers of health insurance.

CHAPTER III

School Choice and Student Mobility from Low-Performing Schools: Evidence from the California Open Enrollment Act (with Hayley Abourezk-Pinkstone)

3.0 Abstract

School choice policies can provide additional educational opportunities to students that would be otherwise constrained to their neighborhood school, but the effects of such policies depend on the spread of take-up. If take-up rates vary systematically across students by race or socioeconomic status because of persisting barriers to access, then school choice policies may change the distribution of students across schools and the level of racial or socioeconomic segregation across schools. Using a difference-in-difference approach and an event study, we empirically examine how the California Open Enrollment Act (2010-2016), which increased public school choice for students attending low-achieving public schools, impacted student enrollment patterns across schools by race and socioeconomic status. Total enrollment at treated schools falls by up to 1.4% as a result of the policy, and this effect persists for several years. Further, we find that Hispanic student enrollment and free-and-reduced price meal (FRPM) eligible student enrollment at treated schools fall by up to 5.9% and 6.7% relative to comparison schools, respectively, in a school's first year of treatment, and that these enrollment effects each tend to grow over time. We discuss how our context relates to recent literature that calls into question the validity of TWFE DiD and event study designs when treatment effects are heterogeneous across treatment cohorts or time. Our results are robust to an exercise aimed at eliminating this potential source of bias. Our findings suggest that the Open Enrollment Act did expand public schooling options for low-income students and minority students attending treated schools, enabling them to switch to higher-performing

public schools.

3.1. Introduction

School choice programs expand the set of schools in which families can enroll their children. A large body of research in the economics of education examines the effects of expanded choice for individual, participating students. Hastings et al. (2012) find that expanded choice results in modest positive effects for "movers" on attendance and test scores. Deming et al. (2014) document positive effects of expanded choice on future graduation rates, post-secondary attendance, and completion rates, and Deming (2011) finds negative effects on crime. Other studies find null effects[1] on similar outcome variables (Cullen et al., 2006), although overall null effects sometimes mask significant effects by race or other subgroup.[2] The existing research does not yet answer how school choice programs change the distribution of students across schools as a result of possibly non-random take-up.[3] Given evidence of how school resources, quality, and peer effects impact educational outcomes (i.e., Sacerdote (2011); Chetty et al. (2011)), the distribution of students and resources[4] across schools matters for student outcomes. Expansions of school choice can change the distribution of students and resources across schools if policy take-up rates are correlated with student characteristics such as socioeconomic status and race. We examine the effects of a recent expansion of public school choice, the California Open Enrollment Act, on differential take-up rates by race and income. In particular, we focus on mobility of minority and low-income students out of low performing schools towards higher performing schools in response to the policy and how this compares to the mobility of their peers.

Similar initiatives to the California Enrollment Act have been relatively widespread across several states in the United States, and as of 2022, 43 states have signed their own

[1]This relationship is complicated in part by the effect of movers' class rank changing; as they enroll in a higher-performing public school, the student's relative rank falls. After accounting for a student's cardinal achievement level, a lower ordinal rank tends to be empirically associated with lower grades and completion rates (Fabregas, 2020) and lower test scores (Murphy and Weinhardt, 2020).

[2]For example, in the context of voucher programs, Howell et al. (2002) finds that expanded choice increases test performance of Black students who move schools in response to the policy, despite the estimated effects for the full sample being insignificant. See also Hastings and Staiger (2006) for a discussion of heterogeneous effects.

[3]See Altonji et al. (2015) for a discussion of the theoretical underpinnings of "cream skimming" associated with school choice programs.

[4]Highly segregated schools experience unequal access to funding resources; a recent report from EdBuild highlights the $23 billion discrepancy in funding between predominantly white (75%-100% white) and predominantly minority (0%-25% white) districts in 2016, despite serving roughly the same number of students nationwide (See https://edbuild.org/content/23-billion). Segregation has been shown to harm students by depressing student performance and exacerbating inequality (Frankenberg et al., 2019).

version of legislation allowing interdistrict open enrollment in the past (Erwin et al., 2022).[5] Despite their commonality, to our knowledge, researchers have not explored the effect of open enrollment policies on school-level enrollment outcomes, making the findings of this study potentially relevant in a broader context. It has, however, been shown that school choice programs increase the ability of families to self-sort across schools. Sorting in schools has been shown to be important both in theory (Epple and Romano, 1998) and in practice (Kane and Staiger 2006; Bayer et al. 2007). Urquiola (2005) provides evidence of increased sorting in metropolitan areas that span many school districts, suggesting that access to inter-district choice increases household sorting behavior. By increasing families' sorting opportunities, expanded public school choice programs like the California Open Enrollment Act could change the distribution of students across schools and could impact racial or socioeconomic segregation. However, the direction of such an effect is not obvious; determining which students are most likely to participate in school choice programs and how this will change the student distribution across schools is not immediately clear in theory.

It is possible that school choice programs might particularly benefit low-income or minority students. In the absence of widespread public school choice programs, families are limited in their autonomy over which schools their children may attend. In California, students are generally assigned to public schools based on residential boundaries. If a family is dissatisfied with their default neighborhood school, they could enroll their child in a private school or a magnet school, or move to a neighborhood with better schools. Yet, each of these alternatives requires families to overcome major financial and institutional hurdles. Therefore, public school choice programs might be especially beneficial to low-income and/or minority families who face barriers (discrimination, financial barriers, language barriers, etc.) associated with moving schools through other channels, such as changing residences. In this sense, an expansion of school choice could provide alternative schooling options for families who might have otherwise been constrained to their neighborhood school, expanding educational options for low-income and minority students.

On the other hand, school choice policies may do little to expand the feasible choice set of schools for low-income or minority students, if the policy does not adequately eliminate barriers associated with moving to better schools. To the extent that barriers still persist when expanded public school choice opportunities are available, such as limited access to reliable transportation, information on school performance, internet access, and language differences, it may be the case that low-income or minority students are less likely to take advantage of

[5]It become more common in the "No Child Left Behind" (NCLB)/school accountability era for individual states to adopt their own open enrollment policies (Mikulecky, 2013). However, many of these state's policies, including the California Open Enrollment Act, are distinct from and are not directly a *result* of NCLB.

school choice expansions than their peers in practice, if these barriers preclude any possible gains from the policies. If this is the case, school choice programs could exacerbate unequal access to educational options across racial or socioeconomic lines, and could result in an equilibrium distribution of students across schools that is even more uneven.

Given this theoretical ambiguity, this paper examines how an expansion of public school choice affects student enrollment patterns across schools by race and socioeconomic status, by identifying which students leave low-performing schools when their choice opportunities expand, and considering how this changes the equilibrium distribution of students across schools. We identify a recent policy, the California Open Enrollment Act (2010-2016), as an event that increased the degree of public school choice for many families in California. The Open Enrollment Act required that the California Department of Education publish a list of the state's 1,000 lowest-performing schools by Academic Performance Index (API) score[6] and mandated that districts provide students at these schools the option to transfer to a higher-performing school of their choice in any district. For affected students, this legislation created both an opportunity to attend a better-performing school and a nudge to do so, in the form of information sent home that their school ranked among California's lowest performing schools (an example letter is provided in Appendix C.1, Figure C.1). The legislation also used two arbitrary cutoffs in its phrasing and implementation. The first is a "district cutoff": the law stipulates that no more than 10% of a given district's schools can appear on the open-enrollment list. The second is a "school type cutoff": the law specifies a fixed ratio of elementary, middle, and high schools that should appear on the list. Together, these cutoffs create some randomness in assignment to treatment, which allows us to assume that it is challenging for schools to anticipate their treatment status with any certainty, thus limiting any short-run behavioral responses in anticipation of treatment that might impact our results. Our paper contributes to the school-choice literature by providing novel estimates of the plausibly causal impact of an expansion of public school choice on school-level enrollment trends by race and socioeconomic status.

We find that this expansion of public school choice in California induces students to leave low-performing schools when they are given the opportunity to do so. All else equal, total

[6]The API ("Academic Performance Index") score was a single score that was assigned to each public school in California, annually or bi-annually, by the California Department of Education. API scores ranged from 200 to 1000, and indicated how well a school performed on standardized tests like the STAR (Standardized Testing and Reporting) and the CAHSEE (California High School Exit Exam). API scores served as a primary measure of school performance and accountability from 1999-00 through 2013-14.

enrollment at treated schools declines by up to 1.4%, relative to comparison[7] schools, in the first year a school is treated. Moreover, the negative effect of the law on total student enrollment in treated schools persists for several years; three years after treatment, total enrollment at treated schools continues to decline by up to 1.6% relative to comparison schools. The persistent effects of treatment over time may be partially explained by the fact that newly enrolled students are allowed to remain at their school of choice regardless of their default school's future open enrollment status, and that siblings of students who have moved schools in response to this opportunity are also allowed to attend the school of choice, regardless of open enrollment status in the year that the sibling is enrolling. Lastly, word-of-mouth and negative reputational effects may matter when considering the long-run impacts of being labelled as a "low-performing school" and experiencing the dissemination of that information to families.

When analyzing the effects of this policy along racial and socioeconomic lines, we find that the effects of the policy for Hispanic students and for low-income students in California are large and persistent over time. For instance, three years after the policy, Hispanic enrollment at treated schools falls by up to 6.3% relative to comparison schools, and enrollment of free- and reduced-price (FRPM) eligible students, our proxy for low-income student enrollment, falls by 8% in treated schools relative to comparison schools. The magnitude of the effects found for white and higher-income students are often lower and sometimes statistically indistinguishable from zero in the longer-run. We take this and other evidence discussed in the paper to suggest that the Open Enrollment Act plausibly expanded schooling opportunities for low-income students and minority students and resulted in movements from low-performing schools to higher-performing schools.

The remainder of the paper proceeds as follows: Section 3.2 outlines key elements of the California Open Enrollment Act for context. Section 3.3 summarizes our data. Section 3.4 provides initial evidence that the policy expands choice for low income and minority students using a difference-in-difference approach with two way fixed effects, which is robust to newer specifications following the literature. Section 3.5 expands on the difference-in-difference approach by including leads and lags in an event study design and discusses assumptions crucial to identification, including treatment effect homogeneity. Finally, Section 3.6 concludes.

[7]Throughout this paper, we refer to schools that do not appear on the Open Enrollment list as "comparison" schools rather than untreated schools, because they may be indirectly impacted by treatment as potential receiver schools for any students switching schools in response to the policy. As such, we recognize that point estimates reported in this paper aimed at capturing enrollment trends at treated schools *relative to* comparison schools might overstate the true effect of the policy. We attempt to size this source of bias. See Appendix C.2 for a detailed discussion.

3.2. Context: The California Open Enrollment Act

The California Open Enrollment Act, added by SBX5 Romero (Ch. 3, Fifth Extraordinary Session, Statutes of 2010), was signed into law on January 7, 2010, as California Education Code sections 48350-48361. The legislation states that the California Department of Education will publish a list of 1,000 low-performing public schools[8] ("Open Enrollment" schools) across California Local Educational Agencies ("LEAs" i.e., school districts) and students enrolled at these schools will be permitted to attend a higher-performing school of their choice. The Open Enrollment list is based on the previous years' Academic Performance Index ("API") scores, and is created using the following algorithm, as specified by publicly-available legislation documents[9]:

> *Creating the list starts with the identification of 1,000 schools in the ratio of decile 1 elementary, middle and high schools (i.e., 687 elementary schools, 165 middle schools, and 148 high schools) that have the lowest API scores. When an LEA on the list has reached its "ten percent" cap, the LEA's schools with the highest API scores are dropped from the list until the LEA has no more than its "ten percent" number of schools on the list. Schools with the next lowest API scores remaining in the pool are then added to create the next list of 1,000 schools that maintains the decile 1 ratio of schools. This process of creating a list in the decile 1 ratio ranked by API score, followed by dropping schools that exceed their LEA's ten percent cap continues until a final list of 1,000 schools is achieved that both maintains the decile 1 ratio of elementary, middle, and high schools and does not exceed any LEA's "ten percent" number of schools. (pp. 1-2)*

The legislation requires that, when a school falls on the open enrollment list, the district will notify families of enrolled and incoming students of the school's open enrollment status. An example of such a family notification letter from Martinez Unified School District is shown in Appendix C.1, Figure C.1. Notice that this letter includes both information about the school's low API score performance, and a description of how to request a transfer to a higher-performing school. We interpret this as two distinct elements of treatment: an opportunity to attend a better school, and a "nudge" to do so (information about the low performance of the current school). This letter highlights several other noteworthy features of the law, such as the fact that the California Open Enrollment Act allows for both intra- and inter-district transfers, that it does not require districts to provide transportation, and

[8]The law does not apply to charter schools, schools that are closed, and schools with fewer than 100 valid test scores, and court, community, or community day schools.

[9]Source: July 10 Item Addendum (Item 33), California Department of Education, Executive Office, SBE-004 (Rev. 06/2008).

that if more applications are received than spaces are available, a lottery determines which student(s) are permitted to transfer.[10]

Open Enrollment Lists were published for the 2010-11 through 2015-16 school years. The first year of the legislation is unique: the Open Enrollment Act was passed in January of 2010 and was intended to be administered in the 2010-11 school year. However, when the California School Boards Association and numerous schools and districts pushed back against this timeline, the California Department of Education agreed that, while an Open Enrollment list would be published for the 2010-11 school year and families would be notified of their school's status, schools need not accept transfer students in adherence with the law during the 2010-11 school year. That is, the Open Enrollment Act did *not* apply to student transfers for the 2010-11 school year, but parents *did* receive a letter notifying them of their school's API performance and open enrollment school status. In the 2011-12 through 2015-16 school years, the policy was in full effect and school districts were required to accept student transfers[11] in adherence with the law.

3.3. Data

3.3.1 Open Enrollment Status

Using open enrollment lists provided directly to us by the California Department of Education (CDE), we create a dataset that encodes each school's open enrollment status for each year that the policy was in effect. The CDE was able to provide the finalized open enrollment lists for the 2012-13, 2013-14, 2014-15, and 2015-16 school years, but the CDE no longer has record of the lists from the 2010-11 and 2011-12 school years. For the 2010-11 school year, we identified a draft list of the 1,000 schools on the Open Enrollment list sent as an addendum to members of the California State Board of Education on July 9, 2010, which we use as a proxy for the finalized list for the 2010-11 school year.[12] For the

[10]Additionally, although this is not mentioned in the attached letter, when 15% or more of the students in a school speak a primary language other than English, notices of a school's open enrollment status must be written in each student's primary language and may be answered by the parent in either language.

[11]Schools were generally required to accept transfers in adherence with the law, but there were exceptions. Districts could adopt "specific written standards" to process applications; these standards could include "consideration of the capacity of a program, class, grade level, or school building or adverse financial impact. The standards may not include consideration of a student's academic achievement, proficiency in English language, family income, or any of the prohibited bases of discrimination." When transfer applications exceeded slots available, schools were instructed to conduct lotteries (Open Enrollment Act Information Sheet).

[12]The OE draft list likely varies from the finalized list through the effect of schools applying for for last-minute waivers, but we expect that a version of the list produced in July (for August enrollment) is a close approximation of the final list, and this intuition about accuracy was confirmed by a staff member of the CDE.

2011-12 school year, we were unable to find any version of a published Open Enrollment list. Instead, using the rule discussed in Section 3.2, we design and run an iterative algorithm that ranks all schools in California by API score, selects the lowest 687 elementary, 165 middle, and 148 high schools, preserves all of the schools falling below the 10% per district cutoff while removing those falling above the 10% rank within their district from the list, and then repeats this ranking process for the remaining pool of schools, ensuring the same ratio of elementary, middle, and high schools, until the list reaches 1,000 schools. Our algorithm produces our best approximation of the actual open enrollment list in the 2011-12 school year. Nonetheless, this list is not exact, for several reasons. First, the final published open enrollment lists are generally comprised of *just under* 1,000 schools, because after a draft of the list is completed, schools are allowed to apply for waivers to be removed from the list under specific circumstances (such as in the case where a high-performing school ends up on the Open Enrollment List due to the 10% rule). Second, in our algorithm, we break ties between schools with identical API scores randomly, just as the CDE did, creating a small amount of random variation in the list.[13]

3.3.2 Treatment Intensity

Table 3.1 shows how many of the ever-treated schools were on the open enrollment list exactly once, exactly twice, and so on, up to the maximum amount possible of six times. A majority, or 1,495 out of 1,978 ever-treated schools, appear more than once. On average, an ever-treated school appears on the list around three times.

[13]Using the official published lists and our algorithm-produced lists for the missing years, we test the accuracy of our algorithm for the years in which we have a draft list (2010-2011) and a finalized list (2012-13 through 2015-16) in Table C.3.

Table 3.1: Count of Schools that Appear on the OE List Once vs. Multiple Times

Number of Appearances on OE List	Count of Unique Schools
1 Time	483
2 Times	474
3 Times	317
4 Times	251
5 Times	190
6 Times	263
Total Ever Treated	1,978
Total Never Treated	14,787

Notes: This table presents the number of unique schools that appear on the open enrollment list exactly once, twice, and so on, up to six times. For our total number of public schools, we use the "Public Schools and Districts Data File", published by the California Department of Education (Accessed on January 8, 2021). This list contains all active and previously active public schools in the state of California as of the access date. We drop all schools that have missing data from this master public school list. Data Source: California Department of Education.

3.3.3 Outcomes and Covariates

We draw school and district-level data from the California Department of Education (CDE) and the American Community Survey (ACS) 5-year estimates, downloaded from IPUMS (Integrated Public-Use Microdata Series). From the CDE, we draw school-level information such as Academic Performance Index data, directory information, total school enrollment, enrollment by race and gender, and enrollment by free- or reduced-price meal (FRPM) status. We also draw district-level information about private schools from the CDE. From the ACS, we select population-level variables for migration by school district (a count of individuals 1 year or older currently residing in the district who moved from out-of-state, and moved from abroad, within the past 12 months), median home values in each school district, and population demographic breakdowns by district including by race, age, and occupational category. Because the ACS is based on 5-year estimates, we set the year equal to the final year in the moving window (for instance, for the ACS 2015-2019 Survey, we would treat this as data for the year 2019).

Table 3.2 presents a descriptive summary of the differences between schools that never appear on the open enrollment list (never treated) and schools that appear on the OE list at least once (ever treated) at the beginning of the final year the policy was in place.[14] On average, these schools also have a higher share of free- and reduced-price eligible students, and they have a larger enrollment share of Hispanic students and a smaller enrollment share of white students.[15]

Table 3.2: School-Level Student Composition Descriptive Statistics

	Never on OE list 10-11 through 15-16	On OE List at least once
Racial shares		
% White	32.2%	16%
% Black	8%	7%
% Hispanic	55.8%	71.9%
% Asian	11.8%	4.4%
Socioeconomic proxies		
% FRPM	57%	80.2%
% Non-FRPM	43%	19.8%
2015-16 Sample size	12,976	1,953

Notes: This table illustrates school-level descriptive differences between treated and comparison schools. The columns separate schools that were ever-treated during the duration of the policy (2010-11 through 2015-16) vs. schools that were *never* treated during the duration of the policy, and presents descriptive statistics for the 2015-16 school year for each of these groups. We exclude charter schools from this table because charter schools are excluded in our main analysis. We only report four racial groups, but there are other (smaller) ethnic/racial groups not reported. Source: California Department of Education.

3.4. Staggered Difference-in-Difference Approach

In the context of the California Open Enrollment Act, increased school choice is not randomly assigned to schools and districts. Rather, it is intentionally assigned to California's lowest performing schools. Because low-performing schools differ systematically from higher performing schools[16], the absence of randomly-assigned treatment necessitates the use of quasi-experimental methods to explore the causal impact of the legislation. To start, we employ

[14]In Table C.4 (Appendix C.4), we present descriptive statistics at a district-level. Schools that appear on the open enrollment list at least once are located in districts with larger fractions of minority residents and districts with lower median income.

[15]The difference in the enrollment share of black students appears negligible, but it is slightly higher in ever-treated schools than never-treated schools. We present the racial share of Asian students in the descriptive summary (Table 3.2), but we exclude reporting this group from our main results because it is not the focus of our research question.

[16]See, for instance, Table 3.2.

a basic TWFE difference-in-difference setup with a single staggered treatment indicator, in which we seek to compare enrollment outcomes in treated and comparison schools before and after treated schools appeared on the list, exploiting variation in treatment timing. Year fixed effects account for year-on-year enrollment trends that are common across schools, and individual school fixed effects account for time-invariant school characteristics that might also impact our outcomes of interest.

Our staggered difference-in-difference framework with two-way fixed effects allows us to account for varying treatment years and intervals, and can help us understand the impact of treatment if necessary assumptions (such as parallel trends) hold. However, as Baker et al. (2021); Callaway and Sant'Anna (2021); Goodman-Bacon (2021) show, the estimated coefficient in a TWFE difference-in-difference design is a weighted average of all possible simple 2x2 DiD estimates comparing a single group that has changed treatment status to another group that has not. Importantly, negative or uneven weights underlying these sub-components of the TWFE DiD estimator can result in biased treatment estimates. Consequently, we recognize that our basic TWFE difference-in-difference approach only yields an unbiased estimate of the average treatment effect on the treated (ATT) if treatment effects are homogeneous across all six treatment cohorts $\in \{2010, 2011, 2012, 2013, 2014, 2015\}$ and across time.

3.4.1 Identification Assumptions

3.4.1.1 Parallel Trends

Our difference-in-difference approach hinges on the assumption that, in the absence of the California Open Enrollment Act, schools that were treated and not treated under the Act would have observed similar changes in our outcomes of interest. While it is not possible to assess whether this assumption about the counterfactual holds, we identify several theoretical violations of the parallel trends assumption that could confound our results. For example, treated schools by definition have lower API scores on average than their untreated counterparts, and this information is publicly available to families who seek it out. Given lower performance, families with children enrolled in treated schools might have more of an incentive to change schools for their children than families with children enrolled in comparison schools, even in the absence of the law. While their legal ability to move to a different public school is limited in the absence of the law, families could enroll their child at a private school or move residences in order to change schools. If families with students enrolled at treated schools pursue these strategies at a different rate than families with students enrolled in comparison schools, this could lead to a violation of the parallel trends assumption.

It is also possible that migration rates into and out of treated and untreated school districts vary over time in systematically different ways. For example, we know that treated schools serve significantly larger Hispanic student populations than untreated schools (Table 3.2). California experiences large flows of Hispanic immigrants from Mexico. If Hispanic families migrate faster into districts with many treated schools than into districts with few treated schools during this time period (i.e. because of existing social networks), this could serve as a potential violation of the parallel trends assumption. Because of this potential threat to the parallel trends assumption, we include district-level information on population in-migration (migration from out-of-state, and migration from abroad) for each district in our analysis.

3.4.1.2 Treatment Effect Homogeneity

As Baker et al. (2021); Callaway and Sant'Anna (2021); Goodman-Bacon (2021) highlight, staggered TWFE difference-in-difference designs may generate biased estimates of the average treatment effect on the treated when the effects of the policy are heterogeneous across treatment cohorts or across time. Baker et al. (2021) suggest the following:

> *Under some conditions–when treatment effects can evolve over time (when there are "dynamic treatment effects")–staggered DiD estimates can obtain the opposite sign compared to the true ATE or ATT, even when the researcher is able to randomize treatment assignment. The intuition is that in the standard staggered DiD approach, already treated units can act as effective comparison units, and changes in their outcomes over time are subtracted from the changes of later-treated units (the treated). (p. 2)*

If treatment effects are homogeneous across cohorts[17] and years, then our estimates of the ATT do not contain this source of bias. However, there may be reasons why the true treatment effect could be heterogeneous across cohorts or years. For instance, earlier treatment cohorts might have experienced less informational salience about the availability of choice and the details of the policy.[18] Also, changing population demographics and migration during the period encompassing the policy (2010-2016) might lead to different demographic compositions across cohorts, possibly resulting in different responses to treatment. In light of these concerns, we provide a robustness check in Appendix C.5.1 where we compute the average treatment effect on the treated separately for each treatment cohort and year, as

[17] Treatment "cohort" g is the set of schools that are first treated in year g.

