Essays on Human Behavior

by

M. Steven Holloway

A dissertation accepted and approved in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in Economics

Dissertation Committee:

Michael Kuhn, Chair

Edward Rubin, Core Member

Jonathan Davis, Core Member

Ryan Light, Institutional Representative

University of Oregon

Summer 2024

DISSERTATION ABSTRACT

M. Steven Holloway

Doctor of Philosophy in Economics

Title: Essays on Human Behavior

This dissertation explores human behaviors that may confound current economic measurements in the environmental and behavioral contexts.

Chapter 2 examines disparate wildfire-smoke avoidance responses across demographic and socioeconomic groups in the western United States. Cellphone location data is combined with census block group level data to uncover heterogeneity in communities' responses to smoke. This heterogeneity is found to be correlated with historical margins of disadvantage. These communities have in many cases been documented to face higher levels of exposure to other pollutants and are more negatively impacted by these exposures. This suggests that those that have a higher benefit of avoiding wildfire-smoke also face higher costs of avoidance, leading to substantial equity concerns in terms of climate adaptation.

Chapter 3 focuses on the presentation and development of theory on time-varying time preferences, results from a novel experiment designed to measure these preferences, and structural estimation and evaluation of models allowing for these preferences. While standard economic models acknowledge the agents' ability to exhibit preference reversals, these reversals are largely attributed to a present- or future-bias in agents' underlying preferences. However, instead of myopic reversals, agents may instead be making decisions according to time preferences that drift across measurement dates. To test for this, an experiment that records subjects' decisions at seven points in time across twelve weeks is conducted. It is found that time-varying preferences are present in the sample studied and that models allowing for these preferences outperform traditional and behavioral discounting models in both aggregate fit and number of subjects best described.

Chapter 4 builds on the analysis conducted in Chapter 3 to determine what margins drive the time-varying behavior observed. Characteristics of subjects such as their demographics,

socioeconomic status, risk preferences, cognition, and other behavioral factors, along with dynamic data on subjects' financial and psychological well-being were collected as part of the experiment conducted in Chapter 3. While little evidence of correlation between demographics or preferences and discounting is found, subjects do accurately perceive their time-variance. This suggests time-variance may itself be a preference that can be modeled, or that subjects can learn about their discounting process over time. In an attempt to capture this, the theory developed in Chapter 3 is used to develop a discount function that asymptotically approaches time-invariance. This model does not outperform the non-asymptotic, time-invariant, behavioral discounting model it approaches in the limit for the aggregate data. However, it does better explain a large portion of our subjects' behavior.

This dissertation includes unpublished co-authored material.

ACKNOWLEDGEMENTS

I would like to thank my advisor, Michael Kuhn, for his constant support in the process of producing this dissertation. His expertise, advise, and shared passion made the latter half of this dissertation possible. I sincerely appreciate the immense time and effort he has committed to our project and my mentorship. His demeanor and mentorship style serve as an inspiration in my academic career.

I would also like to acknowledge my co-author, Ed Rubin, whose time and effort made the production of the second chapter of this dissertation an enjoyable learning experience. I am truly grateful to have gained insights from his expertise and for his kindhearted, supportive presence in my time at the University of Oregon.

In addition, I would like to acknowledge Jonathan Davis and Ryan Light for their time and effort as part of my committee. Their comments have helped to shape this dissertation, as well as projects outside its scope and I am thankful for their feedback.

I also thank my co-author, Connor Wiegand, for his instrumental contributions to our shared work and for his friendship throughout the course of the program and beyond.

Finally, I extend my most heartfelt thanks to my partner, Rachel Simons, whose unconditional love and support has been a constant in my time here. Her encouragement, belief, and excitement continually instill the confidence required to achieve my goals.

TABLE OF CONTENTS

Chapter Page

LIST OF FIGURES

Figure Page

Figure Page

LIST OF TABLES

Table Page

Table	Page

CHAPTER I

INTRODUCTION

Tractability and intuition of economic models often supersedes generalizability in their development. While this leads to digestible, and often accurate, models of aggregate human behavior, they may overlook the finer details of the decision-making processes and the subject-level impacts that drive aggregate results. The literature on Environmental Justice and Behavioral Economics both approach this issue in different ways. Environmental Justice documents the inequity of outcomes of environmental policy and stimuli, while Behavioral Economics augments traditional economic models by integrating psychological and social forces otherwise not considered in economic analysis. This dissertation adds to these literatures by first exploring the heterogeneous impacts of an environmental stimuli and then documenting, addressing, and attempting to ameliorate the ramifications of a longstanding oversight in intertemporal decision-making.