[18] In particular, we might expect the treatment effect of the policy for the 2010 cohort to be different than the treatment effect for other cohorts, since an OE list was published for 2010-11 but districts were not required to accept transfers in adherence with the policy in that year.

suggested in Callaway and Sant'Anna (2021). Results from this robustness exercise mostly align with the patterns identified in this section, suggesting that bias from treatment effect heterogeneity is not having a large impact on our reported estimates.

3.4.2 Model

In our difference-in-difference set-up, the relationship between open enrollment status and outcomes (enrollment, enrollment by race, enrollment by FRPM status, etc.) is expressed as:

$$y_{it} = \beta_0 + \beta_1 (OE\_EverTreated * Post)_{it} + \beta_2 X_{it} + \tau_t + \gamma_i + \varepsilon_{it} \tag{3.4.1}$$

where y_{it} is the outcome variable of interest (total student enrollment, enrollment by race, and enrollment by FRPM status), $(OE\_EverTreated * Post)_{it}$ is a staggered policy dummy that equals one if school i has been treated under the Open Enrollment Act *as of* year t, meaning it equals one if school i ever appears on the open enrollment list *and* this treatment occurred before or in year t, τ_t for $t \in T$ represents time fixed effects by year, where the set of (school) years is $T = \{2002-2003, 2003-2004, ..., 2018-2019\}$, γ_i represents school fixed effects to ensure that we are capturing the within-school effect of the open enrollment policy on student enrollment behavior, X_{it} is a vector of time-varying district-level covariates that affect the outcome of interest, currently made up of two district-level population in-migration variables as discussed above, and ε_{it} is an error term. The coefficient of interest is β_1.

Our choice to set the treatment indicator equal to one if school i has been treated *as of* year t instead of precisely *in* year t, even though technically open enrollment status of a school turns "on" and "off" year-by-year, reflects the fact that students who enroll in a school of choice under the policy in a given year are permitted to *stay* at that school even if the school does not appear on the open enrollment list in subsequent years, and that students' siblings are also permitted to enroll in the same school. That is, if a school was on the open enrollment list in 2013-14 but not in 2014-15, that school would still be experiencing the effects of treatment in 2014-15 despite its status. If we categorized that school as "not treated" in 2014-15 despite it experiencing the effects of treatment, our ability to estimate the effect of expanded school choice on that school's enrollment patterns would diminish, and point estimates would likely underestimate the true effect of the law. Therefore, we set $(OE\_EverTreated*Post)_{it}$ equal to one when the school first appears on the OE list and then hold that value for the remaining years $t \in T$. This is consistent with the setup of staggered difference-in-difference designs presented in recent research on TWFE DiD estimators[19].

[19]For example, see Assumption 1 in Callaway and Sant'Anna (2021) ("irreversibility"), which they describe as the assumption that "units do not forget about the treatment experience."

3.4.3 Staggered Difference-in-Difference Results

Table 3.3 presents the results for the staggered difference-in-difference approach associated with Equation 3.4.1. Results for the year of the policy are presented in the first column. We also include results for when we lag OE treatment dummy by one, two, or three periods. This allows us to provide estimates of initial and longer-run effects of the policy on enrollment at treatment schools. In the initial treatment year total enrollment decreases by up to 1.4% at treated schools, relative to comparison schools. FRPM-eligible student enrollment decreases by up to 6.7% and non-FRPM-eligible student enrollment decreases by up to 6.6% at treated schools.[20] In the longer-run, the effect of the OE policy on total enrollment increases slightly from a 1.4% enrollment drop to as much as a 1.6% drop three years after treatment. FRPM eligible enrollment continues to drop in response to the policy in the long-run, reaching up to a 8.0% drop in enrollment relative to comparison schools three years after initial treatment. This provides suggestive evidence that students from low-income backgrounds take advantage of the expanded school choice policy over time.

Table 3.3 also demonstrates how enrollment falls in response to the policy across different racial groups. In the year of treatment, Hispanic student enrollment declines by as much as 5.9% at treatment schools relative to comparison schools. Hispanic enrollment continues to drop when we consider the lagged, long-run effects of the policy in Table 3.3, columns (3)-(5). From examination, it also appears that the initial fall in white student enrollment (up to -2.8%) at open enrollment schools relative to comparison schools tapers off after a few years and becomes statistically indistinguishable from zero. Black student enrollment at treated schools initially falls by up to 3.2% relative to comparison schools, but by the third year after a treated school first appears on the OE list, there is no statistically significant impact for this group. The magnitudes of these declines in student enrollment in response to the policy are sizeable when compared to existing stylized facts that document demand for expanded school choice over time. For instance, according to a descriptive report on take-up of school choice policies, the percentage of children attending a "chosen" public school other than their assigned public school increased from 11% to 16% from 1993 to 2007.[21] Additionally, although we do not report these in our main results, we find that low-performing, treated schools experience a small boost (slightly over 3% increase, $P < .01$) in their gAPI score in the year of treatment. This could be due to several reasons, including schools responding to falling on the OE list by trying to improve academic outcomes. However, this finding is promising because it provides suggestive evidence that the policy did not harm "stayers" (i.e.

[20]See Appendix C.3 for a more detailed discussion that explains the discrepancy between the percentage decline for sub-groups and overall percentage decline in enrollment.

[21]Mikulecky (2013); Data Source: National Center for Education Statistics.

students who did not respond to the policy).

Importantly, a relevant concern for interpreting the estimates across our empirical approaches is that the policy will simultaneously result in decreased enrollment at treatment schools and lead to increased enrollment at comparison schools because students are transitioning from lower-performing treatment schools to higher-performing comparison schools. When measuring the impact of the policy using enrollment changes in treatment schools *relative to* comparison schools, we might therefore overstate the true magnitude of the effect of treatment. We address the plausible concern that the magnitudes of our estimates may potentially overstate the true effect of treatment in Appendix C.2. We use a back-of-the-envelope calculation to we estimate that the size of any potential bias. Our calculation follows from the intuition that students who leave open enrollment (i.e. treatment schools) are likely to disperse towards other higher-performing public schools that were not on the open enrollment list in a given year. Regardless of the specific recipient schools the students move towards in response to the OE policy, each recipient school *on average* receives a number of students proportional to the number of donors schools divided by the number of recipient (i.e. not on the open enrollment list) schools in a given year.[22] We provide a formal derivation of this in Appendix C.2. Furthermore, we find that from this approach, the potential bias is likely to be relatively small (under a 12% bias). The magnitude of our reported point estimates might therefore contain this bias, but the overall conclusions of our paper remain unchanged.

[22]As we also mention in Appendix C.2, it is possible that students may have substituted to higher-performing schools that were *also* on the OE list in a given year, although we assume that this likely comprises a small number of student transfers. However, even in the extreme case that all treatment students substituted to other "higher performing" open enrollment schools in a given year, this would have a downward, rather than upward bias on our point estimates.

Table 3.3: Staggered Difference-in-Difference Results

(1) Outcome Variable	(2) Year of Treatment	(3) +1 Year	(4) +2 Years	(5) +3 Years	(6) N *(Year of)*
Log Total Enrollment	$-.014^{**}$ (.006)	$-.015^{**}$ (.006)	$-.016^{***}$ (.006)	$-.016^{***}$ (.006)	95,861
Log Black Enrollment	$-.032^{**}$ (.013)	$-.022^{*}$ (.012)	$-.020^{*}$ (.012)	$-.014$ (.012)	85,197
Log Hispanic Enrollment	$-.059^{***}$ (.008)	$-.060^{***}$ (.007)	$-.063^{***}$ (.007)	$-.063^{***}$ (.007)	95,053
Log White Enrollment	$-.028^{**}$ (.013)	$-.010$ (.011)	$-.015$ (.011)	$-.003$ (.011)	93,810
Log FRPM Enrollment	$-.067^{***}$ (.007)	$-.073^{***}$ (.007)	$-.078^{***}$ (.007)	$-.080^{***}$ (.007)	94,058
Log non-FRPM Enrollment	$-.066^{***}$ (.013)	$-.060^{***}$ (.013)	$-.048^{***}$ (.012)	$-.043^{***}$ (.012)	92,292

Notes: This table presents TWFE difference-in-difference results in the short- and longer-run, associated with Equation 3.4.1. Outcome variables are listed in the left-most column, including enrollment, enrollment by race, and enrollment by FRPM (free and reduced price meal) status. In this table, we lag the treatment dummy from Equation 3.4.1 ($OE\_EverTreated * Post$) successively for up to 3 years to study the longer-term effects of treatment on enrollment patterns. All estimates in this table include school and year fixed effects and the covariates discussed in Section 3.4. Standard errors (in parenthesis) are clustered at the school level. The final column lists the regression sample size for each of the outcome variables in the year of treatment, although it is possible for sample sizes to change for the lagged treatment outcomes. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.10$. This figure reports unadjusted point estimates; see Appendix C.2.2 for an approximation of the potential upper bound of bias on the estimates. Source: California Department of Education; CALPADS; IPUMS.

3.5. Event Study

We move away from the difference-in-difference model that estimates a single coefficient for effect of treatment under the Open Enrollment Act towards an event study model that incorporates leads and lags of treatment under the policy. The advantage of the event study approach, over the staggered difference-in-difference approach, is that it allows us to visualize the coefficients on indicator variables for "years relative" to the first year a treated school appears on the open enrollment list.[23] This setup allows us to estimate the pattern of enrollment changes over time in response to the treatment, relative to the year before the first time a school appeared on the OE list. In our specification, we denote $Event_i$ to represent a variable recording the calendar year t when school i appears on the open enrollment list for

[23]However, as we discuss below, our empirical event study potentially faces the same threats to identification that the staggered difference-in-difference model faces; namely, treatment effect heterogeneity from multiple treatment periods may contaminate estimates of pre-trends and pose a threat to the assumption of parallel trends.

the first time. Specifically, we estimate[24]:

$$y_{it} = \sum_{Y=-3}^{-2} \mu_Y(\mathbb{1}[t = Event_i - Y])_{it} + \sum_{Y=0}^{9} \mu_Y(\mathbb{1}[t = Event_i + Y])_{it} + \beta X_{it} + \gamma_i + \tau_t + \varepsilon_{it} \quad (3.5.1)$$

where y_{it} is the outcome variable of interest (total enrollment, enrollment by race, and enrollment by FRPM status), μ_Y are coefficients on indicators for years relative to first time a school appears on the open enrollment list[25], where μ_{-1} is omitted such that all other estimated effects are computed relative to this year, X_{it} is a vector of time-varying school level controls currently containing measures of population in-migration as described below, γ_i are school fixed effects, τ_t are year fixed effects, and ε_{it} is an error term. The coefficients of interest are the μ_Y's and the pattern they take over time.

While there may be additional school-level controls worth including in the vector X_{it}, of particular interest to us is the role that in-migration into districts may have had on outcome variables (student enrollment, enrollment by subgroup) during the period of study. Migration is not captured within time or school fixed effects, because migration is time-varying across school districts. Consequently, we include two population measures of in-migration in X_{it} calculated at the school-district level: the number of individuals age 1 year or older residing in the school district serving school i who, as of time t, had moved into the district within the past 12 months from (a) abroad[26] and (b) from a different state[27]. Our decision to adjust for migration is based on our observation of the significant variability of migration across districts in California. California has a population of 10.5 million immigrants; 46% of children in California have at least one immigrant parent and 49% of California's immigrants were born in Latin America (Perez et al., 2023). In our data set, between the years of 2009 and 2019, cumulative migration from abroad makes up an average of 5.4% of each districts' total population and an additional 10.4% of the district's population is comprised of migrants from different states. Furthermore, average in-migration over the past 10 years as a percent of a school district's total population in 2019 ranges from about 0% to about 67% across different California public school districts, reflecting highly variable in-migration rates across districts. Varying in-migration rates across districts might reflect existing social networks that are clustered in particular school districts, which may systematically correlate with school

[24]We include all possible relative lags (9 years), three relative leads to analyze potential pre-trends, and we do not trim the sample. We include "never treated" units in our analysis.

[25]For example, μ_0 is the coefficient on an indicator that takes a value of one if school i was treated precisely in year t and zero otherwise, μ_3 is the coefficient on an indicator variable that takes a value of one if school i was treated precisely three years before year t, and zero otherwise, and so on.

[26]IPUMS NHGIS, Source Code: RWJE081

[27]IPUMS NHGIS, Source Code: RWJE065

resources and treatment status. Given the large number of immigrant children in the state and the extent to which immigration varies across districts, contemporaneous in-migration could be correlated with both the enrollment outcomes of interest and treatment. Ignoring the effects of migration could thus result in omitted variable bias.

3.5.1 Identification Assumptions

3.5.1.1 Conditional Parallel Trends

For the parallel trends assumption to hold in our setting, it must be the case that in the absence of the Open Enrollment Act, enrollment outcomes at treated schools would have grown at the same rates as enrollment outcomes in comparison schools. As we discuss above, however, we choose to include in-migration as a covariate in X_{it} to address potential omitted variable bias. In this vein, without conditioning on migration, parallel trends could be violated. For instance, if enrollment increased or decreased for migration-related reasons during policy years (2010-2016) in ways that differ across districts with a large number of treatment (or comparison) schools, their student enrollment paths would *not* be parallel in the absence of the policy. Thus, following Roth et al. (2023), we only require the parallel trends assumption to hold conditional on time-varying covariates, which in our case are contemporary measures of in-migration from abroad and from other states. As highlighted by Roth et al. (2023), this requires us to assume that treatment does not impact our measures of in-migration from abroad or other states. We find this assumption to be plausible, considering that individuals in different states or abroad are less likely to know about the open enrollment policy nor base migration decisions on the policy. In contrast, the most likely way the open enrollment policy could impact migration would be movements *within* districts or within the state, which we do not include for this reason.

3.5.1.2 No Anticipatory Behavior

It is plausible to assume, in our setting, that treated schools would not be likely to know with certainty that they would appear on a given years' open enrollment list, precisely because the legislation makes use of arbitrary cutoffs in the number of elementary, middle, and high schools allowed on the list, and the rule that no more than 10% of a district's schools appear on the list in a given year. However, this assumption may not hold in later years of the policy, because we observe serial correlation in treatment status for schools that have appeared on the open enrollment list in previous years. That said, even if a school's administration knew with a high probability that it would appear on a given years' open enrollment list, there are limited *behavioral* responses that a school could employ to prevent an appearance on the list (i.e. it is impossible for a school to impact gAPI scores of other schools in its district and it may be too short a time period for schools to systematically attempt to raise their own standardized test scores in anticipation of the policy, which determines OE list eligibility). Schools are also unlikely to have many behavioral responses that could impact their enrollment outcomes in the short-run, which are largely determined by residential boundaries outside the context of the law. For these reasons, we do not believe that anticipatory behavior is having a major impact on our results.

3.5.1.3 Treatment Effect Homogeneity

As Sun and Abraham (2021) describe, staggered treatment timing in an event study setting may result in biased estimates of the μ_Y's if treatment effects of the policy are heterogeneous across different treatment cohorts or across time. For the purpose of this paper, we assume that treatment effects are homogeneous for successively treated cohorts and across time. If this holds, our results would not contain the source of bias discussed in Sun and Abraham (2021). However, one potential concern about this assumption might be that changes in school composition over time in different districts (i.e. by race, immigration status, etc.) could lead to heterogeneous treatment effects across treatment cohorts, if different subgroups of students respond to the policy differently.

Researchers frequently analyze event study pre-trends as suggestive evidence in support of or against the parallel trends assumption. Sun and Abraham (2021) point out that this is problematic in the case of treatment effect heterogeneity. If the assumption of homogeneous treatment effects across different treatment cohorts and across time does not hold, then our estimate of the pre-trends, such as our coefficient μ_{-2}, could be contaminated, which would pose a risk to any conclusions we draw from observing insignificant leads in the event study. We use the regression-based method that Sun and Abraham (2021) use to determine the underlying weights of the μ_Y coefficients. This method allows us to determine the extent to which treatment effect heterogeneity could be contaminating our estimates of pre-trends. Appendix C.5.2 provides our results from using this approach for our estimate of μ_{-2}. We find minimal evidence of non-zero weights on lags of indicators for treatment factoring into μ_{-2}, unlike the example provided in Sun and Abraham (2021). This provides suggestive evidence that our estimate of μ_{-2} is not particularly sensitive to heterogeneous treatment effects.

3.5.2 Event Study Results

Figure 3.1 demonstrates an expected trend: total student enrollment at treated schools decreases following the first time a treated school appears on the open enrollment list (initially by 0.7%, $P < .10$, and after 9 years by 3.2%, $P < .01$). In other words, the negative effect of the law on overall student enrollment in treated schools persists for several years. As we mention in the DiD results section discussion, the sign and magnitude of individuals leaving treated schools is consistent with existing stylized facts suggesting that demand for school choice over time has increased. The lagged effects of treatment may be partially explained by the fact that newly enrolled students are allowed to remain at their school of choice regardless of their default school's future open enrollment status, and that siblings of students who have moved schools in response to this opportunity are also allowed to attend the school of choice, regardless of open enrollment status in the year that the sibling is enrolling. Lastly, word-of-mouth and reputational effects of the open enrollment schools may matter when we consider the long-run impacts of being labelled as a low-performing school and experiencing the dissemination of that information to families.

Figure 3.1 also presents treatment effects for three racial groups of interest: Hispanic, Black, and white students. In the initial treatment wave, Hispanic enrollment falls by 2.5% ($P < .01$), Black enrollment falls by 2.9% ($P < .05$), and white enrollment falls by 4.6% ($P < .01$). By nine-years post policy, Hispanic enrollment has fallen by 12.5% ($P < .05$), but Black enrollment and white enrollment are not statistically different from zero in this longer-term horizon. In other words, the sustained long-run decline in enrollment as a result of the open enrollment policy only occurs for Hispanic students and is not apparent for other racial groups of interest. Figure 3.1 also presents policy effects on FRPM-eligible student enrollment and non-FPRM-eligible student enrollment at treated schools. The policy results in an initial drop in FRPM eligible student enrollment of 3.4% ($P < .01$), which continues to increase in the years following the first time a school appears on the open enrollment list to a 15.7% drop ($P < .01$). Non-FRPM student enrollment falls by 2.6% ($P < .01$) initially, eventually reaching a 7.5% drop nine years post-initial treatment ($P < .01$)

These results provide suggestive evidence that minority students (particularly Hispanic students) and low-income students appear to leave low-performing schools in response expanded school choice, and that this effect in some cases grows over time. Since the California Open Enrollment Act mandated that students attend a higher-performing school of choice, each student that left a low-performing school in response to this policy moved to a higher-performing one. Furthermore, it is notable that this movement appears to be long-lasting, particularly for of FRPM (low-income) students and Hispanic students, which may be suggestive of the policy lowering otherwise restrictive barriers associated with accessing higher quality schools for these subgroups of students.

Figure 3.1: Estimated Treatment Effect of OE Policy on Enrollment Patterns

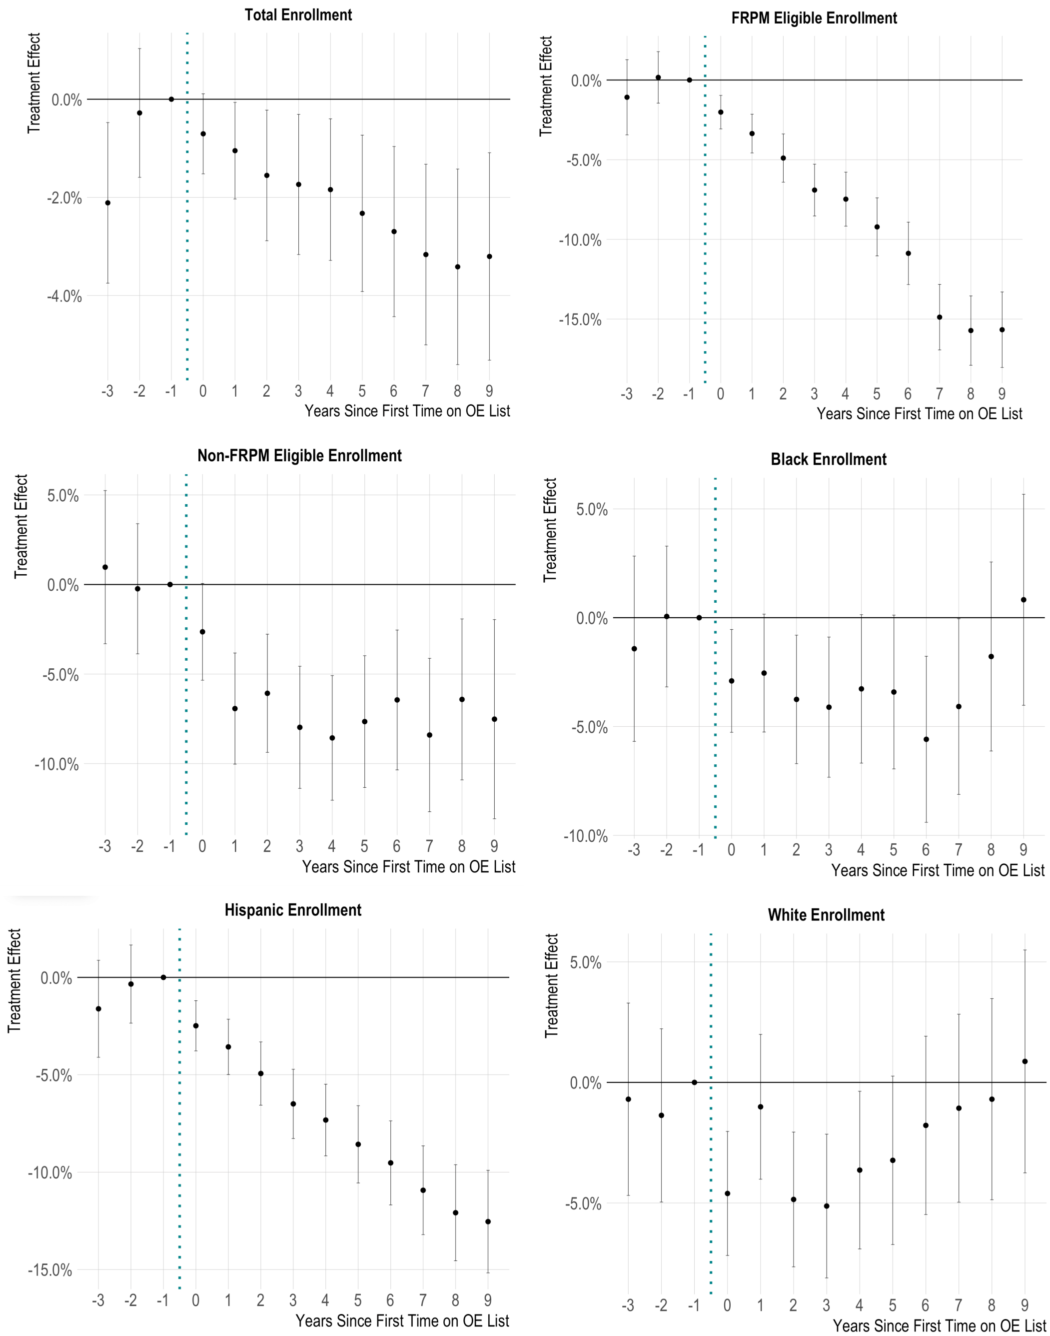

Notes: This figure displays point estimates and 95% confidence intervals from our preferred event study specification (Equation 3.5.1). On the figures, we set the estimate for the omitted period, μ_{-1}, equal to zero. We run the event study on our main outcome variables. We use the natural logarithm of our outcome variables on the LHS of our event study equation, such that this allows us to interpret the treatment effects as "percent change in enrollment due to treatment, relative to the comparison group." This figure reports unadjusted point estimates; see Appendix C.2.2 for an approximation of the potential upper bound of bias on the estimates. Data Source: California Department of Education; CALPADS; IPUMS.

3.6. Conclusion

This paper documents the effects of an expansion of public school choice on enrollment trends by race and socioeconomic status. While the existing school choice literature has primarily focused on student-level academic outcomes as a result of expanded choice for the subset of students who move schools in response to the policy, this paper takes a broader view to examine both the subgroups of students who move when choice opportunities expand, and how this might affect the distribution of students across schools. The distribution of students across schools by race and socioeconomic status, and the extent to which schools are segregated across these lines, matters for student outcomes. For instance, studies show that at least for some minority groups (e.g. Black students), school desegregation significantly increases educational and occupational attainment, reduces the probability of incarceration, and improves adult health status and overall life expectancy (Hahn 2022; Johnson 2015). California K-12 schools have been shown to be particularly segregated along racial lines, with an increasingly large majority of Hispanic students in California attending an intensely segregated[28] school (Frankenberg et al., 2019). Segregation in schools is negatively associated with academic achievement and labor market outcomes for minority students, and contributes to unequal access to funding and financial resources across schools (Frankenberg et al., 2019; Card and Rothstein, 2007). Given this evidence, policy-makers might consider the ways that school choice policies interact with enrollment trends by socioeconomic status and race.

We identify the California Open Enrollment Act (2010-2016) as a convenient context for the study of school choice, because, in each year of the policy, some schools in a given district are treated while others are not. Furthermore, treatment is partially determined by two arbitrary cutoffs written into the legislation. We exploit variation in treatment timing to first set up a staggered difference-in-difference approach with TWFE. Given recent empirical evidence suggesting that TWFE DiD estimators may be biased estimates of the average treatment effect on the treated when treatment effects are heterogeneous across treatment cohorts (g) or years (t), we compute disaggregated $ATT(g, t)$ parameters as suggested by Callaway and Sant'Anna (2021) that allow us to avoid this source of bias. Our results are robust to this approach, even without conditioning on covariates. While the TWFE DiD estimator can provide us with a single estimate of the impact of the treatment, we also use an event study to visualize enrollment trajectories in treated schools over time. The event study also allows us to visually identify whether treatment and comparison schools experience differential enrollment growth in the years before the policy, which they do not appear to do.

[28] In Frankenberg et al. (2019), an "intensely segregated" school is one that enrolls 90-100% non-white students or 90-100% white students. In 2016, 58% of Hispanic or Latino students in California attended a school that qualifies as intensely segregated.

While heterogeneous treatment effects could contaminate the estimated coefficients used to test for pre-trends in the event study, our analysis of the weights underlying the key coefficient of interest for capturing pre-trends (μ_{-2}) in accordance with suggested robustness checks from Sun and Abraham (2021) suggests that the pre-trends estimated in our event study are not particularly susceptible to this source of bias.

Overall, the treatment effects we identify using these two approaches produce relatively consistent results. Total enrollment declines at treated schools relative to comparison schools as a result of a school first appearing on the Open Enrollment list. Hispanic student enrollment falls significantly at treated schools relative to comparison schools in response to the policy. The magnitude of this effect is consistently at least as large or larger than our estimated treatment effect for white student enrollment. The effect of the policy on Hispanic student enrollment appears to persist over time. Enrollment of free- and reduced-price meal (FRPM)-eligible students also falls significantly relative to comparison schools in response to the policy, and the magnitude of this effect is consistently as large as or larger than our estimated treatment effect for non-FRPM eligible students. The effect of the policy on FRPM student enrollment also persists strongly over time. We take this as suggestive evidence that minority students and low-income students might respond to the OE policy as much as, if not even more than, their peers.

Our findings could be consistent with these groups having faced barriers associated with changing schools through other means in the absence of the policy, and the policy mitigating some of these barriers. Given that treated schools disproportionately serve Hispanic students and FRPM-eligible students, we consider that the especially large and persistent enrollment declines of each of these groups at low-performing, treated schools in response to the policy could decrease racial and/or socioeconomic segregation over time. At a minimum, the policy provides students who might otherwise be constrained to a low-performing neighborhood school with opportunities to attend better performing schools. Future research may attempt to link student-level data to the enrollment trends we identify as a result of the California Open Enrollment Act, which could enable the study of longer-term academic, occupational, and health outcomes at the individual-level for students who respond to expansions of public school choice.

APPENDICES

APPENDIX A

Appendix to Chapter 1

A.1. Additional Tables and Figures Appendix

Figure A.1: Survey Characteristics

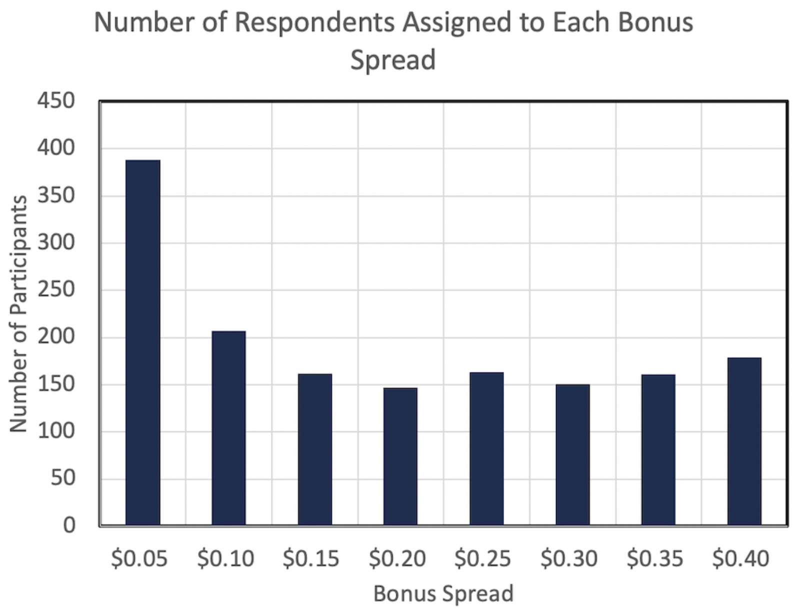

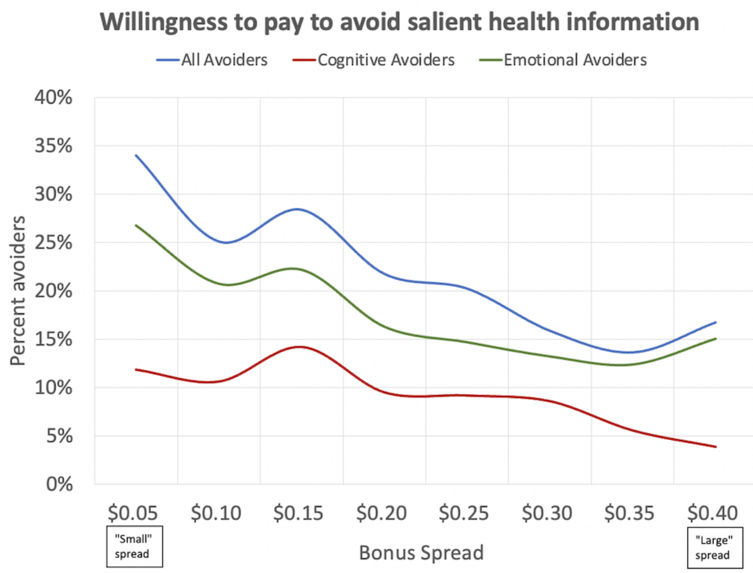

Note: The top graph shows the number of respondents (both avoiders and non-avoiders) I experimentally assigned to each bonus spread condition. The spread is the bonus paid for viewing health information minus the bonus paid for viewing non-health information. For example, in the spread of $.05, all individuals select between preferring salient health information for $.15 or non-health information for $.10. The bottom graph plots the relationship between avoidance and the bonus spread. As the incentive to view salient health information increases, a smaller percent of respondents prefer to avoid the information.

Table A.1: Balance Table: Overall Treatment versus Control

Covariate	(1) Treatment	(2) Control	(3) (t-test) Difference
N	823	735	
Proportion Uninsured (health)	.1701	.1755	−.0032 (.0192)
Proportion Male	.4738	.4217	.0535** (.0253)
Age	34.3997	35.0204	−.5015 (.4798)
Proportion Democrat	.5565	.5632	−.0082 (.0252)
Ethnicity			
Proportion White	.7411	.7551	−.0133 (.0221)
Proportion Black	.0643	.0707	−.0061 (.0127)
Proportion Hispanic	.0680	.0707	−.0038 (.0129)
Proportion Asian	.0959	.0666	.0279** (.0139)
Education			
Proportion w/ at least high school degree	.7618	.7551	.0069 (.0218)
Proportion w/ at least college degree	.2162	.2353	−.0225 (.0212)
Annual Income			
Less than $20,000	.2839	.2514	.0365 (.0229)
Between $20,000 and $35,000	.2015	.1991	.0029 (.0208)
Between $35,000 and $60,000	.2509	.2711	−.0201 (.0228)
Between $60,000 and $100,000	.2002	.2104	−.0137 (.0209)
Above $100,000	.0633	.0677	−.0056 (.0128)
Pre-treatment survey characteristics			
Pre-treatment survey time (seconds)	311.1485	303.9981	8.2195 (9.2246)
Pre-treatment survey clicks	30.7193	29.9129	.7969 (.9757)
F-test of joint significance (F-stat)			0.89 (0.5868)
F-test, number of observations			1,475

Note: Column 3 reports the difference in means between treatment and control groups including a dummy for randomization strata (whether an individual chose to avoid information or not). The value displayed for the joint test of significance are the F-statistics and in parentheses under is the p-value associated with the F-test. The values in the parentheses for the reported values in the *Difference* columns are the standard errors.
*, **, and *** denote significance at the 10, 5, and 1 percent level.

Table A.2: Balance Table: Treatment versus Control

	Only Avoiders			*Only Non-Avoiders*		
	(1) Treatment	(2) Control	(3) (t-test) Difference	(4) Treatment	(5) Control	(6) (t-test) Difference
N	169	202		654	533	
Proportion uninsured (health)	.1834	.2079	−.0244	.1666	.1632	.0034
			(.0415)			(.0216)
Proportion male	.4911	.4356	.0554	.4694	.4165	.0529*
			(.0520)			(.0289)
Age	35.3491	36.6089	−1.2597	34.1544	34.4183	−.2639
			(1.0865)			(.5305)
Proportion Democrat	.5502	.5396	.0106	.5581	.5722	−.0141
			(.0520)			(.0289)
Ethnicity						
Proportion White	.7573	.7524	.0049	.7370	.7560	−.0190
			(.0449)			(.0254)
Proportion Black	.0946	.0495	.0451*	.0565	.0787	−.0222
			(.0265)			(.0145)
Proportion Hispanic	.0414	.0693	−.0278	.0749	.0712	.0036
			(.0241)			(.0152)
Proportion Asian	.0591	.0693	−.0101	.1055	.0656	.0398**
			(.0257)			(.0164)
Education						
Proportion w/ at least high school degree	.7218	.7920	−.0701	.7721	.7410	.0310
			(.0444)			(.0249)
Proportion w/ at least college degree	.1420	.2277	−.0857**	.2354	.2382	−.0027
			(.0406)			(.0248)
Annual Income						
Less than $20,000	.3641	.2722	.0919*	.2631	.2437	.0194
			(.0494)			(.0258)
Between $20,000 and $35,000	.1975	.2146	−.0171	.2025	.1934	.0091
			(.0433)			(.0293)
Between $35,000 and $60,000	.2530	.2722	−.0191	.2503	.2707	−.0203
			(.0471)			(.0260)
Between $60,000 and $100,000	.1481	.1780	−.0298	.2137	.2224	−.0087
			(.0396)			(.0245)
Above $100,000	.03703	.0628	−.0257	.0701	.0696	.0005
			(.0235)			(.0151)
Pre-treatment survey characteristics						
Pre-treatment survey time (seconds)	325.7464	313.1723	12.5741	307.3763	300.5212	6.8550
			(21.5942)			(10.0572)
Pre-treatment survey clicks	31.7988	28.8217	2.9770	30.4403	30.3264	.1139
			(2.1356)			(1.0927)
F-test of joint significance (F-stat)			1.14		1.02	
			(0.3145)		(0.4284)	
F-test, number of observations			353		1,140	

Note: The value displayed for the joint test of significance are the F-statistics and in parentheses under is the p-value associated with the F-test. The values in the parentheses for the reported values in the *Difference* columns are the standard errors.
*, **, and *** denote significance at the 10, 5, and 1 percent level.

Table A.3: Sample Characteristics of the Primary and Secondary Outcomes

	(1) Avoiders	(2) Non-Avoiders	(3) Difference
Primary outcomes			
Switch to generous plan	50.834	47.6313	3.2027** (1.4289)
Switch to plan	17.9784	15.7030	2.2754** (1.0549)
Secondary outcomes			
Count no insurance	2.1013	1.8099	.2914** (.1377)
Count high-deductible plan	3.9702	3.6849	.2853* (.1608)
Count generous plan	4.9283	5.5051	−.5767*** (.1737)

Note: The **primary outcomes** are measured in percentage points and can take on a value between 0% (indicates an individual switches to a more generous plan option when the probability of having a health event is 0%) and 100% (indicates an individual is waits until the health event is certain to switch to a more generous plan). The **secondary outcomes** are measured in counts and are the number of times an individual selects each of the three plan options for the 11 potential options (ranging 0% to 100% in increments of 10% in the health simulation). *Note:* I report means and the difference column after including bonus spread as a control.

Table A.4: Association between Survey Time and Choice of Insurance Plan for Treated Respondents

	(1) Switch to generous plan	(2) Switch to plan	(3) Count generous plan	(4) Count high-deductible plan
Time Spent on Treatment (minutes)	$-.90^{**}$	.03	$.11^{**}$	$-.10^{*}$
	(.44)	(.35)	(.06)	(.05)
Total Clicks on Treatment	$-.10^{**}$	.02	$.01^{**}$	$-.01^{**}$
	(.05)	(.04)	(.01)	(.01)

Note: This table reports the results from running an OLS regression with for outcomes listed on the top row on the time individuals spent on health information (in minutes) and the number of clicks (a proxy for engagement with the material). This is carried out for only individuals assigned to treatment (N = 823).

*, **, and *** denote significance at the 10, 5, and 1 percent level.

Table A.4 shows that descriptively, treated individuals who spend more time on treatment appear to select more generous plans (columns 1 and 3). There is also suggestive evidence that individuals who spend longer on pre-treatment (learning and being quizzed about insurance terms) tend to select "no insurance" more (column 4). This points to a potential underlying explanation behind a lack of the anticipated, measurable treatment effect that would result in individuals selecting more generous health plans. In particular, this suggests that perhaps individuals in the treatment arms for this experiment did not spend sufficient time engaging with the material. Alternatively, this may be suggestive that a larger "dose" of health consequence information could be more conducive of individuals spending a longer duration on the treatment material, resulting in a detectable effect size.

A.1.1 Avoidance for Particular Subgroups

Another potential aspect could be that avoidance matters more for explaining plan risk preferences for certain groups of individuals. Figure A.2 presents results on relationship between avoidance and plan choice preferences. The relationship is more pronounced for certain population groups. For instance, non-white avoiders appear to switch to the generous plan when their health risk is 6 percentage points higher, male avoiders appear to switch when their health risk is 4.6 pp higher, and "Republican and Independent" avoiders switch when the health risk is 6.5 pp higher. Identifying these patterns of higher risk seeking behavior for particular population groups is plausibly helpful from a cost perspective. For instance, targeted interventions may ideally focus on addressing avoidance within certain groups of individuals for whom avoidance appears to significantly play a role for plan choice.

Figure A.2: Heterogeneity in How Avoidance Helps Explain Plan Choices

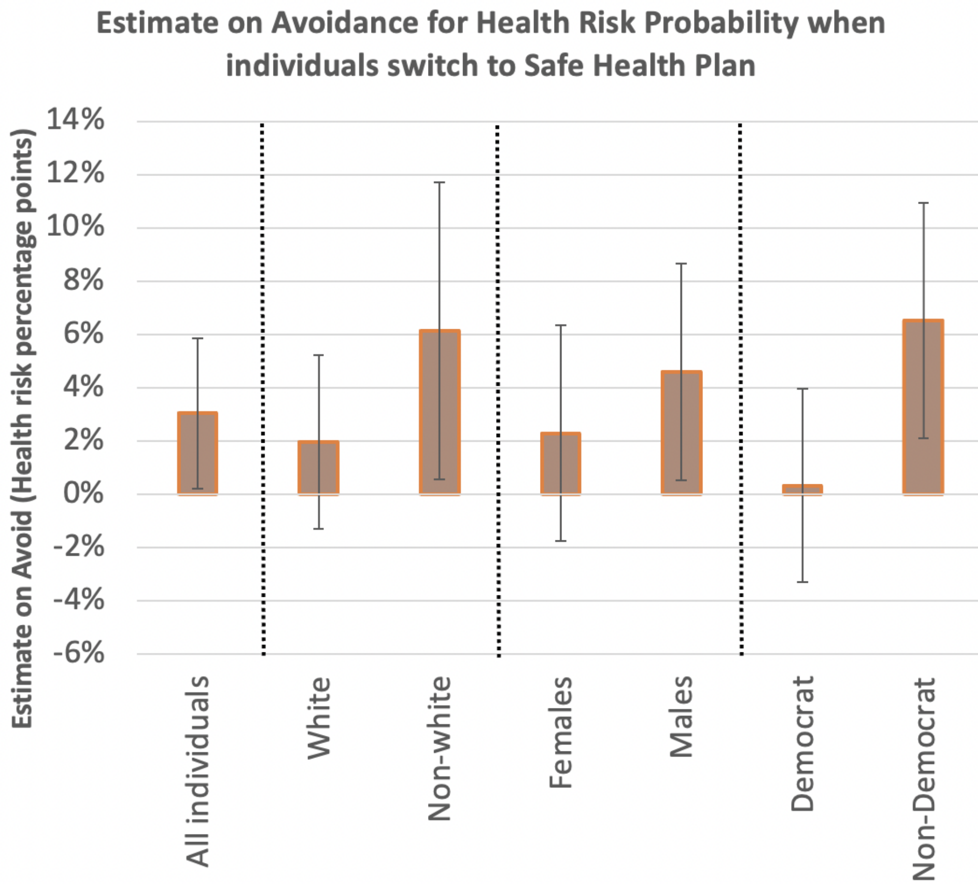

Note: The bars report the point estimate on the indicator for *Avoid*, where the outcome is "Switch to the generous plan" (i.e., equation 1.3.1. The 95% confidence intervals use standard errors on the estimate for each group of individuals listed on the bottom axis.

As table A.5 represents, avoiders in my sample who suggest that their primary reason for avoiding health information is due to "cognitive" reasons appear to wait until health risks are significantly higher than avoiders due to emotional and non-cognitive reasons. This subgroup of avoiders also appears to perform significantly fewer calculations and perform worse on the initial quiz about insurance terms.

Table A.5: Cognitive Avoiders versus Non-Cognitive Avoiders

	(1) Cognitive Avoiders	(2) Non-Cognitive avoiders	(3) Difference
N	149	222	
Preferences for Health Insurance			
Switch to generous plan	54.22	48.58	5.63**
			(2.47)
Switch to plan	19.99	16.30	3.69*
			(2.47)
Demographic Characteristics			
% Uninsured (health)	21.47	18.47	3.01
			(4.21)
% Earning Above $20,000	71.83	66.35	5.48
			(5.04)
% w/ at least college degree	17.44	19.81	−2.37
			(4.15)
% w/ 1 or less acq. hospitalized	67.11	61.26	5.85
			(5.09)
Average Age	37.01 $yrs.$	35.37 $yrs.$	1.63
			(1.10)
Knowledge and Choice Characteristics			
% Performed Calculations	19.46	33.78	14.32***
			(4.71)
Pre-treatment health risk beliefs (out of 100)	46.04	49.18	−3.14*
			(1.87)
Insurance term quiz 1 score (out of 4)	2.4228	2.7792	−.3564**
			(.1492)
Insurance term quiz 2 score (out of 4)	3.6375	3.7612	−.1236
			(.0777)
Post-treatment health knowledge (out of 100)	80.00	81.49	−1.49
			(3.49)

Note: Column 1 presents average characteristics of individuals who avoid the survey health information due to cognitive reasons (because they find extra information about health tiresome or tedious) alone or due to cognitive *and* emotional reasons and Column 2 presents average characteristics of individuals who avoid the survey health information *only* due to emotional aversion to health information . Column (3) reports the difference between avoiders and non-avoiders after controlling for bonus spread. *, **, and *** denote significance at the 10, 5, and 1 percent level.

A.1.2 Characteristics of Avoiders

Figures A.3, A.4, A.5 and A.6 provide a comparison across a number of characteristics and traits for avoiders and non-avoiders in my sample (control and treatment pooled).[1] I also include these observed characteristics as control variables in section 1.3.2.

Figure A.3 panel (a) suggests that self-defined avoiders are initially less knowledgeable about insurance terminology. Yet, avoiders learn at a faster rate due to catching-up: controlling for the bonus spread, OLS estimates suggest on average, the percent of avoiders who learn[2] insurance terms is 6.1% ($p < .05$) higher than non-avoiders. After teaching all individuals the term definitions, there is no discernible difference in scores between avoiders and non-avoiders. It is plausible this could be due to past or systemic tendencies of this group to avoid other forms of information that may impact their knowledge about health insurance.

In addition, Figure A.3 panel (b) demonstrates that avoiders tend to have less accurate beliefs about health risks (pre-treatment). Avoiders assigned to treatment still appear to incorrectly calibrate health risks relative to non-avoiders, despite being taught this information. Importantly, I assume these survey measures contain both concurrent and predictive validity: that is, I document that at least a portion of knowledge relating to health insurance competency and health risks is linked to avoidance in this survey. Yet, it seems likely that this knowledge gap would likely also translate to outside of an experimental setting.

Figure A.4 panel (a) sheds light on the fact that a significantly smaller proportion of avoiders earn more than \$20,000, \$35,000, and \$60,000 annually compared to non-avoiders and panel (b) suggests that avoiders are on average 1.9 ($p < .01$) years older than non-avoiders. Moreover, I document a strong positive association between respondents earning more than \$20,000, \$35,000, \$60,000, and \$100,000 annually and having health insurance ($p < .01$). As I report below in section 1.3.2, including income as a covariate in OLS regressions between avoidance and risk preferences for plans decreases the magnitude of point estimates on avoidance.

Figure A.5 documents several other attributes about avoiders: there is suggestive evidence

[1] *Baseline health risk accuracy* is a measure of how accurate an individual was about two initial questions that asked about how frequently individuals visited the doctor and the number of ER visits in the past year. *Beliefs about health risks* uses these same questions, but measures whether individuals thought these risks were low or high. *Health information knowledge (post-treatment)* is a variable representing how well a treated individual performed on a post-treatment quiz about the health information in the treatment. *All insurance terms correct (before)* is a measure of whether individuals got all insurance terms correct on their first try, *All insurance terms correct (after)* on their second try, *Learned insurance terms* is an indicator for whether an individual learned between quiz 1 and quiz 2, and *Willingness to learn* is an individuals' self-reported willingness to spend an additional minute learning about plans prior to selecting one.

[2] Learning is defined here as an individual getting all terms correct on the second insurance term quiz after getting at least one wrong on the first quiz.

that they are more likely uninsured in real life, know fewer acquaintances who have been hospitalized in the past year, are less likely to hold an advanced degree, and importantly, are less likely to have performed calculations during the survey. Notably, avoiders do not appear to be more likely to be male or to hold a high school degree. There is no statistically significant difference in the total survey time spent between avoiders and non-avoiders, and yet as Figure A.6 shows, these two groups allocate their time differently within the survey. Avoiders spend significantly less time during the portion of the survey where they are required to select plans, which could indicate a distaste for thinking about health insurance choice.