Chapter 2 investigates historically marginalized communities in the western United States' responses to an increasingly prevalent pollutant, wildfire-smoke. Using cellphone location data, combined with census block group level (CBG) Census data and twice-daily smoke plume geometry, we find that avoidance of wildfire smoke is heterogeneous across communities and correlated with race and income. In particular those from communities with above median average incomes, as well as those with above median shares of White population, exhibit more activity away from home and travel further when exposed to smoke relative when not exposed to smoke. This pattern does not hold for those with below median average incomes and those with above median shares of Black or Hispanic populations. Instead, these populations do not respond on average, or exhibit even less activity away from their home CBG when exposed to smoke relative when not exposed to smoke. The interaction of this inequity and the well-documented more negative impact of other pollutants on these communities suggests a compounding effect that leaves disadvantaged communities with higher benefits of escaping wildfire-smoke while facing even higher costs. Without the acknowledgement of the costs of disadvantage or the potential for heterogeneous impact of commensurate pollution exposure, this disparity would go unobserved.

In Chapter 3, we document and further develop the theory of intertemporal discount functions that allow for systematic time-variance of preferences and present the results of a novel longitudinal experiment designed to measure these preferences. We develop an extension of the standard exponential discount function that can accommodate time-variance, present-bias, and novel combinations of the two in the interest of generalizing the modeling of intertemporal discounting with limited up-take cost. The empirical contribution is an unprecedented measurement of individual time preferences: we survey individuals seven times over 12 weeks to observe if time-variance is present, and to compare to the limited previous results of studies interested in time-variance that based their measurements on only one or two points in time. We characterize time preferences of both the aggregate subject pool and individuals themselves using both reduced form and structural models. We find that aggregate preferences and a majority of individuals' preferences are best described by discount functions that allow for time-variance, and that in particular, subjects become more impatient over time. This possibility is assumed away in traditional economic models, including existing behavioral economics models, underlining the need for more general models of time preference that are more inclusive of valid choice profiles. This is most evident when we consider the relationship between "time-inconsistency" – a specific manifestation of present bias that is often equated with it—and time-variance. We find that nearly three-quarters of time-inconsistency measured in the data can be explained by time-variance, and that if the decreasing patience observed in our study is also present in the subjects of similar studies measuring time-inconsistency, that the prevalence of present-bias in the literature may be overstated on both the intensive and extensive margins.

In Chapter 4, I continue to study time-variance in the sample from the experiment in Chapter 3. I investigate the margins on which time-variance differs across subjects and the potential sources of this phenomena. Data collected as part of the experiment included static measurements of subjects' demographics, socioeconomic status, risk preferences, patience, and other behavioral characteristics, as well as dynamic data on financial and psychological well-being. While financial standing seems to be the biggest driver of discounting behavior, poorer finances only result in very minor economic impact in terms of the discount factor applied to a decision. However, subjects do appear to be aware of whether they act time-varyingly. This suggests that, given the existing data, time-variance may be best described via a preference-parameter that allows a subject's discount function to evolve over time. As linear or exponential growth of a discount function, traditionally bound between 0 and 1, is not tractable, An intuitive way to include this parameter is to create a model that is asymptotically time-invariant. I leverage the theory developed in Chapter 3 to develop a discounting model that exhibits time-variance, but in the limit can be considered time-invariant. This model is also consistent with learning about one's preferences as a driver of time-variance. I find that this model better fits the data compared to both the standard exponential and analogous time-invariant discount model.

This dissertation contributes to both the Environmental Justice literature and the Behavioral Economics literature while drawing insights from experimental design and economic modeling, with special thought given to the subjects of our analysis. The insights drawn from each of these chapters have potential for impacts beyond their own fields, including in broader environmental sciences, ecology, sociology, psychology, and public policy. By acknowledging the complexity of human nature and opening doors to its inclusion in economic models this dissertation aims to aid in creating more complete, inclusive, pictures of the economic systems it studies.

CHAPTER II

UNEQUAL AVOIDANCE: DISPARITIES IN SMOKE-INDUCED OUT-MIGRATION

This chapter was co-authored with Assistant Professor Edward Rubin. I contributed to the development of the core idea and analysis methods used, and wrote the first full draft of the paper. Professor Rubin performed the coding to generate results for the full sample after I performed them at smaller scale, created figures and tables, and revised and edited the paper after my first draft.