Another finding about avoiders is that they report placing more weight on the plan premium than non-avoiders (Table A.6) and they tend to underweight other plan characteristics such as the deductible, coinsurance, and out-of-pocket maximum, or post-premium plan payments. This result is congruent with existing research that finds that older beneficiaries overweight premiums (Abaluck and Gruber, 2011). Furthermore, this finding may suggest that the manner in which these plans characteristics are emphasized to customers could significantly matter.

Table A.6: Relative Importance of Plan Features For Selection

	(1) Avoiders	(2) Non-Avoiders	(3) Difference
Premium	.4633	.3855	$.0731^{**}$ (.0301)
Deductible	.3079	.3613	$-.0560^{*}$ (.0292)
Coinsurance	.0395	.0449	$-.0064$ (.0125)
Out-of-pocket maximum	.1158	.1340	$-.0130$ (.0206)
Post-premium payments	.4633	.5402	$-.0755^{***}$ (.0306)

Note: Respondents were asked what the most important plan characteristic was that factored into their decision when they were selecting plans in this survey (for the health simulation). This table shows the proportion of Avoiders and Non-Avoiders who selected each listed plan feature as being the most important to them for their choices. *Post-premium payments* refers to whether an individual selected deductible, coinsurance, or out-of-pocket maximum as their most important feature. Column 3 reports the difference controlling for bonus spread.

*, **, and *** denote significance at the 10, 5, and 1 percent level.

Figure A.3: Average insurance term knowledge and health risk knowledge across avoider types

(a)

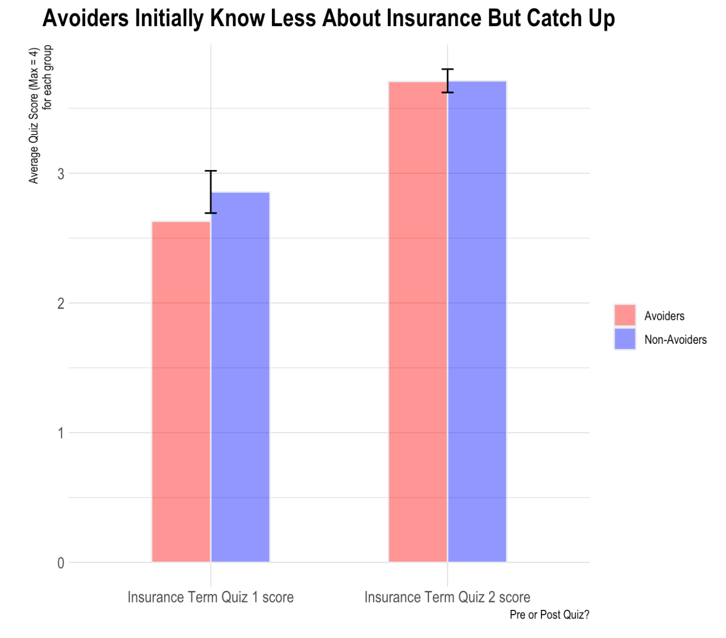

(b)

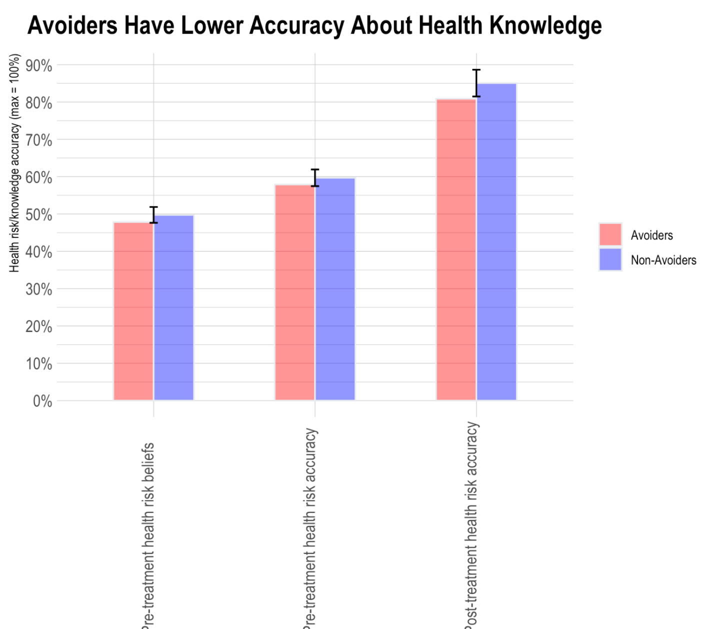

Note: Colored bars report average scores and quiz accuracy for avoiders versus non avoiders. I calculate the bar height for avoiders and non-avoiders after controlling the bonus spread in an OLS regression between avoid and the characteristic of interest. The 95% confidence intervals are based on standard errors on the "avoid" coefficient after controlling for bonus spread.

Figure A.4: Income and age characteristics of avoiders versus non avoiders

(a)

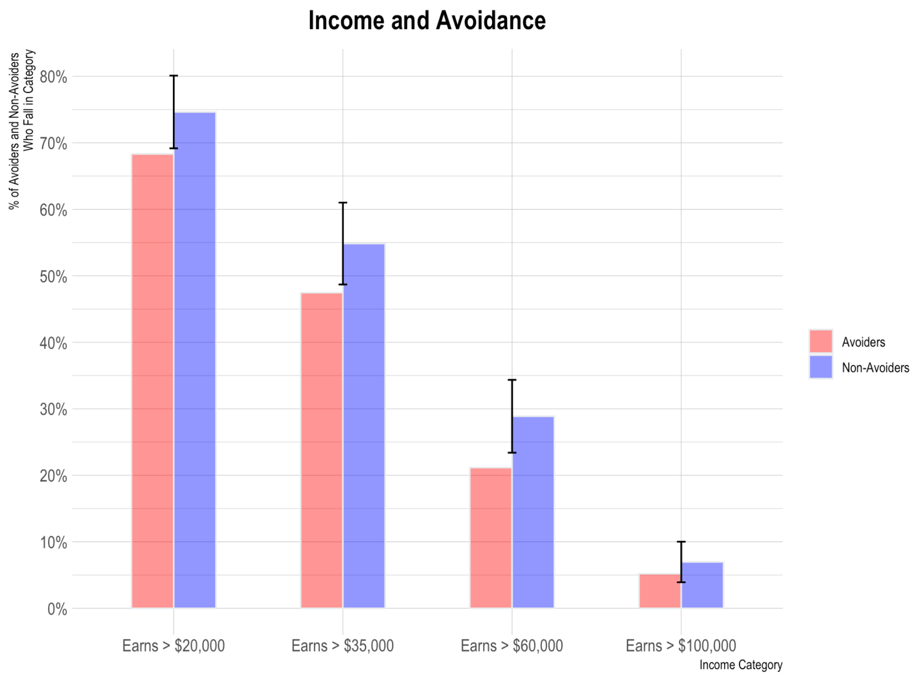

(b)

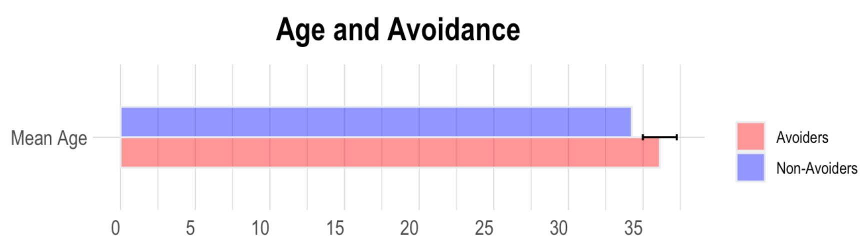

Notes: Colored bars report the average proportion of avoiders and non-avoiders in the sample who fall into each income bucket and the average age of avoiders versus non-avoiders. I calculate the bar height for avoiders and non-avoiders after controlling the bonus spread in an OLS regression between avoid and the characteristic of interest. The 95% confidence intervals are based on standard errors on the "avoid" coefficient after controlling for bonus spread.

Figure A.5: Proportion of avoiders and non avoiders in each characteristic category

(a)

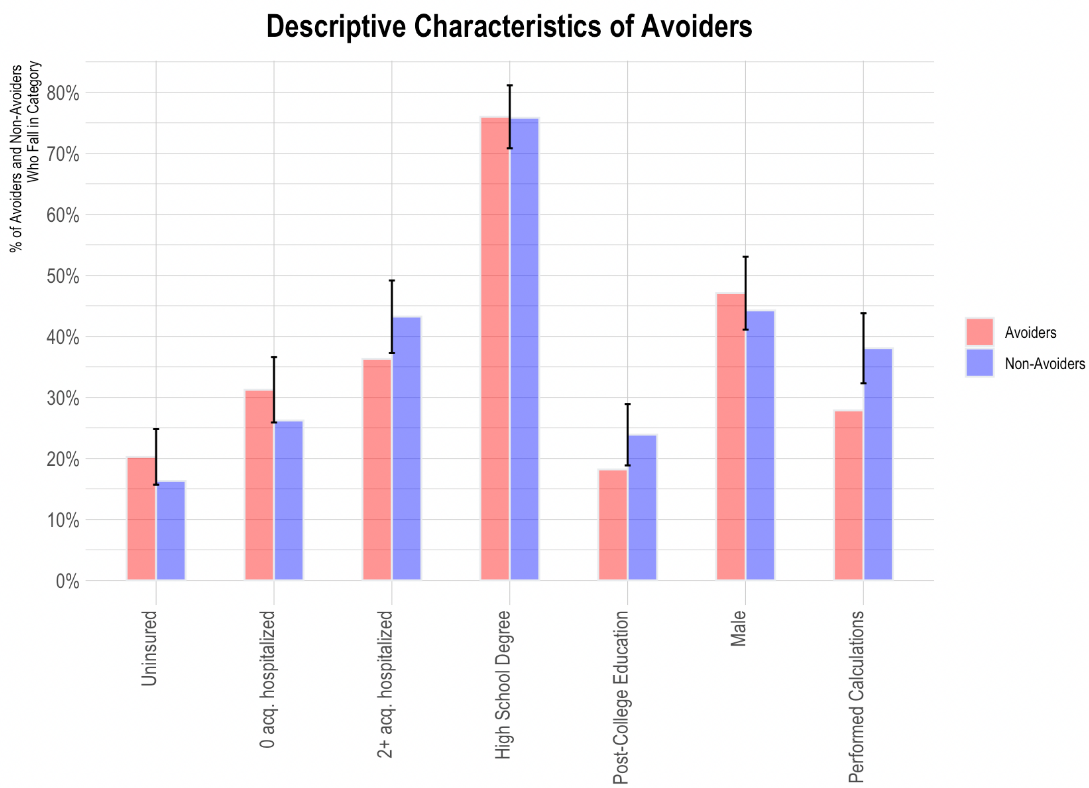

Notes: Colored bars report the average proportion of avoiders and non-avoiders in the sample who fall into each characteristic/trait category. I calculate the bar height for avoiders and non-avoiders after controlling the bonus spread in an OLS regression between avoid and the characteristic of interest. The 95% confidence intervals are based on standard errors on the "avoid" coefficient after controlling for bonus spread.

Figure A.6: Time spent on different survey parts

(a)

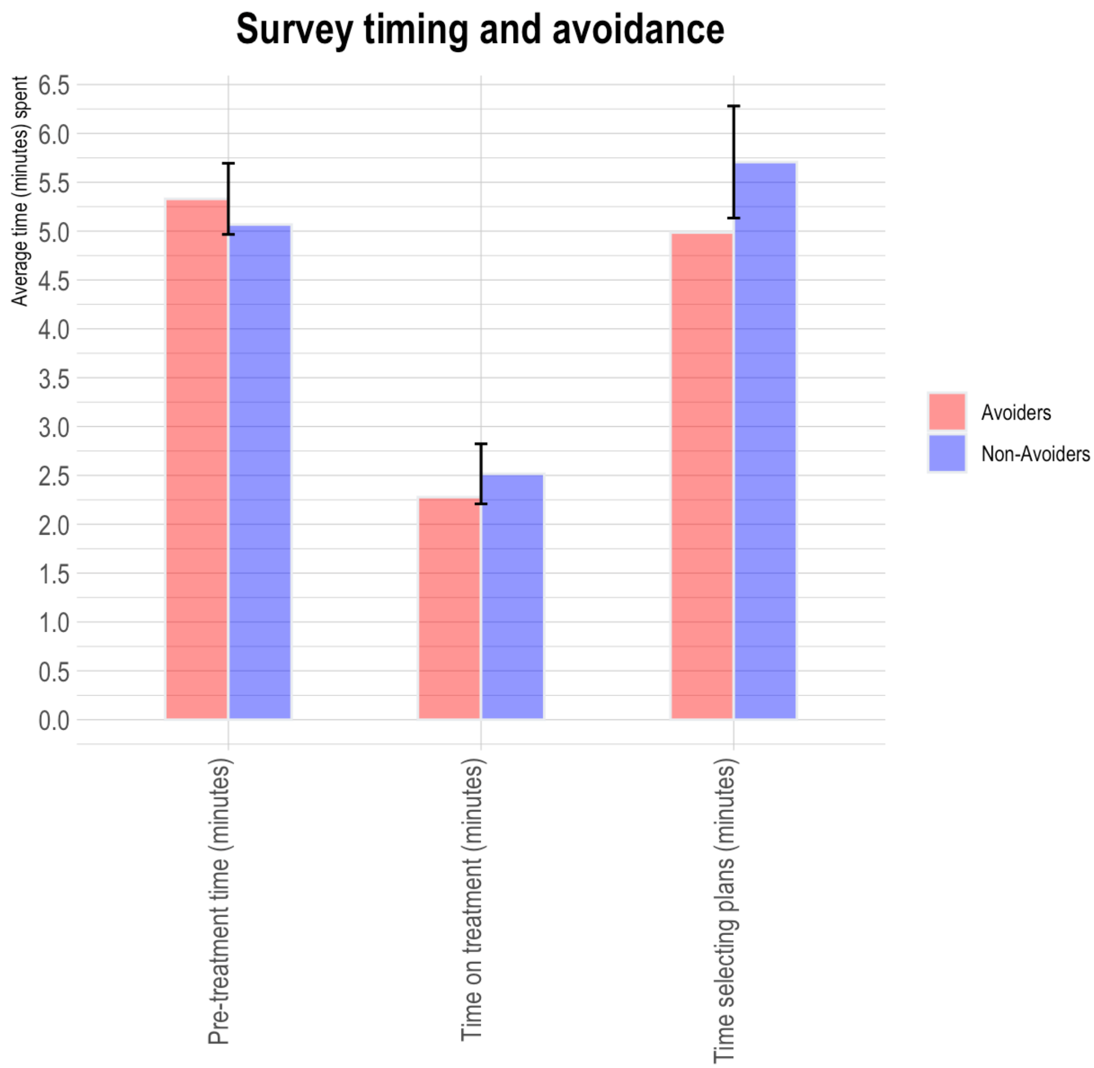

Note: Colored bars report the average minutes that avoiders and non-avoiders in the sample spent on a few parts of the survey. "Time on treatment" bars only report means for avoiders and non-avoiders who were assigned to treatment. For all measurements, I calculate the bar height for avoiders and non-avoiders after controlling the bonus spread in an OLS regression between avoid and the characteristic of interest. The 95% confidence intervals are based on standard errors on the "avoid" coefficient after controlling for bonus spread.

A.1.3 Avoiders Prefer Riskier Plans

Table A.7 presents results from equation 1.3.1 when I separate treatment and control individuals. The results suggest that avoidance is associated with riskier plan choices, regardless of treatment assignment. This is most noticeable from observing the point estimate on the "count generous plan" outcome measure.

Table A.7: Control for Differences in Information Content to Study Whether Avoidance as a Characteristic Helps Explain Plan Choices

	(1) Switch generous	(2) Switch to insurance	(3) Count no insurance	(4) Count high-deductible	(5) Count generous
Only treatment group individuals					
No controls	3.9874**	2.8736*	.3523*	.2587	−.6110**
	(2.0203)	(1.5468)	(.2020)	(.2252)	(.2443)
Control for covariates X_i	3.0621	1.6159	.1385	.3651	−.5036**
	(2.1030)	(1.5505)	(.2020)	(.2398)	(.2538)
Only control-group individuals					
No controls	2.3496	1.4207	.2469	.2904	−.5374**
	(1.9830)	(1.4033)	(.1837)	(.2259)	(.2431)
Control for covariates X_i	2.1262	1.3441	.0834	.3184	−.4019*
	(2.0618)	(1.4094)	(.1802)	(.2351)	(.2457)

Note: I separate out "only treatment-group" and "only control-group" individuals. The first row in each cell represents the point estimate on "Avoid" from from only controlling for the bonus spread, and the second row in each cell represents the estimate from including demographic, survey, and knowledge/learning attributes. It appears that there is a tendency for information avoiders to select more risky health plans and less generous plans, even after separating them out based on treatment status. There is suggestive evidence of this both for "only treatment" and for "only control" avoiders.

*, **, and *** denote significance at the 10, 5, and 1 percent level.

A.2. Decomposition

I further decompose changes in the magnitude of β_1 by studying the relative impact of each of the added covariates. I find that when I only include income, education level, and initial insurance term quiz score results as controls, the coefficient on avoid drops to a value even smaller than that reported in column 3 with all covariates (Table 1.3). With the inclusion of only these three covariates, avoiders switch to the generous plan when their health risks are 2.36 pp higher than non-avoiders, with a p-value equal to .10. This suggests that income, education, and initial insurance term knowledge appear to be highly informative in helping explain avoidance patterns on plan risk preferences.

Table A.8 presents results from a Blinder-Oaxaca decomposition in the difference in the coefficient on avoid for each of the outcome variables. Of particular interest is row (2), which presents results for the portion of the mean difference in risk preferences between avoiders and non-avoiders attributable to different distributions between the two groups based on observable characteristics. For four out of five of the main outcomes, the explained component is small and statistically insignificant when compared to the unexplained component. This is suggestive of the fact that a portion of the magnitude of the coefficient on avoidance likely corresponds to unobserved traits. For instance, it could be the case that information avoiders are significantly riskier along other dimensions of their lives, and this is reflective of a deeper personality trait of a willingness to take on higher risk. Yet still, there may be social and/or cultural biases that relate to both a desire to avoid salient information as well as insurance choice preferences. Still another reason may be that there are other dimensions of knowledge relating to insurance and hospital bills, of which information avoiders lack pre-existing knowledge.

Table A.8: Blinder-Oaxaca Decomposition, Portion of Avoidance Explained by Covariates

	(1) *Switch generous plan*	(2) *Switch plan*	(3) *Count generous plan*	(4) *Count high-deductible plan*	(5) *Count no insurance*
[1] Difference (Avoid - Non-avoid)	2.9178**	1.8027	−.5462***	.3209*	.2253
	(1.4198)	(1.1359)	(.1783)	(.1704)	(.1458)
[2] Explained	.2531	.6164	−.1131	−.0416	.1548**
	(.5733)	(.5125)	(.0718)	(.0546)	(.0673)
[3] Unexplained	2.6647*	1.1862	−.4330**	.3625**	.0704
	(1.3693)	(1.0415)	(.1704)	(.1698)	(.1363)

Note: Results are from a Blinder-Oaxaca decomposition that uses the coefficients from a pooled model over both avoiders and non-avoiders as the reference coefficients. Additionally, the decomposition includes controls for group indicators. This table reports the results from decomposing differences in outcomes between avoiders and non-avoiders by explained and unexplained.
*, **, and *** denote significance at the 10, 5, and 1 percent level.

A.3. Additional Details about the Experiment

A.3.1 Internet Usage and Health Insurance

A potential concern underlying my approach is that my experiment emphasizes the relationship between avoidance and health insurance choices among a population of individuals who are "internet-loving." Intrinsically, any individual taking my survey is, at the very least, connected to information through either the internet or their cell-devices from data plans, because these are the two platforms through which they are allowed to complete my Prolific survey. This would skew my sample towards "more-informed" individuals. So, any conclusions I draw about avoidance in this study would consequently be internally valid among individuals with an affinity for using the internet. I use data from the 2020 American Community Survey to study the associations between health insurance coverage and internet accessibility.

A few facts suggest that validity may be a relatively small concern. Overall in 2020, 91.8% of individuals overall subscribed to the internet using a cellular data plan and 92.3% of individuals in the U.S. had access to the internet. Furthermore, 89.8%[3] of individuals without health insurance had access to internet and 93.2% of those with health insurance had access to internet. The difference in the proportion with access to internet, albeit small, is statistically significant ($p < .001$). This could be suggestive of the fact that internet access, to a degree, plays a role in helping explain underinsurance. One possibility for why this may be the case, among several others, could be that not having internet access is reflective of avoidance because that it suggests a distaste for seeking out or being exposed to certain types of information. However, if anything, this may potentially result in a downward, rather than upward bias, of the point estimates I report between the association between avoidance and underinsurance for health.

[3]11% of subsidy-eligible uninsured people did not have access to internet access at home, compared to 6% of the non-elderly US population. These numbers are higher for the free bronze eligible uninsured (13%) and the free silver eligible uninsured (15%) (Cox and McDermott, 2021).

A.3.2 Details about experimental implementation

In my experiment, I ask respondents to select between health plan options for different hypothetical probabilities that they will experience a negative health issue in the following year. However, this is potentially a challenging task for individuals. Some people who have not explicitly considered their health risks while they selected health insurance in the past or may have a low level of numeracy. Moreover, a risk in soliciting health plan preferences on an online experiment is that respondents have no direct monetary incentive to select health plans in the scenario that they would actually select during a real open enrollment period. While these are valid concerns that are difficult to fully address, I explore a few ways to mitigate these concerns in the survey design.

First, to address the concern that selecting plans is potentially challenging, prior to asking individuals about their health insurance preferences, I present them with key definitions of the following health insurance terms: deductible, premium, coinsurance, and out-of-pocket maximum (after an initial quiz on these terms to measure learning). On the following survey page, to keep participants engaged and make sure they have learned the terms correctly, I ask respondents to match the health terms to their respective definitions. These four financial features appear in the hypothetical health plans I present to respondents in the simulation. In addition to teaching individuals these terms and measuring learning rates for respondents, this also provides me with a measure of the attentiveness of the participants during the experiment.

Second, a concern could be that in the health insurance simulation, individuals are too driven by giving a "correct" answer the health plans they are required to select for the different health risk probabilities, as opposed to providing us with their preferences. Consequently, prior to the preference solicitation task, I tell respondents: "This is an opinion survey with no right or wrong answers. Please answer the following question based on on your own preferences."

Once respondents have completed the simulation asking them for their health plan preferences, I include an open-ended question for respondents that asks them to qualitatively describe how they selected the health plan they preferred for the hypothetical chances of the major health event in the scenario with the following prompt: "Please explain what information or past experiences you drew upon, or any other reasons that explain why your selections for the health insurance plans above reflect how you would actually select for insurance. Particularly thoughtful and insightful responses will receive an additional bonus of $0.50." Ideally, this prompt will incentivize individuals who "rushed" through the health simulation to return to choose responses that are more in line with their actual preferences. Moreover, the degree to which individuals respond to this question indicates provides a measure of care they took in going through the survey for selecting an insurance plan in the scenario. I find that a significant proportion of respondents thoroughly fill in this question bank in the survey, which suggests that this is an effective way to ensure individuals are thinking carefully and retrospectively about their health insurance preferences during the simulation. Moreover, survey responses with short or empty responses can be filtered out for measures of robustness. I use short answers responses for a text analysis.

A final concern is that, contrary to how my experiment models behavior, individuals may not necessarily view their health risks in terms of a "lottery" or mentally represent health risks in terms of probabilities when they actually select for health plans. Instead, it seems plausible that individuals might incorporate rougher signals/estimates of their health risks when determining the most optimal health plan to select. Moreover, some individuals may already have a broad sense of their health risks for the upcoming year. While this is a valid concern, I take steps to motivate the use of a health lottery by having individuals think of an illness that is similar, but not the same as, COVID-19.

A.3.3 Additional data collection procedures

In the summer of 2018, there was an increased concern that online participant platforms, particularly Amazon's Mechanical Turk, were susceptible to private servers fraudulently gaining access to studies (Kennedy et al., 2020). Social science researchers since then have attempted to allay concerns about quality of responses in a few ways, including attentiveness checks soliciting requiring written responses, which I employ in the experimental design. I also follow previous economic papers that rely on online experiments (e.g. Fisman et al. (2020)) to screen out potential bots by making it required that all participants fill out a "captcha" (an identification task where individuals have to select boxes that have images embedded in them, which are difficult for computers to interpret).

A.3.4 Health Insurance Choice Simulation

The experimental health simulation task is designed to simulate a particular years' open enrollment period for health insurance. The health plans are designed such that the PPO Gold is a relatively more "generous" plan, because, despite having a high premium, there is a low deductible and low co-insurance. The PPO Bronze plan is relatively less generous (low premium, high deductible, and high co-insurance). Economists have used a similar analytical approach (Holt and Laury., 2002) to identify risk aversion by studying at which probability of risk individuals switch from the generous option to the high-deductible one. This health lottery setup allows me to analytically deduce a coefficient of risk aversion to establish a

Figure A.7: Health Insurance Plan Financial Characteristics for the Lottery

	Monthly Premium	Deductible	Coinsurance	Out of pocket max
PPO Bronze	$265.83	$8,000	20%	$8,550
PPO Gold	$480.17	$0	5%	$8,150

Note: This figure presents the 4 financial characteristics that are given to respondents to help them select between the PPO Bronze and PPO Gold plans in the health simulation in the survey.

Figure A.8: Health Insurance Plan Total Expected Costs

	Sick state	Healthy state
PPO Bronze	-$11,739	-$3,189
PPO Gold	$7,012	-$5,762

Note: This table presents potential payoffs in sick and healthy states, under the PPO Bronze and PPO Gold plans presented to survey respondents in my survey instrument.

baseline risk level.

Respondents may select between PPO Bronze, PPO Gold[4] (see characteristics from Table A.7), or no insurance for any of the 11 health risk probabilities, for a potential adverse health event (whose probability of occurring ranges from 0% to 100% in increments of 10%). The plans are constructed such that, for a medical issue costing $25,000, survey respondents could apply the definitions of a deductible, coinsurance, premium, and out of pocket maximum (which they are taught in the beginning of the experiment) to find their total medical expenditures for the year under each plan for each health state (sick or not sick). The total cost (including premium, deductible, and co-insurance) for each state under each hypothetical health plan is shown below in Table A.8.

The payoffs from Table A.8 facilitate the construction of a "health risk" lottery based

[4]The Bronze plan is also presented as a "low premium, high deductible plan" and the Gold plan is also referred to as a "high premium, low deductible plan" to survey participants. I intentionally include the terms "Bronze" and "Gold", because this is a common decision aid individuals rely on even when selecting between insurance plans, and even though individuals are told the specific financial characteristics, this may be a simpler way for them to remember the plans they are choosing between in the health lottery. I also deliberately make it so that both plans are preferred provider organization (PPO), because this ensures that the health lottery is only based on the financial characteristics provided (premium, deductible, coinsurance, and out-of-pocket maximum), which is important for converting these into lottery payoffs as presented in Table A.8.

Figure A.9: CRRA Bounds

# PPO Gold Choices	Wealth: $15k	Wealth: $50k	Wealth: $500k	Risk classification
0-1	-3.52 > r	-18.7 > r	-212.6 > r	Extremely risk loving
2	-2.47 > r > -3.52	-13.4 > r > -18.7	-152.7 > r > -212.6	Highly risk loving
3	-1.66 > r > -2.47	-9.18 > r > -13.4	-105 > r > -152.7	Very risk loving
4	-.97 > r > -1.66	-5.43 > r > -9.18	-62.3 > r > -105	Risk loving
5	-.31 > r > -.97	-1.79 > r > -5.43	-20.5 > r > -62.3	Risky tendency
6	.36 > r > -.31	2.1 > r > -1.79	23.9 > r > -20.5	Risk neutral
7	1.13 > r > .36	6.6 > r > 2.1	76.7 > r > -23.9	Risk averse
8	2.2 > r > 1.13	13.2 > r > 6.6	153.5 > r > 76.7	Very risk averse
9-10	r > 2.2	r > 13.2	r > 153.5	Stay in bed

Note: This figure presents different potential upper and lower bounds based on three potential initial wealth holdings for an individual who selects the PPO Gold Plan in the experiment a specified number of times (corresponding to rows). These calculations for bounds are based on a CRRA utility: $U(x) = \frac{x^{(1-r)}}{1-r}$ for three potential starting wealth amounts. Furthermore, these calculations assume that there individuals consistently select plans (sequentially).

on 11 ranked health risk probabilities ranging from 0% to 100%, in increments of 10%.[5] Following Sydnor (2010), Table A.9 allows for differing levels of lifetime wealth, x, to calculate lottery payoffs.

As Sydnor (2010) discusses, the choice of lifetime wealth has different interpretations. For individuals in my sample, $500,000 is a conservative estimate for total lifetime wealth. Lifetime wealth of $50,000 and $15,000 provide comparable estimates of the CRRA parameter for measuring annual or quarterly income, which facilitates comparison of risk aversion bounds in the health insurance context to other insurance context. Unlike Sydnor (2010), it is likely more appropriate in the health insurance context to benchmark lifetime wealth against annual or quarterly income because the decision to insure is (whether passively or actively) made on an annual basis by consumers. If one wanted to study risk aversion in this context relative to total lifetime wealth, we would have to impose assumptions on the flow of payments over several years into the future based on future expected plan prices. This type of calibration for risk aversion would likely make sense when studying the subset of consumers driven by inertia, locking into a plan for several years; however, my study does not focus on separating out passive versus active insurance consumers.

[5] When designing lotteries, there is an inherent tradeoff between power and feasibility. For instance, a health lottery could be constructed in sparser increments (e.g. increments of 20%), but a concern is that this would be *too* crude and would not allow me to detect any substantial variation in preferences between information avoiders and information seekers. Alternatively, I could have construct a health lottery with much more frequent health risk probabilities (increments of 5%, for instance), but a concern with this approach is that it would be too taxing a task for individuals to fill in 20 lottery preferences.

A.3.5 Pre-specified Decision Rules for Outcome Measures

Inconsistent choices

An aspect that I discussed in my pre-analysis plan was that certain individuals may display inconsistent choice patterns by switching non-monotonically. For example, an individual would select *PPO Gold* for low health risks and then switch to *PPO Bronze* for high health risks. I pre-specified that I would drop these individuals from my outcomes. I find that 90 individuals, or 5.5% of the sample, displayed this form of inconsistent non-monotonic switching.

My sample does not include these individuals. I was concerned by the potential that individuals with this choice pattern would either represent individuals trying to manipulate their outcomes or not being truthful and not taking the lottery seriously. I validated that this was the correct decision to make by looking at the association between the length of blurb individuals were incentivized to write explaining why they selected their insurance choices for the 11 lottery options and whether or not an individual displayed this inconsistent choice pattern. I find on average, individuals with the inconsistent choice pattern had blurbs that were 162 characters shorter ($p < .01$) than non-inconsistent individuals.

Switching Patterns

I hypothesized that a majority of survey respondents would switch from choosing no insurance to PPO Bronze (the high-deductible plan) before selecting PPO Gold (the generous plan) for high probabilities of experiencing the health issue. Yet, a minority of individuals in the pilot survey demonstrated the pattern of jumping directly from selecting no insurance to PPO Gold, skipping the PPO Bronze plan altogether.

In my pre-analysis plan, I decided on the decision rule that if more than 5% of individuals in the sample switched directly from no insurance to PPO Gold, my two primary outcome variables would be a measure of (1) health risk probability when an individual switches to *any* insurance plan (either high-deductible or generous), and (2) health risk probability when an individual switches from any less generous plan option to the most generous plan option (PPO Gold). I find that 120 respondents (7.7% of my sample) display this form of switching, which passes my decision rule.

A.4. Heterogeneity and Additional Analyses

Table A.9: T1 (Facts Only) Effect on Health Risk Switch Point for Health Insurance

	Switch to Generous Plan		Switch to Plan	
	(1)	(2)	(3)	(4)
Treatment	.3544	$-.86428$	$-.1277$	-1.2137
	(1.7264)	(1.9568)	(1.2624)	(1.4362)
Avoid	3.5630**	2.3496	2.4320**	1.4207
	(1.7237)	(1.9524)	(1.2305)	(1.3858)
Treatment*Avoid		5.4836		4.7501
		(4.1507)		(3.0036)

Note: Columns 1 and 3 report treatment effects without including an interaction between *Treatment* and *Avoid*. Columns 1 and 2 correspond to treatment effects for primary outcome 1 (the health risk probability when an individual switches to the generous insurance plan), and 3 and 4 are for outcome 2 (the health risk probability when an individual switches to an insurance plan).
*, **, and *** denote significance at the 10, 5, and 1 percent level.

Table A.10: T2 (Facts and Images) Effect on Health Risk Switch Point for Health Insurance

	Switch to Generous Plan		Switch to Plan	
	(1)	(2)	(3)	(4)
Treatment	1.1473	1.1532	1.2852	1.2591
	(1.3091)	(1.4816)	(.9795)	(1.1150)
Avoid	2.3389	2.3496	1.4656	1.4207
	(1.5457)	(1.9821)	(1.1402)	(1.4623)
Treatment*Avoid		$-.0273$		.1146
		(3.1687)		(2.3368)

Note: Columns 1 and 3 report treatment effects without including an interaction between *Treatment* and *Avoid*. Columns 1 and 2 correspond to treatment effects for primary outcome 1 (the health risk probability when an individual switches to the generous insurance plan), and 3 and 4 are for outcome 2 (the health risk probability when an individual switches to an insurance plan).
*, **, and *** denote significance at the 10, 5, and 1 percent level.

A.5. Conceptual Model

I provide a simple model that distinguishes between information about health risks and consequences, two information forms relevant in health insurance choices. Individuals take in two signals that contribute to their knowledge about their health. The first is an individual-specific component of information about an individuals' health *risks*, $\pi^i \in \{0,1\}$, or information about the likelihood of experiencing an illness. The second component is information on the *consequences* of an illness/not seeking out appropriate care on individuals' long-term health outcomes such as mortality, $\pi^c \in \{0,1\}$.[6] Either component, π^i or π^c, will take on a value of 0 when an individual avoids information and a value of 1 when an individual sees the information.

Individuals will place different emphasis on the risk aspect or the consequence component of information. For instance, it may be that an individual only focuses on information about their own health risks, only on consequences, or a combination of both aspects. Individuals weight the individual-specific risk component by $\alpha_i \in [0,1]$ and the consequence component by 1-α_i. Together, these will form an individual-specific information parameter for person i as follows:

$$\pi_i = \alpha_i \pi_i^i + (1 - \alpha_i)\pi_i^c \tag{A.1}$$

I hold fixed the individual component of information and I focus on studying avoidance of the health information on consequences of being underinsured and experiencing health issues on overall outcomes such as mortality. In this setting, π_i^c may take on a value of 0 or 1, depending on whether an individual avoids group information or not. In different experimental settings or in the real-world, the definition of avoiders could be expanded to include those who avoid the component of information relating to their own individual risks. I hold individual information constant across all respondents by providing all of them the same information on the likelihood they will experience a health issue in the upcoming year when they select between health insurance plans. Consequently, in this setting, $\pi_i^i = 1$ for all individuals.

Individuals will then select between health insurance plans in the experiment by considering the information they have, relying on an "information mapping" function $f_i(\cdot)$. When individuals do not avoid health consequence information, $\pi_i^c = 1$. An individuals' insurance decision will reflect full information available to them and they will select health insurance

[6] An example of individual-specific risk information is "Based on your demographics, there is an X% chance something will happen to your health" and an example of health consequence information is "Individuals who fail to seek out care or take out insurance will have an X% chance of being hospitalized and a Z% chance of passing away".

level H* = $f_i(1)$. However, for avoiders of health consequence information, $\pi_i^c = 0$, and they will have a level of information equivalent the amount they weight their own individual health risks by: $\pi_i = \alpha_i$. If it is the case that avoiders only focus on individual health risk information and do not rely on health consequence information, their $\alpha_i = 1$ and so H* = $f_i(1)$, which will be equivalent to the case where they did not avoid health consequence information.

However, the question this study attempts to shed light on is whether avoiding health consequence information potentially results in a different level of health insurance. This might occur if an avoider relies on both individual and group components of information, then $\alpha_i < 1$ and so $\pi_i < 1$, and the level of insurance an avoider selects equals $f_i(\alpha_i)$. There are three different cases that may hold true when $\alpha_i < 1$, which depends on the individual information mapping function:

Case 1: Even after avoiding health consequence information, individual i selects the same level of health insurance: $f_i(1)$ - $f_i(\alpha_i) = 0$. In other words, the foregone information does not impact an individual's plan coverage.

Case 2: Individuals who avoid health consequence information end up with *more generous* insurance coverage, thereby overinsuring for medical care: $f_i(1)$ - $f_i(\alpha_i) < 0$. A potential underlying explanation could be that individuals may tend to believe health consequences are more common than they actually are, and so additional information may lower this subjective belief.

Case 3: Avoidance results in person i underinsuring: $f_i(1)$ - $f_i(\alpha_i) > 0$. This is potentially because these individuals may have underestimated health consequence risks by avoiding relevant health information.

Experimental variation in treatment assignment helps me study whether case 3 holds empirically. As section 1.2 details, my experiment classifies two types of individuals: avoiders and non-avoiders. There are N_1 avoiders and N_2 non-avoiders in the sample, and I randomly assign both avoiders and non-avoiders to the treatment where they read through health consequence information and will take out insurance coverage. Based on the above, if an individual is assigned to see health consequence information, $\pi_i^c = 1$, so $\pi_i = 1$ and they will select a health insurance level of $f_{i1}(1)$. However, if an individual is assigned to view non-health related information in my control group $\pi_i^c = 0$, so $\pi_i = \alpha_i$ and they will take out insurance coverage of $f_{i0}(\alpha_i)$ (allowing for α_i to equal 1). I hypothesized that the treatment effect of assignment to health consequence information would be greater for avoiders than for non-avoiders and significantly greater than zero. The first term in equation A.2 corresponds to the average treatment effect for avoiders and the second term is the average treatment effect for non-avoiders, under this hypothesis:

$$\underbrace{\frac{1}{N_1}\sum_{i=1}^{N_1} f_{i1}(1) - f_{i0}(\alpha_i)}_{\text{Average treatment effect for avoiders}} \quad > \quad \underbrace{\frac{1}{N_2}\sum_{i=1}^{N_2} f_{i1}(1) - f_{i0}(\alpha_i)}_{\text{Average treatment effect for non-avoiders}} \quad \geq 0 \tag{A.2}$$

A.6. Additional Experiment Details

Figure A.10: Experimental Design

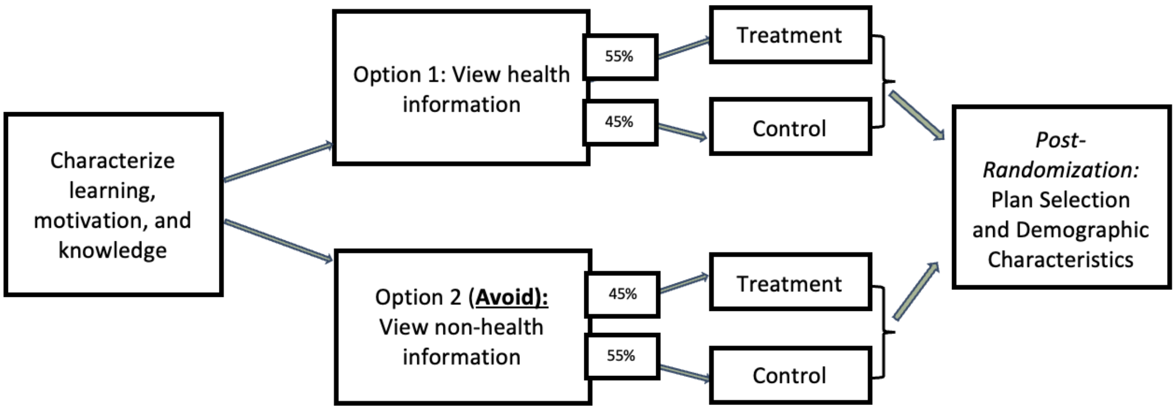

This flowchart illustrates the experimental design. The first part of the experiment elicits individuals' existing insurance term and health risk knowledge, teaches individuals the terms, and quizzes them on the terms. Respondents then express their preferences for avoiding or viewing salient health information. I induce random variation by assigning 55% of respondents to their preferred option. For individuals assigned to treatment (health information), 25% are assigned to facts only and 75% are assigned to facts and images. After individuals go through treatment or control material, they are redirected to select health insurance plans in a hypothetical health insurance simulation. Finally, the experiment concludes after individuals respond to demographic questions.

A.6.1 Additional Excerpts from Online Survey

Figure A.11: Teaching and quizzing individuals about insurance terms

To give us a sense of your familiarity with health insurance, we will quiz you on a few common health insurance terms.

We will give you 2 attempts for the quiz. You will take the quiz for a first time on this page. Then, we will then teach you the definitions of the terms. You will then take the quiz again after you have learned the terms. You will receive a bonus of $.10 if you get all the terms correct on <u>either of the two attempts</u>. Try your best on both quizzes because this will maximize your chances of getting the bonus.

Timing

These page timer metrics will not be displayed to the recipient.

First Click	0 seconds
Last Click	0 seconds
Page Submit	0 seconds
Click Count	0 clicks

Question 1: The most you have to pay for covered services in a year, after which your health plan pays 100% of the costs.

Premium

Deductible

Coinsurance

Out of pocket maximum

Question 2: The percent of costs of a covered health care service you pay after you've paid your deductible.

Premium

Deductible

Coinsurance

Out of pocket maximum

Figure A.12: Teaching and quizzing individuals about insurance terms

Question 3: The amount you pay for covered health care services in a plan year before your health insurance plan begins to contribute.

- Premium
- Deductible
- Coinsurance
- Out of pocket maximum

Question 4: The monthly amount you pay for your health insurance plan.

- Premium
- Deductible
- Coinsurance
- Out of pocket maximum

Figure A.13: Teaching and quizzing individuals about insurance terms

Below, we will provide you with the terms from the previous page:

Please carefully review the following key health insurance terms before continuing:

Premium: The monthly amount you pay for your health insurance plan (**e.g. $350/month**).

Deductible: The amount you pay for covered health care services in a plan year before your health insurance plan begins to contribute (**e.g. you pay $6000, after which the insurance plan starts paying**).

Coinsurance: The percent of costs of a covered health care service you pay *after* you've paid your deductible (**e.g: a coinsurance of 10% means your insurance covers the other 90% of medical costs**).

Out-of-pocket maximum: The most you have to pay for covered services in a year in the form of deductibles and coinsurance (**e.g. $8000**), after which your health plan pays 100% of the costs **(Note: this is separate from the premium)**.

Figure A.14: Treatment (with images)

Click to reveal the percent of adults in the United States had a visit with a doctor or other health care professional in the past year:

Reveal information

84.9% of adults in the United States.

The total population of the United States is around 320 million people. Click to see the number of annual visits to the emergency room in the United States:

Reveal information

There are **130 million annual visits** to the emergency room in the United States.

Figure A.15: Treatment (with images)

Annually, how many people are hospitalized for a day or longer in the United States in an intensive care unit?

Reveal information

30 million people are hospitalized annually. Without seeking out timely care in a hospital, you could possibly incur life-changing medical issues, or even worse, **you may die**.

Figure A.16: Treatment (with images)

Click to reveal the annual number of deaths in the United States due to being underinsured:

Reveal information

There are nearly **45,000 deaths in the U.S. annually due to being underinsured** (either from having no health insurance or having a plan that does not adequately cover out-of-pocket expenses).

Click to reveal the top reasons people are hospitalized:

Reveal information

Some of the top reasons that people are hospitalized include: **pneumonia, stroke, chest pain, infections, and fractures. COVID-19** is another potentially serious illness that may result in life-threatening hospitalization or death. These are all issues that can occur to anyone at any given time.

Figure A.17: Treatment (with images)

1 in 5 (20%) uninsured adults say that they went without needed health care in the past year because of cost, compared to 3% of adults with private insurance. This can have serious consequences on your health.

Consider Susan Finley, a woman in the United States who worked in retail. She was recovering from pneumonia, but when she returned to work in May 2016, Susan lost her job and also lost her health insurance coverage.

During the next three months, Susan avoided going to see the doctor or to a hospital about her flu-like symptoms, because she did not have health insurance.

Click to see what happened to Susan:

Reveal information

Susan's parents went to check on her, but "they couldn't get into her apartment. They got the landlord to open it up, **went in and found that she had passed away.** It came as a complete surprise to everybody," said her son.

Figure A.18: Treatment (with images)

In reality, health insurance coverage can be complicated because it may depend on several factors including a plan's network coverage and intricate negotiations between providers and health plans. Despite this, most health plans will cover **essential health benefits**.

Click to see what health insurance plans typically cover:

Reveal information

1. Urgent care visits
2. Emergency transportation and room services
3. Hospitalization (like surgery, transplants, and overnight stays)
4. Regular doctors' visits
5. FDA-approved prescription drugs
6. Laboratory services (Diagnostic tests, MRIs, CT scans)
7. Pregnancy, maternity, and newborn care
8. Mental health and substance abuse disorder services

Click to learn about what is not covered by most health insurance plans:

Reveal information

1. Cosmetic procedures (for instance: plastic surgery, tattoo removal)
2. Prescription drugs for off-label uses (non-FDA approved drugs)
3. Dental care

Figure A.19: Control

You have been assigned to **Option 2**.

We will teach you how to play a game called Sudoku. There are a few questions included as part of the walk-through. You are not required to get the questions correct, but they may help you with learning the game.

This 2x2 sudoku game starts with four blocks (upper right, lower right, upper left, and lower left quadrants). You are also provided with four starting letters, one for each block.

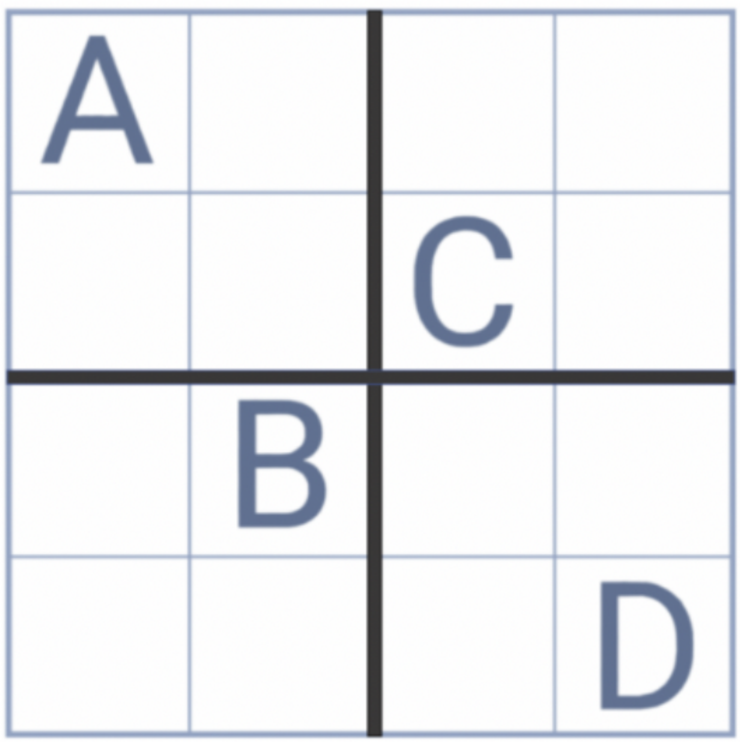

The goal is to fill in all of the squares using only the letters A, B, C, or D. There are a couple key rules to fill in the rest of the game:

Rule 1: The letters A, B, C, and D must occur only once in each row and column.

Rule 2: The letters A, B, C, and D can only occur once in each block of 2x2.

Rule 3: You are not allowed to delete or change the four starting letters (those are fixed).

Figure A.20: Health Insurance Choice Simulation

The following is an opinion question with no right or wrong answers. Please answer the following question based on your own preferences.

In 2020 and 2021, we experienced a global pandemic. There was uncertainty for many about whether they would be hospitalized with COVID-19. In this scenario, we want you to consider different probabilities that you may fall sick with a serious health issue in the next year. This health issue would require hospitalization, testing, rehabilitation, and expensive prescription drugs, resulting in a bill costing you **$25,000** if you are uninsured. These would be the only medical costs you will incur next year.

The open enrollment period is taking place. You can select between the two following health insurance plans, which are based on real plans.

PPO Bronze (High deductible, low premium plan)

Monthly premium	$265.83
Annual deductible	$8,000
Coinsurance (for all services)	20%
Out of pocket maximum	$8,550

PPO Gold (Low deductible, high premium plan)

Monthly premium	$480.17
Annual deductible	$0
Coinsurance (for all services)	5%
Out of pocket maximum	$8,150

*Both plans offer a national network, accidental injury, and emergency out of service area coverage

Figure A.21: Demographic Questions

Below, we will ask you a few questions about personal characteristics to help us understand more about insurance choices.

Please rate your current mood along the following characteristics:

0 10 20 30 40 50 60 70 80 90 100

Worried
(0 = not worried, 100 = extremely worried)

Surprised
(0 = not surprised, 100 = extremely surprised)

Alert
(0 = not alert, 100 = extremely alert)

Did you perform any calculations during this survey to help you with choosing health plans?

- Yes
- No

Do you currently have health insurance?

- Yes
- No

The national open enrollment period for selecting a health insurance plan is currently taking place. Are you planning on either signing up for a new plan or switching to a different plan this period?

- Yes
- No

How often do you delay tasks that you have to do?

- Almost always
- Frequently
- Sometimes
- Rarely
- Never

Figure A.22: Demographic Questions

Over the last 2 weeks, how often have you felt down, depressed, or hopeless?

- Not at all
- Several days
- More than half the days
- Nearly every day

On a scale from 0 to 100, how would you rate your physical health in the past year?
(0 = Very Poor Health, 100 = Excellent Health)

0 10 20 30 40 50 60 70 80 90 100

How many relatives, friends, or acquaintances do you know who had a serious health issue or who were hospitalized in the past year?

Please indicate which range your annual income falls into:

What is your highest level of education completed/in the process of obtaining?

Elementary school equivalent	Middle school equivalent	High school degree	Associate's degree	College/bachelor's degree	Master's degree	PhD

What is your gender?

- Male
- Female
- Nonbinary
- Prefer not to answer

What is your racial background?

What is your political affiliation?

What state do you live in?

APPENDIX B

Appendix to Chapter 2

B.1. Sample Construction and Empirical Setup

B.1.1 Sample Selection

The HRS contains distinct cohorts of individuals that entered the longitudinal survey at different times. In particular, I exclude the first three survey waves of the AHEAD cohort (individuals who were born before 1924), because these specific waves for this cohort were out-of-sync with the remainder of the HRS. A small minority of referent individual-wave pairs (0.19%) in the sample (after applying specified sample restrictions) appear with a missing spousal hospital admission variable due to specific reasons: missing due to refusal, a missing web interview, for another "missing" reason, or for an unknown reason. 3.1% of individual-wave pairs have a "spousal hospital admission" in a given survey wave coded as missing due to a spouse not responding to the HRS for that particular survey wave. It is possible that there is a selection issue that could arise from dropping these observations *or* setting these observations equal to zero in their respective wave. I adopt the latter approach of setting the *transformed* indicator variable of whether a spouse was hospitalized equal to zero for these respondents, but given the small proportion of the sample that is affected by refusals or spouses not responding for a particular survey wave, it is unlikely that this decision would affect the main results substantially.

B.1.1.1 "Individual-Oriented" Sample

My sample is "individual-oriented", rather than "spouse-oriented". Particularly, this means that I allow for unmarried individuals to exist in the sample. This implies a few following cases, all of which hold true in my sample. First, an individuals' spouse may pass away or separate from an individual at any particular survey wave. These individuals are permitted to continue on in the sample for successive survey waves, even after their initial spouse has passed away or divorced them. Second, an individual *may* or *may not* re-marry. If an individual re-marries after an initial spouse passes away, and their new spouse is hospitalized, this can potentially be counted as an "initial spousal hospital admission" (as long as their initial spouse was never hospitalized). Exposure to treatment, in other words, is solely defined as the *first time in the observable survey waves that individual i was exposed to a spousal hospital admission event.*

B.1.1.2 "Symmetric Sample"

Furthermore, the sample follows a "symmetric property of relations". Consider individual i, who may be present in one or more survey wave in the HRS. Furthermore, individual i may have a spouse, whom I refer to as i_s. The symmetric property of relations suggests that individual i_s may *also* be present as a distinct individual in the survey, with individual i as their spouse. Following the symmetric property, this implies that both an individual and their spouse are considered to be distinct individuals in the sample construction. I allow for this property to avoid arbitrarily selecting one individual per "household pairing" to focus on. For instance, it is possible that both individual i and individual i_s are hospitalized in the observable survey waves. In these cases, both an individual and their spouse would "experience a spousal hospitalization event". It is important to note under this property that previous sample restrictions still apply: for instance, individual i is excluded from the sample if individual i_s is hospitalized in individual i's first survey wave, and the opposite holds true, as well.

One inherent property of the symmetric sample is that individuals who encounter their own hospitalization event, prior to (or following) their spouse's hospital admission, may themselves be hospitalized. It is possible that an individual whose spouse was admitted to the hospital *after* their own hospital stay might experience heterogeneous treatment effects that differ in comparison to treated individuals who were not initially self-hospitalized. Although I assume relatively homogeneous treatment effects in this paper, I also show that my main findings are robust to potential heterogeneous treatment effects in Appendix B.3.3.

B.1.2 Empirical Details

As I mention in Section 2.3, I cluster standard errors in my analyses at the respondent level instead of the at the household level because the unit of analysis is at the individual level. This decision to cluster at the individual level follows from the symmetric property of the panel dataset that I impose and discuss above in Section B.1.1.2. In other words, because multiple individuals within the same household may be focal individuals, I decide to cluster the standard errors at the individual level. In my analyses, I estimate equation 2.3.1 using weighted least squares using person-level analysis weights from the HRS. The HRS dataset provides person-level weights for the first 13 waves; I reproduce the analogous weight for individuals in Wave 14 using the survey weight from Wave 13.[1] I use person-level weights instead of household analysis weights because the unit of treatment is defined at the individual level. The weights ensure population-level representation by accounting for differential selection probabilities by race/ethnicity and birth cohort and to further correct for differential non-response (Lee et al., 2021).

[1]Note that the 2018 weights were added in the subsequent version of the RAND HRS Longitudinal file.

B.2. Additional Tables, and Figures

Table B.1: Average "Initially Uninsured" Sample Characteristics

Characteristic	Spouse Ever Hospitalized	Spouse Never Hospitalized
Age	66.91	65.53
% White (not Hispanic)	55%	42%
% Above 65	55%	44%
Assets (Checking + Savings)	$18,147	$8,672
Self-hospitalized	27%	30%
Individual Holds Private Plan	26%	16%
Total Annual Private Plan Premium	$452	$217
Individual Covered by Spouse	3.0%	1.0%
Individual Uninsured Second Wave	67%	67%
Number of Unique Respondents	2,579	8,565

Notes: This table presents averages across individual survey wave pairs for the sample of individuals who did not hold a private plan in their first HRS survey wave, allowing for the same individual to be repeated up to 14 waves, using the universe of data between 1992 to 2018. The proportions represent proportions of outcomes across all of the survey waves represented by all individuals. Note that the sample reflects (a) dropping individual-wave pairs with missing weights, (b) dropping individuals whose spouses were hospitalized in the first wave that an individual appeared in the survey, and (c) limiting to individuals who entered the HRS without a private health plan. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.10$. *Data Source:* Health and Retirement Study.

Table B.2: Statistics on Individual-Wave Sample Composition for Main and Heterogeneity Analyses

Characteristic	Initially Uninsured Sample		Full Sample	
	Num. Ever Treated	Num. Never Treated	Num. Ever Treated	Num. Never Treated
Overall sample	20,238	38,447	101,692	118,416
Heterogeneity Analyses				
A. Racial/Ethnic Groups				
Black sample	3,224	11,188	11,621	25,210
White sample	11,658	16,716	78,307	73,444
Hispanic sample	1,133	2,957	2,408	5,412
B. Spousal Insurance Status				
Spouse uninsured sample	2,714	38,447	4,933	118,416
Spouse insured sample	15,762	38,447	91,097	118,416
C. Marriage Length				
Short marriage sample	4,001	38,447	17,497	118,416
Medium marriage sample	6,104	38,447	37,285	118,416
Long marriage sample	6,596	38,447	37,234	118,416
D. Initial Asset Holdings				
Lower Assets	12,040	38,447	35,867	118,416
Higher Assets	6,689	38,447	61,697	118,416

Notes: This table provides sample compositions (individuals are counted multiple times for multiple waves in the survey) for the main regression analysis ("Overall Sample" row) and for different groupings of data for the heterogeneity analyses carried out across different groups. The first two columns present the number of unique individual-wave combinations for "ever-treated" and "never-treated" groups for those who were privately uninsured in their first survey wave. The second two columns present the analogous number of "ever-treated" and "never-treated" individual-wave combinations who were present in the main sample (regardless of initial private health insurance status). I rely on these these samples, respectively, to make different inferences in my analyses. Note that these respective sample numbers take into account dropping individual-wave pairs with missing weights and individuals whose spouses were hospitalized in the first wave that an individual appeared in the survey. I include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. *Data Source:* Health and Retirement Study.

Figure B.1 reports results for individuals who did not have a private plan in (a) their first two and (b) their first three survey waves in the HRS. I interpret these individuals as those with a successively lower propensity to have private insurance. The effect size appears to decrease in magnitude as I limit the sample to only include these individuals. However, there is still a significant increase (albeit slightly delayed) effect for these groups. For instance, there is a 4.6pp increase in private plan take-up one wave post-treatment ($P < 0.05$) and a 4.7pp increase in private plan take-up two waves post-treatment ($P < 0.01$) for individuals who were initially without a private plan in the two and first three survey waves they appeared in the data, respectively. This provides suggestive evidence that observing "salient" health events could result in people who have been initially uninsured for *several* waves to eventually take-up a form of insurance.

Figure B.1: Main: Private Insurance for individuals with increasingly less marginal attachment to being initially privately insured

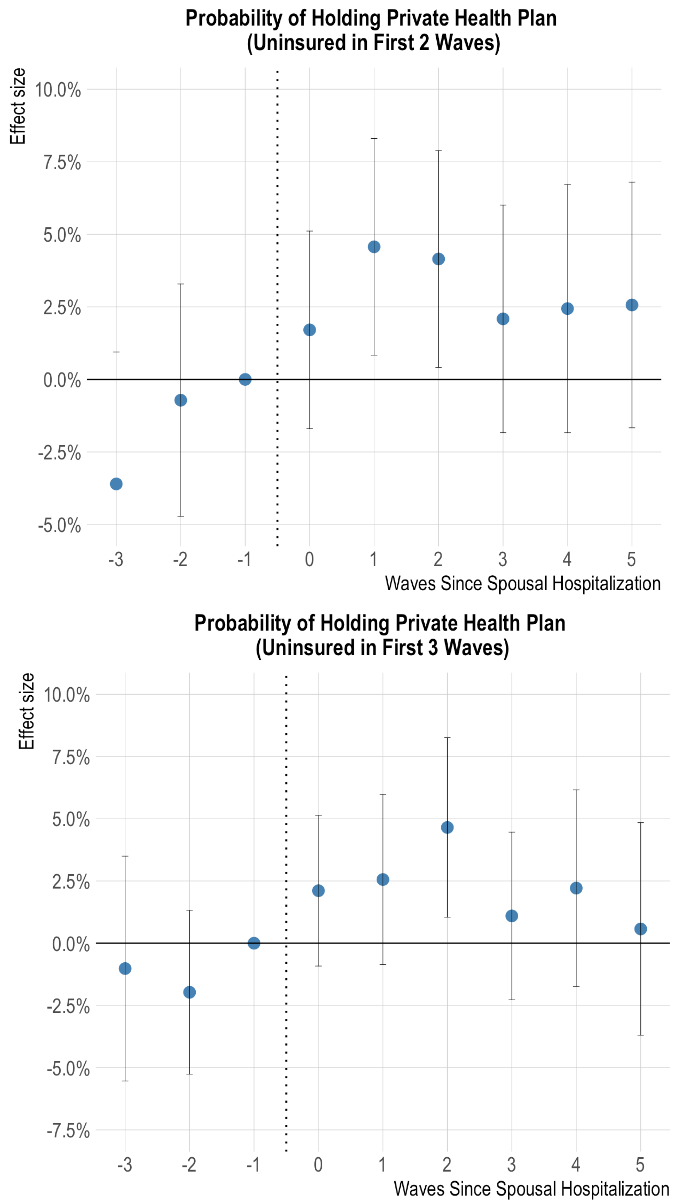

Notes: This figure presents estimates using equation 2.3.1 when I subset the sample to only include individuals who were uninsured in their first *two* survey waves and individuals who were uninsured in their first *three* survey waves. I include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. *Data source:* Health and Retirement Study.

I explore the effect of private insurance take-up after accounting for age, particularly whether an individual is above the age of 65. Figure B.2 presents results from dividing the sample of "initially uninsured" respondents into two distinct sub-samples: one where individuals i are always below the age of 65 ("Less than 65 years" sample, censored once individuals pass the age of 65) and a second where individuals i are always above the age of 65 ("Over 65 years" sample). I find that take-up of private plans increases by 11pp for treated individuals under the age of 65 years in the survey wave corresponding to the spousal hospital admission. This effect persists for two additional survey waves after the spousal event. For individuals above the age of 65 years, there is an increased 5.7pp take-up in private plans in the corresponding year. This suggests that the overall effect is not solely driven by individuals under the age of 65. The effect remains slightly significant for two additional waves. This is particularly noteworthy, considering that individuals above the age of 65 are also eligible for Medicare. Finally, Figure B.3 shows that by five waves following treatment (for the initially uninsured sample), liquid assets have fallen overall ($P < 0.05$).

Figure B.2: Private Insurance Take-up, Censored Above and Below Age 65

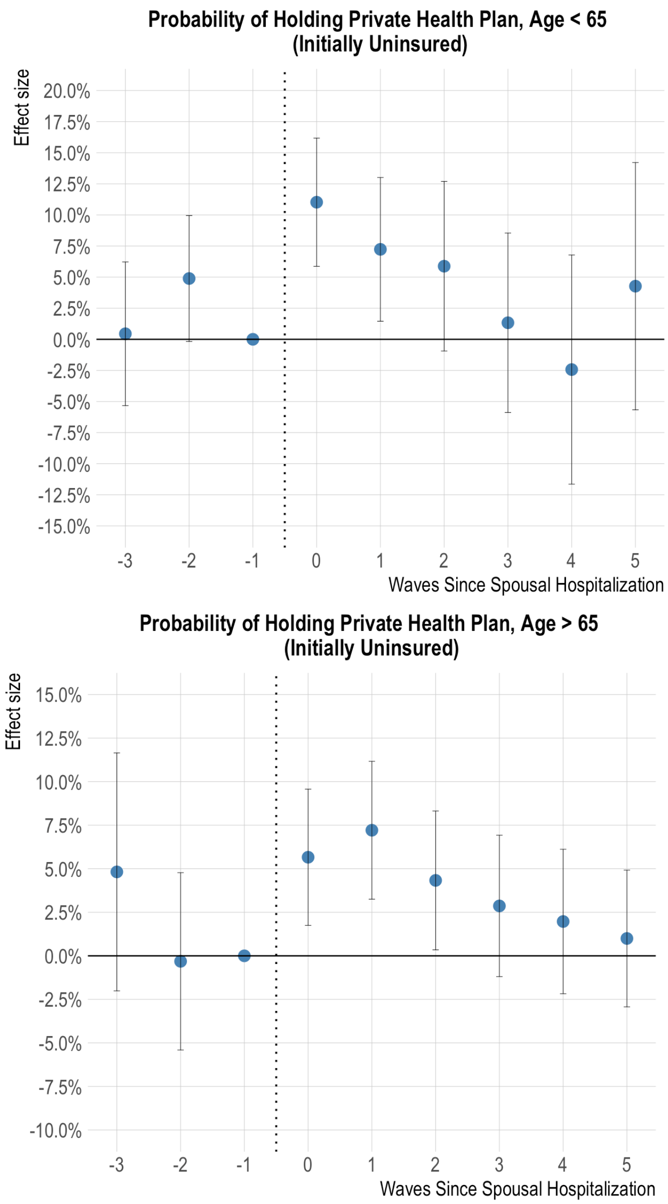

Notes: This figure presents estimates using equation 2.3.1 when I subset the sample to only include individual-wave observations that are (a) below the age of 65 (top panel) and (b) above the age of 65 (bottom panel). I include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. *Data source:* Health and Retirement Study.

Figure B.3: Household Assets After Spousal Hospitalization Event (Individuals Uninsured in their first Survey Wave)

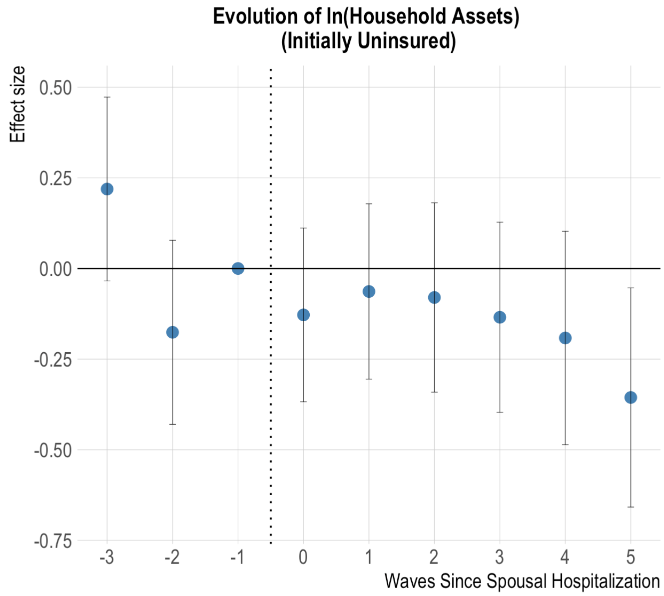

Notes: This figure presents the effect size following a spousal hospitalization event on household assets. This graph reports estimates for individuals who were uninsured in their first survey wave. I include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. *Data source:* Health and Retirement Study.

B.3. Additional Robustness

B.3.1 Balanced Panel

Table B.3: Robustness to Alternative Specifications and Sample Restrictions

	Balanced panel (all lags included)
Wave 3 Pre-treatment Effect	-1.7 (2.7)
Wave 2 Pre-treatment Effect	-0.3 (2.3)
Wave 0 Effect	3.9 (2.4)
Wave 1 Effect	4.0^{*} (2.3)
Wave 2 Effect	3.0^{**} (2.5)
Wave 3 Effect	3.0 (2.5)
Wave 4 Effect	5.0^{*} (3.0)
Wave 5 Effect	5.6^{*} (3.2)
Unique respondents in sample	1,057
Unique observations in sample	10,610

Notes: This table presents point estimates and standard errors in parentheses from running my baseline specification, Equation 2.3.1. This table presents event study results from a balanced sample. To define this sample, all respondents must be present in survey waves -2 and 2. However, I include all of the lags, and there may be attrition for individuals in different waves following wave 2. I omit the wave prior to treatment (μ_{-1}) Furthermore, all individuals are "ever-treated," "initially uninsured" individuals. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.10$. *Data Source:* Health and Retirement Study.

B.3.2 Restricting Sample to "Ever Treated" Individuals

Figure B.4: Main: Private Insurance for only "ever treated" respondents

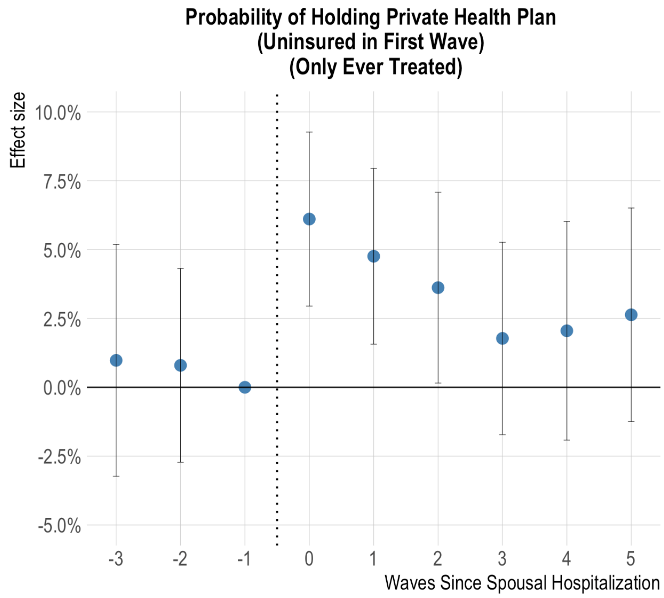

Notes: These show estimates using equation 2.3.1 when I subset the sample to only include individuals who were ever treated. I include 95% confidence intervals corresponding to the treatment effect in each year following a spousal hospitalization. Standard errors are clustered at the individual level. The model fits weighted least squares using person-level analysis weights from the HRS. *Data source:* Health and Retirement Study.

B.3.3 Alternative Event Study Specification Using Interaction Weighted Estimators

I follow Sun and Abraham (2021) and implement their Stata package *eventstudyinteract* to re-run the event study design in a way that is robust to potential heterogeneous treatment effects across cohorts. This approach estimates "interaction-weighted" estimators, formed by estimating separate cohort-specific average treatment effects on the treated (CATT) l periods from initial treatment. The individual $CATT_{e,l}$'s represent the average treatment effect l periods from the initial treatment period for each treatment cohort e. The interaction-weighted estimators method then averages estimates of $CATT_{e,l}$ across cohorts e at every given period l relative to the spousal hospitalization. The specification I use to estimate individual $CATT_{e,l}$'s follows from Sun and Abraham (2021):

$$CATT_{e,l} = Y_{it} = \alpha_i + \lambda_t + \sum_{e \notin C} \sum_{l \notin -1} \delta_{e,l}(\mathbb{1}\{E_i = e\} \cdot D_{it}^{l}) + \epsilon_{it} \tag{B.1}$$

In the equation above, E_i is the time for unit i to initially receive a binary absorbing treatment, α_i and λ_t are unit and time fixed effects, and D_{it}^{l} is an indicator for unit i being l periods away from initial treatment at calendar time t. The equation interacts relative period indicators with cohort indicators and excludes indicators for cohorts from some set C. As defined in Sun and Abraham (2021), $E_i = \infty$ for never-treated units. In this setting, because there is a never-treated cohort (i.e. individuals whose spouses are never hospitalized in the observable survey window), $\infty \in supp(E_i)$. This allows me to set C = $\{\infty\}$ in my specification.

The presence of a never-treated group allows me to estimate equation B.1 on all observations. Weights are estimated by sample shares of each cohort in each relevant period and the interaction-weighted estimator is formed through a weighted average of estimates for $CATT_{e,l}$ from Equation B.1.[2] I present the results from the alternative estimation method proposed by Sun and Abraham (2021), which is robust to treatment effects heterogeneity, in Figure B.5.

[2]Sun and Abraham (2021) discuss the specific weighting scheme that the *eventstudyinteract* package employs. Particularly, the interaction-weighted estimator associated with relative period l is "guaranteed to estimate a convex average of $CATT_{e,l}$ using weights that are sample share of each cohort e." Furthermore, the weights are guaranteed to be convex, sum to one, and be non-negative (unlike the standard event study approach).

Figure B.5: Interaction-Weighted Estimators

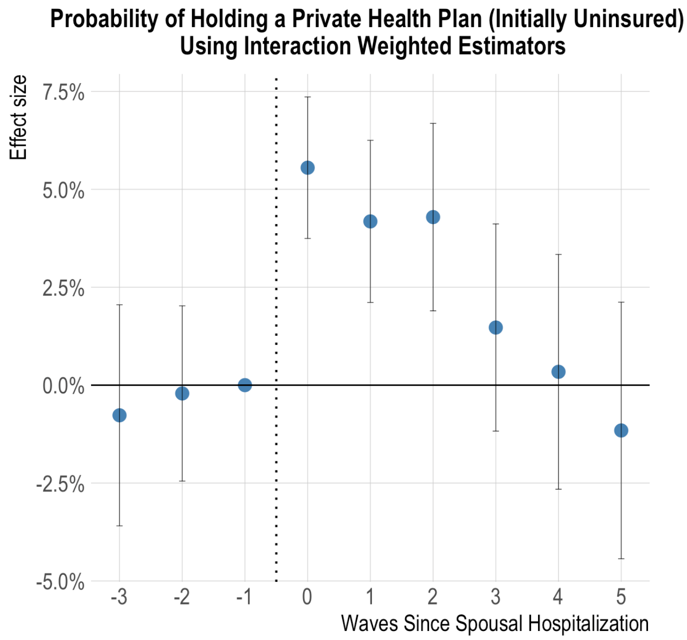

Notes: This figure presents results using the *eventstudyinteract* package prepared by Sun and Abraham (2021). I bin the relative times before lead period 4 and lag period 5 and assume constant treatment effects within those two respective bins. In the *eventstudyinteract* specification, I use individuals whose spouses were never hospitalized in the observable sample to be set as the control cohort. I plot 95% confidence intervals corresponding to each point estimate and standard errors are clustered at the individual level. *Data source:* Health and Retirement Study.

APPENDIX C

Appendix to Chapter 3

C.1. Further Details: Example Letter

Figure C.1: Letter Informing Parents of Open Enrollment Status

5/24/2021 GAMUT Online : Martinez USD : Open Enrollment Act Transfers E 5118

Martinez USD | E 5118 Students

Open Enrollment Act Transfers

PARENTAL NOTIFICATION:

OPTION TO TRANSFER

[Date]

To the parents/guardians of students at __________________ School:

On January 7, 2010, the California Legislature signed into law the Open Enrollment Act, which is known as the Romero Bill and referenced in Education Code Sections 48350-48361. This letter is to notify you that for the [School Year] [School Name] will be listed on the "Open Enrollment" list from the state.

The Open Enrollment Act requires the California Department of Education (CDE) to annually create a list of 1,000 schools ranked by their Academic Performance Index (API). A school's API is a number ranging from 200 to 1000 and is calculated using the results for each school's students on statewide tests. 800 is the API target set by the state of California. For additional information about how the Open Enrollment List is established, please visit the CDE web site: http://www.cde.ca.gov/sp/eo/op.

Having [School Name] on the Open Enrollment list means all students eligible for enrollment at [School Name] during the 2011-2012 school year have the option to apply for transfer to a school with a higher Academic Performance Index (API). [School Name] current API is [Current Year Growth API]. The school to which you request a transfer must have a higher API than [Current Year Growth API].

Students may request transfers to schools within the Martinez Unified School District or outside the district. Schools within the district eligible for incoming transfer requests include: [Eligible Schools]. Parents are responsible for all transportation to and from any new school resulting from Open Enrollment transfers.

Students applying for transfer within Martinez Unified School District will be placed based on space availability. If more applicants are received than available spaces exist, a lottery will be held to determine which student's requests for transfer will be accepted.

The application window for requesting a transfer to another school within Martinez Unified School District will be from March 10, [Current School Year] to April 10, [Current School Year]. All applications will be reviewed after the close of the application window. If you wish to request a transfer to another school in the state of California under the Open Enrollment Act you will need to apply for a release from Martinez Unified School District. Applications for release from Martinez Unified School District are available at any time. These applications are available through the Director of Student Services Office at the district office.

There is no action needed if you would like to remain enrolled in [Current School].

If you are interested in applying for transfer to schools outside Martinez Unified School District with a higher Academic Performance Index (API), please contact those school districts directly for information about their application requirements and timelines. Additional information regarding schools eligible for incoming transfer can be found through the CDE web site: http://api.cde.ca.gov/reports.

Sincerely,

[Director's Name]

Director, Student Services

Exhibit MARTINEZ UNIFIED SCHOOL DISTRICT

version: December 13, 2010 Martinez, California

revised: June 23, 2014

Notes: The above is an example of the type of letter that was required by the law to be sent to parents of students who were attending low-performing schools that were included on the list of 1,000 open enrollment schools in a given school year. The legislation required schools to reach out to parents in their own native language if 15% or more of that schools' population spoke a language other than English.

C.2. Adjustments for Overcounting Bias

A valid concern in interpreting the point estimates throughout this paper is that students who leave treated schools presumably switch to other, higher performing public schools within their district or in another district, providing these receiver schools with a positive bump in enrollment. Any estimation approach measuring changes in enrollment in treated schools *relative to* comparison schools would thus capture not just the true average treatment effect on the treated (ATT) associated with students leaving treated schools, but also an indirect, secondary effect of students joining comparison schools. The combination of these two effects could lead us to overestimate the true effect of the policy.

For example, consider a simplified setting where there is exactly one school on the open enrollment list and exactly five schools not on the open enrollment list ("comparison" schools). Suppose the true effect of the policy is that the school on the open enrollment list loses ten students. Regardless of which of the remaining five school(s) these ten students move to, assuming all ten students remain in the public school system, the average comparison school gains two students. Since the treated school lost ten students and each comparison school gained two students (on average), when we estimate the effect of the policy by attempting to measure how much enrollment has declined at treated schools relative to comparison schools, we might measure this effect as a decline in enrollment of -12 students instead of -10 students. Notice that the size of the bias (here, a 20% bias away from zero) is directly proportional to the ratio of treated schools to comparison schools. We formally derive this source of bias in a simple 2x2 difference-in-difference setting below, but the same logic applies to more generalized models. In reality, because the number of comparison schools is substantially larger than the number of treated schools under the California Open Enrollment Act, we determine that this source of bias is small. Our approximations are given in section C.2.2 below, after deriving the source of bias.

C.2.1 Derivation of Bias Adjustment Calculation

For simplicity, our derivation of point-estimate adjustments associated with correcting for "double counting" are presented in a 2-period DiD framework below. For a discussion the separate source of bias that can arise from multi-period DiD TWFE estimation (i.e. Goodman-Bacon (2021); Callaway and Sant'Anna (2021)), please see Appendix C.5.

Let $Y_{g,t}$ be an observed outcome variable of interest for group $g \in \{T, C\}$ (Treatment or Control) at time $t \in \{1, 2\}$. Group T switches from not treated to treated as we move from time $t = 1$ to $t = 2$, while Group C is not treated. Let $Y_{g,t}(0)$ and $Y_{g,t}(1)$ denote *potential* (or *counterfactual*) outcomes of group g in time t without and with the treatment, respectively.

To measure the causal effect of treatment, we are interested in identifying the average treatment effect for group T (ATT):

$$ATT = \mathbb{E}\left[Y_{T,2}(1) - Y_{T,2}(0)\right].$$

The standard (2-period) difference-in-difference (DiD) estimator is:

$$\widehat{DiD} = \left(Y_{T,2} - Y_{T,1}\right) - \left(Y_{C,2} - Y_{C,1}\right).$$

In order for the DiD estimator to be an unbiased estimate of the ATT, the parallel trends assumption must hold. That is, we must have:

$$\mathbb{E}\left[Y_{T,2}(0) - Y_{T,1}(0)\right] = \mathbb{E}\left[Y_{C,2}(0) - Y_{C,1}(0)\right].$$

Under this assumption, we can then write:

$$\begin{aligned}
\mathbb{E}\left[\widehat{DiD}\right] &= \mathbb{E}\left[\left(Y_{T,2} - Y_{T,1}\right) - \left(Y_{C,2} - Y_{C,1}\right)\right] \\
&= \mathbb{E}\left[\left(Y_{T,2}(1) - Y_{T,1}(0)\right) - \left(Y_{C,2}(0) - Y_{C,1}(0)\right)\right] \\
&= \mathbb{E}\left[Y_{T,2}(1) - Y_{T,2}(0)\right] + \mathbb{E}\left[Y_{T,2}(0) - Y_{T,1}(0)\right] - \mathbb{E}\left[Y_{C,2}(0) - Y_{C,1}(0)\right] \\
&= \mathbb{E}\left[Y_{T,2}(1) - Y_{T,2}(0)\right]
\end{aligned}$$

where the final step follows from the parallel trends assumption stated above. This shows that, under the parallel trends assumption, $\widehat{DiD}$ is an unbiased estimator of the ATT.

However, in our context and as highlighted in the example above, our observed value $Y_{C,2}$ differs from the potential/counterfactual outcome $Y_{C,2}(0)$ associated with non-treatment in $t = 2$ for Group C, because Group C *is* indirectly impacted by treatment; when Group T schools lose one student on average, any given Group C school should expect to gain $\frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}$ students, on average. Therefore, if treatment impacts the enrollment outcome at Group T schools by $Y_{T,2}(1) - Y_{T,2}(0)$, it should impact enrollment at Group C schools by $\left(\frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)\left(Y_{T,2}(1) - Y_{T,2}(0)\right)$ in the *opposite* direction. So, our observed value $Y_{C,2}$ includes both $Y_{C,2}(0)$, *and* a bias term that is directly (negatively) proportional to the ATT and which captures the average *gain* in students caused by the movement of students from Group T to Group C schools:

$$Y_{C,2} = Y_{C,2}(0) + (-1)\left(\frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)\left[Y_{T,2}(1) - Y_{T,2}(0)\right]$$

Plugging this in for $Y_{C,2}$ in our expression for the expected value of our DiD estimator we have:

$$\begin{aligned}
\mathbb{E}\left[\widehat{DiD}\right] &= \mathbb{E}\left[\left(Y_{T,2} - Y_{T,1}\right) - \left(Y_{C,2} - Y_{C,1}\right)\right] \\
&= \mathbb{E}\left[\left(Y_{T,2}(1) - Y_{T,1}(0)\right) - \left(Y_{C,2}(0) + (-1)\left(\frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)\left[Y_{T,2}(1) - Y_{T,2}(0)\right] - Y_{C,1}(0)\right)\right] \\
&= \mathbb{E}\left[Y_{T,2}(1) - Y_{T,2}(0)\right] + \mathbb{E}\left[Y_{T,2}(0) - Y_{T,1}(0)\right] - \mathbb{E}\left[Y_{C,2}(0) - Y_{C,1}(0)\right] \\
&+ \left(\frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)\mathbb{E}\left[Y_{T,2}(1) - Y_{T,2}(0)\right] \\
&= \mathbb{E}\left[Y_{T,2}(1) - Y_{T,2}(0)\right] + \left(\frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)\mathbb{E}\left[Y_{T,2}(1) - Y_{T,2}(0)\right] \\
&= \left(1 + \frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)\mathbb{E}\left[Y_{T,2}(1) - Y_{T,2}(0)\right]
\end{aligned}$$

From this, we can see that to obtain an unbiased estimator of $ATT = \mathbb{E}\left[Y_{T,2}(1) - Y_{T,2}(0)\right]$, we must divide our DiD estimator by $\left(1 + \frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)$. Then we have:

$$\mathbb{E}\left[\frac{\widehat{DiD}}{\left(1 + \frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)}\right] = \mathbb{E}\left[Y_{T,2}(1) - Y_{T,2}(0)\right]$$

which is the effect we are attempting to isolate (ATT).

Since $ATT = \mathbb{E}\left[Y_{T,2}(1) - Y_{T,2}(0)\right]$, we can then write:

$$ATT = \mathbb{E}\left[\frac{\widehat{DiD}}{\left(1+\frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)}\right] \implies$$

$$\boxed{\mathbb{E}\left[\widehat{DiD}\right] = ATT * \left(1+\frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)} \tag{C.1}$$

Therefore, the point estimate above contains a bias away from zero of the size $\left(\frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right) * ATT$. To correct for this source of bias, we can scale our point estimates down by $\left(\frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)$, which is a decimal. In section C.2.2 below, we compute the size of this source of bias for each year in our treatment period using counts of open enrollment schools and comparison schools during policy years.

We also recognize that, while most students who left treated public schools in response to the policy likely joined comparison public schools that were not on the OE list, some families may have responded to the policy by enrolling their child(ren) in a charter school, a private school, or homeschooling them. Because this expands the set of potential receiver schools, this channel reduces the size of the bias, suggesting that estimates of the bias containing only public schools are likely an overestimate of the true size of the bias. We therefore also compute the ratio of treated to comparison (receiver) schools using this broader definition of receiver schools (including private and charter), for comparison.

C.2.2 Calculation of Error Bounds

Table C.1 presents underlying counts of schools we use to construct the bounds on this source of bias. To find the size of the potential bias described above under the definition of comparison schools as only "other public schools", we divide the average number of treated schools across policy years from Column (2) by the average number of public comparison schools in Column (3). To find the error bound using a less stringent definition of comparison schools, we calculate an estimate for the number of private schools in Column (4). The CDE only provides us with total private school student enrollment, and so we divide total enrollment each year by the average private school enrollment in California in 2022, 188 students. Because we only have the number of charter schools between years 2010-2013, we estimate the number of charter schools in 2014 and 2015 by assuming a linear growth rate in charter schools between the years 2010-2013 and applying this growth rate to estimate the number of charter schools in 2014 and 2015 (Column 5). Column 6 sums Column (3) + Column (4) + Column (5). We divide the average number of treated schools (Column 2) over the average number of comparison schools (Column 6) to come up with an alternative measure of bias, presented in Table C.2.

Table C.1: Calculation Underlying Error Bounds in Table C.2

(1) Year	(2) # Treated Schools	(3) Comparison Schools	(4) Estimate for # Private Schools	(5) Estimate for # Charter Schools	(6) Estimate for *All* Comparison Schools
2010	966	8061	2740	738	11539
2011	954	8098	2644	819	11561
2012	951	8080	2746	900	11726
2013	955	8037	2720	962	11719
2014	962	9041	2679	1037	12756
2015	958	9077	2663	1111	12852
Average (Rounded)	958	8399	2699	928	12025

Notes: This table provides the counts used for underlying calculations of potential error bounds presented in Table C.2. Column 1 lists all the years the OE policy was in place. Column 2 lists the number of treated schools in each year (the number of schools that appeared on a given years' open enrollment list. Column 3 lists the number of public schools *not* on the OE list in a given year. Columns 4 and 5 list an estimate of the number of private and charter schools in the state for each policy year. Column 6 sums the number of public non-OE schools with the estimate of charter and private schools. The final row takes the average for each column.

Our estimates of the size of bias are likely overestimates, because we only include only "non-OE" public schools in the denominator, when in fact the law only stipulates that students leaving a treated school must attend a higher-ranked school (allowing them to technically switch to other OE schools in a given year, although movement through this channel seems unlikely). We also cannot account for families who home-school their children in response to the policy. This is why we consider our error estimates in Table C.2 to be upper bounds on the size of this source of bias. Recall that to correct for the bias, we need to scale our point estimates down by $\left(\frac{\text{Number of Treated Schools}}{\text{Number of Comparison Schools}}\right)$. On average across years, there are 958 treated schools and 8399 comparison schools. This leads to an average bias size of 958/8399 = .114 or 11.4%. When using the broader definition of comparison/receivers schools, this denominator grows to 12025, suggesting a bias size of .079, or 7.9%. These computations are reflected in Table C.2 below.

Table C.2: Calculated Error Bounds: How much Bigger is $\widehat{DiD}$ than ATT?

Definition of Comparison school	Size of Bias (%)
Only Public Non-OE Schools	11.4%
Comparison, Private, Charter	7.9%

Notes: The table above presents potential upper bounds on how biased our point estimates $\widehat{DiD}$ may be, relative to the true ATT (abstracting from other sources of bias). To compute the percentages in the table, we divide the number of schools on the open enrollment list in each year the policy was in place by the relevant total number of comparison schools (defined in the two ways described above and counted in columns (3) and (6) of the table above) for each year. The bias is different each year as the number of schools in each group changes from year to year. We report the *average* of this ratio in the table. Our preferred interpretation of number of non-treatment schools is row 2, which includes non-treated public schools as well as charter schools and private schools (which by definition are also non-treated). These together represent a set of potential recipient institutions for students leaving treated schools.

Based on these back-of-the-envelope calculations, we find that our point estimates might be up to 11.4% (or 7.9%) larger than the true ATT (abstracting from other types of potential bias) because of this "double-counting" issue. Therefore, when interpreting point estimates, consider that the true effect of the policy might be the stated point estimate scaled down by 11.4% (or 7.9%).

While the examples provided above are described in levels, our outcome variables are measured in logs. Because percent change is a linear approximation of log differences, we interpret our point estimates as percentage declines in enrollment in treated schools relative to comparison schools. To provide an analogous example to the one above, in a world where there is one treatment school and five comparison schools, if the true treatment effect on the treated is that enrollment falls at treated schools by 10%, then each comparison school can expect an average 10/5= 2% enrollment bump *only if* comparison schools are the same size (in terms of the number of enrolled students) as treatment schools, on average. We test this using paired t-tests for each year that the California Open Enrollment policy is in effect; with the exception of the 2010 open enrollment year (when the policy was not yet in full effect), the difference in average school size between open enrollment and comparison schools is not significant at conventional levels. We therefore assume that the bias size adjustment is also appropriate for our logged outcome variables.

C.3. Discussion on Interpreting Magnitudes

Our outcomes are log-transformed, which implies that OLS results using our empirical specifications will tell us the policy effect on the percentage change in enrollment variables (our outcomes). By construction, the FRPM and non-FRPM enrollment variables sum to the total enrollment at each school in each year. However, we should not expect that the percentage decline in total enrollment will equal the average of the percentage decline of its subgroups, because subgroup attrition could vary at the school level, such that average total enrollment effects may be weaker or canceled out altogether. For instance, a likely scenario is that a particular school that loses FRPM students as a result of the policy may simultaneously experience gains in non-FRPM students, or may lose non-FRPM students at a lower rate than it loses FRPM students. Similarly, the same could be true for schools that lose non-FRPM students: these same schools could experience smaller proportional losses in FRPM enrollment or even gains in FRPM student enrollment as a result of the policy. The result of this would be that total enrollment at treated schools would still decline in the same "level" amount, but in percentage terms, it would appear to be a smaller change, because at a school level, the non-uniform decreases and increases in FRPM and non-FRPM enrollment would cancel themselves out, to a large extent. Furthermore, variation in the overall school enrollment plays an important role in determining average percentage declines.[1] This is consistent with what we observe in our empirical results: we observe a 1.4% decline in total enrollment as a result of the OE policy, but the percentage declines in FRPM and non-FRPM enrollment are both larger (6.7% and 6.6% declines). The event study similarly presents larger estimates of FRPM and non-FRPM enrollment, relative to total enrollment.

[1]Note that no matter how the entry and exit of FRPM and non-FRPM students due to the policy distributes *across* schools, the average *level* effect of enrollment will always equal the sum of the average level effect of the policy on FRPM and non-FRPM students. This does not necessarily hold for percentage changes, however.

C.4. Additional Tables

Table C.3: Accuracy Scores for our Open Enrollment Algorithm

Open Enrollment Year	Num. Schools on OE List	Accuracy
2010-11 List*	1,000	91.5%
2011-12 List	NA	NA
2012-13 List	978	93.6%
2013-14 List	979	93.2%
2014-15 List	984	94.9%
2015-16 List	974	95.2%

Notes: This table presents algorithm accuracy scores, defined here as the number of open enrollment schools determined by the algorithm that also fall on the actual corresponding (published) open enrollment list in year X, divided by the number of schools on the actual published open enrollment list in year X. The accuracy score for the 2011-12 List is N/A because we did not receive the actual published list from the California Department of Education (CDE). Additionally, the 2010-11 list is based on a draft list provided by the CDE, and so the number of schools on the open enrollment list is slightly higher than other years because we do not have information on schools that applied for waivers in this year. As such, the preferred rows for interpreting accuracy are the rows corresponding to the 2012-13 through 2015-16 Lists.

Table C.4: District Population Characteristics by School Treatment Status

	All schools	Never on OE list	On OE list at least once
	(1)	(2)	(3)
A. Income and wealth variables			
Number of observations	105,057	88,908	16,149
Median House Value	\$428,497.70	\$440,357.70	\$363,219.10
	(233511.8)	(240336.4)	(178031.5)
GDP per capita	\$29,825.55	\$30,532.54	\$25,933.26
	(12557.84)	(12966.26)	(9078.011)
Median income	\$36190.38	\$36886.93	\$32355.62
	(12609.44)	(13101.58)	(8483.218)
B. Racial shares			
Share whites	.427	.4337	.3926
	(.232)	(.2312)	(.2351)
Share Hispanic	.369	.3595	.4225
	(.2153)	(.2119)	(.2261)
Share blacks	.0525	.0525	.0528
	(.0567)	(.0565)	(.0577)
C. Occupations			
Share Agriculture	4.026	3.68	5.92
	(7.93)	(7.31)	(10.52)
Share Construction	6.62	6.59	6.81
	(2.54)	(2.55)	(2.49)
Share Manufacturing	9.17	9.22	8.91
	(4.27)	(4.27)	(4.27)
Share Professionals	11.63	11.84	10.48
	(4.47)	(4.51)	(4.02)
Share Educational Services	20.29	20.38	19.78
	(4.21)	(4.21)	(4.20)

Notes: This table presents descriptive statistics for the overall population in school districts (matched to treated and comparison California schools in our open enrollment data). Column (1) shows the means and standard deviations computed over all year-school observations in our panel dataset. (2) and (3) divide the sample into California schools that were on the open enrollment list at least once between 2010 to 2015 and those that never appeared on the open enrollment list.

C.5. Recent Developments in Difference-in-Difference and Event Study Models using TWFE Estimation: Robustness Checks

C.5.1 Estimating Separate ATT(g,t)'s by Treatment-Cohort Group and Year

To address the issues posed by potential treatment effect heterogeneity, we follow Callaway and Sant'Anna (2021) by producing disaggregated estimates of the average treatment effect on the treated for each cohort g (defined as the set of schools first appearing on the open enrollment list at time period g), in year t, denoted $ATT(g,t)$, and referred to as a "group-time average treatment effect." This parameter is defined as:

$$ATT(g,t) = \mathbf{E}[Y_t(g) - Y_t(0)|G_g = 1], \text{for } t \geq g$$

where $Y_t(g)$ denotes the potential outcome that a school would experience at time t if they were first treated in year g (i.e. if they were a member of cohort g), $Y_t(0)$ is the unit's potential outcome at time t if they remain untreated throughout the period of study, G_g is a dummy variable equal to one if a unit is first treated in period g (belongs to cohort g) and $t \geq g$ clarifies that we can only compute group-time average treatment effects on the treated for cohort g in the years *after* cohort g is treated.

The purpose of estimating disaggregated ATT(g,t) parameters, as Callaway and Sant'Anna (2021) articulate, is to avoid issues arising from treatment effect heterogeneity due to the aggregated TWFE difference-in-difference estimators comparing "later treated units with earlier treated units". We estimate each ATT(g,t) by sub-setting the data to only contain observations in periods t and $g-1$, and to contain only units with $G_g = 1$ (first treated in g) and never treated units. We then estimate the following:

$$Y^{g,t} = \alpha_1^{g,t} + \alpha_2^{g,t} G_g + \alpha_3^{g,t} \mathbb{1}(T = t) + \beta^{g,t} G_g * \mathbb{1}(T = t) + \varepsilon^{g,t} \tag{C.1}$$

where $Y^{g,t}$ is an outcome of interest (enrollment, enrollment by race, enrollment by FRPM status), $\varepsilon^{g,t}$ is an error term, all other variables are defined as above and where the coefficient of interest is each $\beta^{g,t}$ which, as Callaway and Sant'Anna (2021) show, estimates $ATT(g,t)$. We estimate these alternative treatment effects for all of our main total enrollment, racial, and socioeconomic enrollment outcome variables.

C.5.1.1 Empirical Results from Aggregating ATT(g,t)'s

After computing $ATT(g,t)$ for each g and t, we can aggregate these estiates in various ways. We present the results from two different aggregation approaches below. The first approach aggregates ATT(g,t) estimates into a *single* simple weighted-average of all ATT(g,t)'s. We present the results from this approach in Table C.5. Following a schools' appearance on the OE list, total enrollment falls by 3pp, FRPM enrollment falls by 9pp, and Hispanic enrollment falls by 7pp. The second approach is to report the average effect of participating in treatment for the group of schools that have been on the open enrollment list for exactly e time periods (years). We present the results from this approach in Figure C.2. For total enrollment as an outcome, the average ATT in the initial treatment year is a -1.6% decline in total enrollment, relative to enrollment at comparison schools (P < .01), and the ATT grows in magnitude over successive years. The average ATT in the initial treatment year is -2.9% for Hispanic enrollment (P < .01) and -2.3% for FRPM student enrollment (P < .01), with effect sizes growing in magnitude for both groups over successive years.

Table C.5: ATT Using Alternative TWFE Estimation Procedure

Outcome	Weighted-Average of all ATT(t,g)
Total Enrollment	$-.03^{***}$
	(.01)
FRPM Enrollment	$-.09^{***}$
	(.01)
Non-FRPM Enrollment	$-.06^{***}$
	(.02)
Hispanic Enrollment	$-.07^{***}$
	(.01)
Black Enrollment	$-.04^{***}$
	(.01)
White Enrollment	$-.03^{*}$
	(.01)

Notes: This table presents a simple weighted-average of all ATT(g,t) pairs using the following weighting scheme used by Callaway and Sant'Anna (2021): $\theta_W^{simple} = \frac{1}{\kappa}\sum_{g\in G}\sum_{t=2}^{T}\mathbb{1}\{g \leq t\}ATT(g,t)P(G=g|G\leq T)$ where $\kappa = \sum_{g\in G}\sum_{t=2}^{T}\mathbb{1}\{t \geq g\}P(G=g|G\leq T)$, which ensures that the weights on ATT(g,t) in the second term sum up to one. We produce these results using the *csdid* command in Stata. We allow for "not-yet-treated units" to be included in the comparison group. The method we employ is a doubly robust DiD based on IPW and OLS. $^{***}p < 0.01,^{**}p < 0.05,^{*}p < 0.10$. This figure reports unadjusted point estimates; see Appendix C.2.2 for an approximation of the potential upper bound of bias on the estimates. Source: California Department of Education; CALPADS; IPUMS.

Figure C.2: Event Study Graphs Using Aggregation of Individual ATT(g,t) pairs

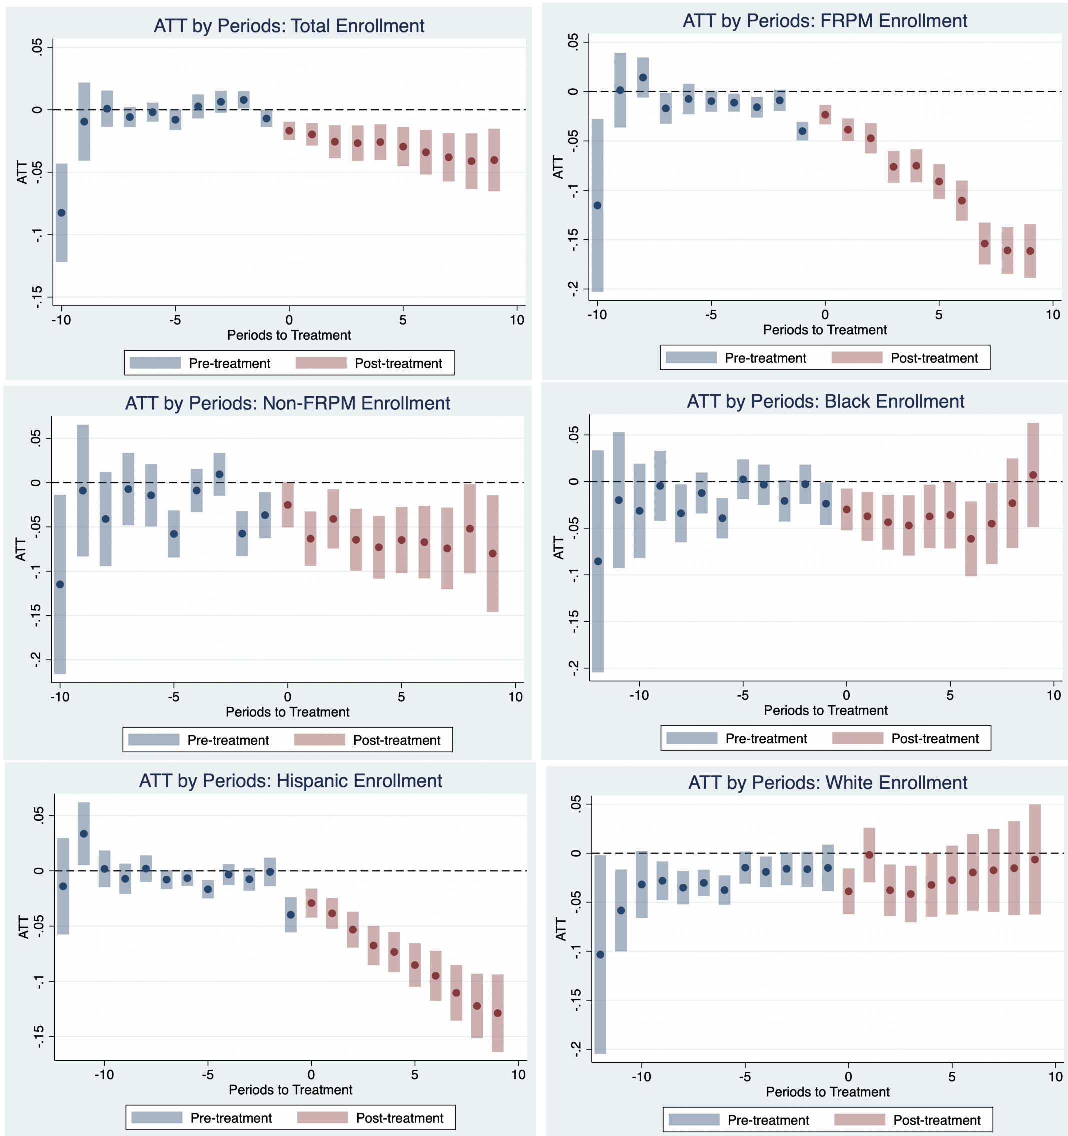

Notes: The graph displays the average treatment effect of participating in treatment for the group of schools that had been exposed to treatment for exactly e periods (in this case, measured in years). We produce these results using the *csdid* command in Stata. We produce these results using the *csdid* command in Stata. We allow for "not-yet-treated units" to be included in the comparison group. The method we employ is a doubly robust DiD based on IPW and OLS. Shaded bands around point estimates represent 95% confidence intervals. Source: California Department of Education; CALPADS; IPUMS.

C.5.2 Weights Underlying Event Study Pre-Trends

As suggested by Sun and Abraham (2021), event study coefficients of interest (in our context, the μ_Y's in Equation 3.5.1 above) can be decomposed into a linear combination of effects from both period Y and other periods. This means that the μ_Y coefficients on indicator variables for *leads* of treatment (i.e. those where $Y < 0$) are actually affected by both pre-trends *and* treatment effects from later periods. Sun and Abraham (2021) note that when treatment effects are homogeneous (and other assumptions hold), these effects from later periods cancel out, but when treatment effects are heterogeneous, the effects do not cancel and this contaminates the estimates. This result calls into question the validity of the common practice of using insignificant pre-trends as evidence in support of the parallel trends assumption.

Sun and Abraham (2021) show that one can calculate the weights underlying the linear combination of effects for each μ_Y. Their Stata package *eventstudyweights* automates this, and we below provide a graphical example of these underlying weights for our coefficient of interest μ_{-2}, arguably the most important coefficient we might use to test for pre-trends (recall μ_{-1} is omitted such that all other estimates are relative to this period).

To clarify terms, Sun and Abraham (2021) first define the "cohort-specific average treatment effect on the treated (CATT) l periods from initial treatment" as:

$$CATT_{e,l} = \mathbb{E}\left[Y_{i,e+l} - Y^{\infty}_{i,e+l}\right]$$

where e is the treatment cohort (first time of treatment), l is the number of periods from initial treatment, $Y_{i,t}$ is the actual outcome of interest for unit i at time t and $Y^{\infty}_{i,t}$ is the potential outcome of unit i a at time t if i were to never receive treatment.

In Sun and Abraham (2021) Propositions 1 through 3, the authors show how μ_Y can be broken down into a linear combination of the cohort-specific average treatment effects ($CATT_{e,l}$'s) with corresponding weights $\omega^{Y}_{e,l}$. As general properties (see Proposition 1), the weights should add to 1 in period Y, add to -1 across omitted periods, and add to zero in all other periods. Figure C.3, top panel, plots these weights for our estimated μ_{-2} coefficients, and compares this to an example plot (bottom panel) that Sun and Abraham (2021) replicate this for the coefficient μ_{-2} in the event study performed by Dobkin et al. (2018). As Sun and Abraham (2021) describe, if the weights on lags of treatment for μ_{-2} are nonzero, then μ_{-2} is particularly sensitive to heterogeneous treatment effects and does not accurately capture pre-trends (because it is being influenced by lags of treatment), whereas if these weights are close to zero, μ_{-2} is not particularly sensitive to potentially heterogeneous treatment effects.

Since the calculated weights underlying μ_{-2} for lags of treatment are close to zero in our setting, we take this as suggestive evidence that our estimated coefficient μ_{-2} is not being significantly contaminated by this form of bias.

Figure C.3: Graphical Illustration of Weights Underlying μ_{-2} Coefficient

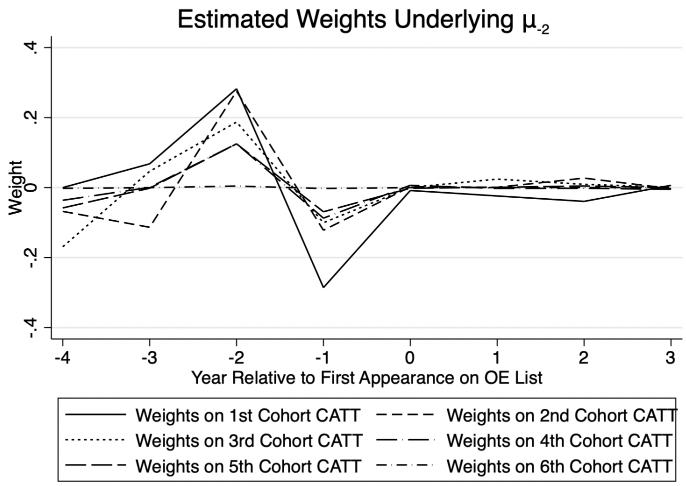

Figure 3: Estimated Weights $\widehat{\omega}_{e,\ell}^{-2}$ Underlying μ_{-2}

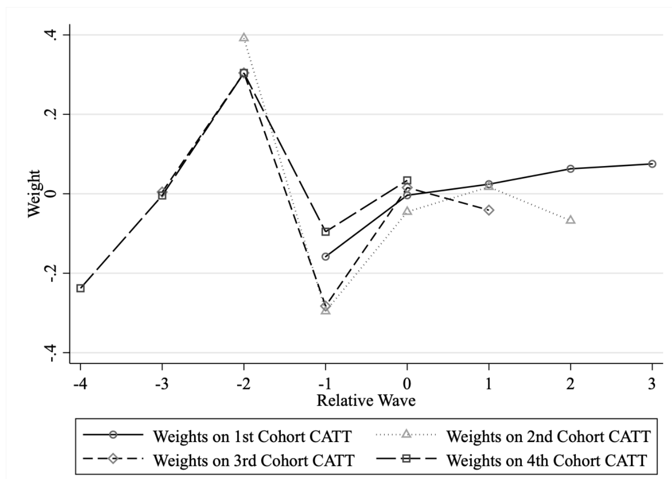

Notes: The top figure shows the weights underlying our event study coefficient, μ_{-2}, which attempts to measure pre-trends in our data (after we exclude charter schools from the sample). The bottom figure shows the weights underlying the event study coefficient μ_{-2} in Figure 3 in Sun and Abraham (2021), using data and setup from Dobkin et al. (2018) as an empirical illustration. Sun and Abraham (2021) demonstrate that μ_{-2} is a linear combination of the cohort-specific average treatment effect on the treated (CATT) across l periods from the initial treatment year and across each cohort, e. The weights presented follow properties discussed by Sun and Abraham (2021): the 6 relative weights from relative year $l = -2$ sum to one, the weights for the other relative years $l \in$ {-3, 0, 1, 2, 3} that we include in the graph sum to zero, and the weights from excluded relative years ($l \in$ {-4, -1} in the figure by Sun and Abraham (2021), only $l = -1$ in our setup) sum to negative one across the excluded waves. We produce our figure using the Stata package *eventstudyweights* by Sun and Abraham (2021).

BIBLIOGRAPHY

BIBLIOGRAPHY

Abaluck, J., M. Bravo, P. Hull, and A. Starc (2020): "Mortality Effects and Choice Across Private Health Insurance Plan," NBER Working Paper.

Abaluck, J. and J. Gruber (2011): "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," American Economic Review, 101.

——— (2016): "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance," American Economic Review, 106.

Ahn, W. and K. Mermin-Bunnell (2022): "It's Time to be disgusting about COVID-19: Effect of disgust priming on COVID-19 public health compliance among liberals and conservatives," PLOS One, 17.

Altonji, J., C. Huang, and C. Taber (2015): "Estimating the Cream Skimming Effect of School Choice," Journal of Political Economy, 123, 266–324.

Baker, A., D. Larcker, and C. Wang (2021): "How Much Should We Trust Staggered Difference-In-Differences Estimates?" European Corporate Governance Institute – Finance Working Paper, 736.

Barbour, J., L. Rintamaki, J. Ramsey, and D. Brashers (2012): "Avoiding Health Information," Journal of Health Communication, 17, 212–229.

Bhargava, S., G. Loewenstein, and J. Sydnor (2017): "Choose to Lose: Health Plan Choices from a Menu with Dominated Option," Quarterly Journal of Economics, 132.

Borusyak, K., X. Jaravel, and J. Spiess (2023): "Revisiting Event Study Designs: Robust and Efficient Estimation," arXiv Working Paper.

Brot-Goldberg, Z., T. Layton, B. Vabson, and A. Wang (2021): "The Behavioral Foundations of Default Effects: Theory and Evidence from Medicare Part D," NBER Working Paper 28331.

Cairns, C., J. Ashman, and K. Kang (2021): "Emergency Department Visit Rates by Selected Characteristics: United States, 2018," National Center for Health Statistics, CDC online.

Callaway, B. and P. Sant'Anna (2021): "Difference-in-Differences with multiple time periods," Journal of Econometrics, 225, 200–230.

Card, D., C. Dobkin, and N. Maestas (2008): “The Impact of Nearly Universal Insurance Coverage on Health Care Utilization: Evidence from Medicare,” American Economic Review, 98.

Card, D. and J. Rothstein (2007): “Racial segregation and the black-white test score gap,” Journal of Public Economics, 91, 2158–2184.

CDC (2020): “Reasons for Being Uninsured Among Adults Aged 18–64 in the United States, 2019,” National Health Interview Survey.

Chetty, R., J. Friedman, N. Hilger, E. Saez, D. W. Schanzenbach, and D. Yagan (2011): “How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project Star,” The Quarterly Journal of Economics, 126, 1593–1660.

Cohen, R., M. Martinez, A. Cha, and E. Terlizzi (2021): “Health Insurance Coverage: Early Release of Estimates From the National Health Interview Survey, January–June 2021,” .

Congressional Research Service (2023): “U.S. Health Care Coverage and Spending,” Congressional Research Service, In Focus, online.

Cox, C. and D. McDermott (2021): “A Closer Look at the Uninsured Marketplace Eligible Population Following the American Rescue Plan Act,” .

Cullen, J., B. Jacob, and S. Levitt (2006): “The Effect of School Choice on Participants: Evidence From Randomized Lotteries,” Econometrica, 74, 1191–1230.

Deming, D. (2011): “Better Schools, Less Crime?” The Quarterly Journal of Economics, 126, 2063–2115.

Deming, D., J. Hastings, T. Kane, and D. Staiger (2014): “School Choice, School Quality, and Postsecondary Attainment,” American Economic Review, 104, 991–1013.

Dobkin, C., A. Finkelstein, R. Kluender, and M. Notowidigdo (2018): “The Economic Consequences of Hospital Admissions,” American Economic Review, 108, 308–352.

Epple, D. and R. Romano (1998): “Competition between Private and Public Schools, Vouchers, and Peer-Group Effects,” The American Economic Review, 88, 33–62.

Erwin, D., B. Kelley, and G. Silva-Padron (2022): “50-State Comparison: Open Enrollment Policies,” Education Commission of the States.

Exley, C. and J. Kessler (2021): “Information Avoidance and Image Concerns,” NBER Working Paper.

Fabregas, R. (2020): “Trade-offs of Attending Better Schools: Regression Discontinuity Evidence from Mexico,” R&R The Economic Journal.

Fisman, R., K. Gladstone, I. Kuziemko, and S. Naidu (2020): "Do Americans want to tax wealth? Evidence from online surveys," Journal of Public Economics, 188.

Frankenberg, E., J. Ee, J. Ayscue, and G. Orfield (2019): "Harming Our Common Future: America's Segregated Schools 65 Years After Brown," .

Geifer, K. and M. King (2021): "One in Six Older Americans Received Needs-Based Assistance Even Before Pandemic," America Counts: Stories.

Goldin, J., I. Lurie, and J. McCubbin (2021): "Health Insurance and Mortality: Experimental Evidence from Taxpayer Outreach," The Quarterly Journal of Economics, 136.

Golman, R., D. Hagmann, and G. Loewenstein (2017): "Information Avoidance," Journal of Economic Literature, 55.

Golman, R., G. Loewenstein, A. Molnar, and S. Saccardo (2021): "The Demand for, and Avoidance of, Information," Management Science.

Goodman-Bacon, A. (2021): "Difference-in-differences with variation in treatment timing," Journal of Econometrics, 225, 254–277.

Handel, B. (2013): "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," American Economic Review, 103, 2643–2682.

Handel, B., J. Kolstad, T. Minten, and J. Spinnewijn (2020): "The Social Determinants of Choice Quality: Evidence from Health Insurance in the Netherlands," Working paper.

Hastings, J., C. Neilson, and S. Zimmerman (2012): "The Effect of School Choice on Intrinsic Motivation and Academic Outcomes," NBER Working Paper.

Hastings, J., T. K. and D. Staiger (2006): "Preferences and Heterogeneous Treatment Effects in a Public School Choice Lottery," NBER Working Paper.

Hertwig, R. and C. Engel (2016): "Homo Ignorans: Deliberately Choosing Not to Know," Perspectives on Psychological Science, 11.

Holt, C. and S. Laury. (2002): "Risk Aversion and Incentive Effects," The American Economic Review, 92.

Howell, W., P. Wolf, D. Campbell, and P. Peterson (2002): "School Vouchers and Academic Performance: Results from Three Randomized Field Trials," Journal of Policy Analysis and Management, 21, 191–217.

Institute of Medicine (2001): Coverage Matters: Insurance and Health Care, National Academy Press.

Kairies-Schwarz, N., J. Kokot, M. Vomhof, and J. Wessling (2017): "Health insurance choice and risk preferences under cumulative prospect theory - an experiment," Journal of Economic Behavior Organization, 137, 374–397.

Kennedy, R., S. Clifford, T. Burleigh, P. Waggoner, R. Jewell, and N. Winter (2020): "Consumer Behaviors Among Individuals Enrolled in High-Deductible Health Plans in the United States," Political Science Research and Methods, 8, 614–629.

Kiefer, M., J. Silverman, B. Young, and K. Nelson (2014): "National Patterns in Diabetes Screening: Data from the National Health and Nutrition Examination Survey (NHANES) 2005–2012," J Gen Intern Med, 30, 612–618.

Kullgren, J., B. Cliff, C. Krenz, H. Levy, B. West, A. Fendrick, J. So, and A. Fagerlin (2019): "A Survey Of Americans With High-Deductible Health Plans Identifies Opportunities To Enhance Consumer Behaviors," Health Affairs, 38.

Lee, S., R. Nishimura, P. Burton, and R. McCammon (2021): "HRS 2016 Sampling Weights," The University of Michigan Health and Retirement Study.

Liu, Y., Y. Hao, and Z. Lu (2022): "Health Shock, Medical Insurance and Financial Asset Allocation: Evidence from CHFS in China," Health Economics Review, 12.

Loewenstein, G. a. J. F., B. McGill, S. Ahmad, S. Linck, S. Sinkula, J. Beshears, J. Choi, J. Kolstad, D. Laibson, B. Madrian, J. List, and K. Volpp (2013): "Consumers' misunderstanding of health insurance," Journal of Health Economics, 32, 850–862.

Mikulecky, M. (2013): "Open Enrollment is on the Menu–But Can You Order It?" .

Miller, S. (2012): "The Effect of Insurance on Emergency Room Visits: An Analysis of the 2006 Massachusetts Health Reform," Journal of Public Economics, 96, 893–908.

Murphy, R. and F. Weinhardt (2020): "Top of the Class: The Importance of Ordinal Rank," Review of Economic Studies, 87, 2777–2826.

Oster, E., I. Shoulson, and E. Dorsey (2013): "Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease," The American Economic Review, 103.

Perez, C., M. Mejia, and H. Johnson (2023): "Immigrants in California," .

Polyakova, M. (2016): "Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D," American Economic Journal: Applied Economics, 8.

Roth, J., P. H. C. Sant'Anna, A. Bilinski, and J. Poe (2023): "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," arXiv, 2201.01194.

Sacerdote, B. (2011): "Peer Effects in Education: How Might They Work, How Big Are They and How Much Do We Know Thus Far?" Handbook of the Economics of Education, 3.

Samek, A. and J. Sydnor (2020): "Impact of Consequence Information on Insurance Choice," NBER Working Paper.

Schram, A. and J. Sonnemans (2011): "How individuals choose health insurance: An experimental analysis," European Economic Review, 55, 799–819.

Sohn, H. (2017): "Racial and Ethnic Disparities in Health Insurance Coverage: Dynamics of Gaining and Losing Coverage over the Life-Course," Popul Res Policy Rev, 36.

Sun, L. and S. Abraham (2021): "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," Journal of Econometrics, 225, 175–199.

Sweeny, K., D. Melynk, W. Miller, and J. Shepperd (2010): "Information Avoidance: Who, What, When, Why," Review of General Psychology, 14, 340–353.

Sydnor, J. (2010): "(Over)insuring Modest Risks," American Economic Journal: Applied Economics, 2, 177–199.

Tolbert, J., P. Drake, and A. Damico (2022): "Key Facts about the Uninsured Population," .

Urquiola, M. (2005): "Does School Choice Lead to Sorting? Evidence from Tiebout Variation," American Economic Review, 95.

Wagstaff, A. (2005): "The Economic Consequences of Health Shocks," World Bank Policy Research Working Paper.