2.1 Introduction

In the face of imminent danger, organisms often flee.[1] Yet nothing guarantees individuals have equal access to this strategy. Given recent work on inequality in the United States, one may suspect opportunities for avoidance are unequal: A substantial literature documents marked disparities along numerous dimensions—life expectancy and mortality, health, healthcare, pollution exposure, transportation, educational opportunities, and employment outcomes.[2] Indeed, many of these dimensions of inequality represent mechanisms that affect individuals' abilities to flee (liquidity, job security/benefits, access to transit)—and/or consequences of unequal opportunities to relocate (health and life expectancy). If less-privileged communities are less able to avoid hazards, existing inequities may worsen.

A contemporary case sheds light on this phenomenon. Wildfire smoke is an increasingly important and pervasive hazard with potentially avoidable health damages—yet communities' abilities to avoid these damages may diverge. The size of the US-wildfire problem is enormous: between 2018–2021, wildfires cost more than $62 billion (Smith, 2022). In the same years, *every* Census Block Group on the West Coast faced wildfire smoke during at least 14 weeks (Figure 1). As wildfires ravaged larger and larger swaths of the US,[3] individuals increasingly confronted the choice to face fires' smoke or flee.

[1] Avoidance is half of *fight or flight*. Recent IPCC reports highlight the importance of such adaptation on humanity's road to mitigating climate-change damages (IPCC, 2022).

[2] Examples follow. **Life expectancy and mortality:** Alsan, Chandra, and Simon (2021); Chetty, Stepner, et al. (2016); Currie and Schwandt (2016a); Harper, Lynch, Burris, and Smith (2007); Schwandt et al. (2021). Currie and Schwandt (2016b) provides a helpful overview. **Health:** Aizer and Currie (2014); Almond and Chay (2006); Currie (2009). **Healthcare:** Chandra, Kakani, and Sacarny (2020); Dieleman et al. (2021); Singh and Venkataramani (2022). **Transportation:** Chung, Myers, and Saunders (2001). **Education:** Autor, Figlio, Karbownik, Roth, and Wasserman (2019); Bertocchi and Dimico (2012); Blanden, Doepke, and Stuhler (2022); Elder, Figlio, Imberman, and Persico (2021). **Employment:** Aneja and Avenancio-Leon (2022); Aneja and Xu (2021); Chetty, Hendren, Kline, and Saez (2014); Chetty, Hendren, Lin, Majerovitz, and Scuderi (2016); Davis and Mazumder (2022).

[3] Baylis and Boomhower (2022); Burke et al. (2021); Goss et al. (2020); O'Dell, Ford, Fischer, and Pierce (2019); Radeloff et al. (2005, 2018); Westerling, Hidalgo, Cayan, and Swetnam (2006) document this phenomenon.

The literature documents the considerable health consequences of wildfire-smoke exposure (mortality, morbidity, and birth outcomes)—and the particulate matter associated with the wildfire smoke.[4] Temporary relocation can avoid smoke altogether—reducing smoke-related health costs. However, temporary and unexpected relocation may only be feasible for some (likely more-privileged) households. This inequality in individuals' abilities to relocate may generate unequal damages from smoke exposure—exacerbating existing inequality.

We estimate the effect of wildfire smoke on short-term out-migration and how this smoke-induced migration varies by communities' racial, ethnic, or income compositions—shedding light on whether communities equally access avoidance strategies. For this analysis, we combine remotely sensed smoke plumes, cellphone-based movement data, and CBG-level demographic data for all recorded wildfire smoke in the US's West Coast during 2018–2021. Our results show that, on average, individuals temporarily relocate when they face wildfire smoke. Residents travel farther and are more likely to leave their home counties when facing smoke.

However, these first results hide substantial heterogeneity. We find that historically marginalized populations—Black, Hispanic, and low-income communities—are significantly less likely to out-migrate when facing the same wildfire smoke as more privileged populations. This heterogeneity mirrors other inequality in individuals' travel habits when smoke is absent—and many already-documented inequalities. Together, these results illustrate fundamental inequities in society's ability to respond to major risks/damages and suggest potentially fruitful avenues for policy.

Our results contribute to three strands of the literature: (1) environmental justice/inequality, (2) defensive investments and avoidance/adaptation, and (3) the economics of wildfires.

First, our results bring new insights into the burdens facing lower-income, Black, and Hispanic communities. A large environmental justice (EJ) and inequality literature has documented numerous dimensions along which historically marginalized communities face worse environmental quality—e.g.in exposure to toxic-release facilities (Bullard, Mohai, Saha, & Wright, 2008; Mohai & Saha, 2007), air pollution (Clark, Harris, Apte, & Marshall, 2022; Colmer, Hardman, Shimshack, & Voorheis, 2020; Hsiang, Oliva, & Walker, 2019), and noise (Casey et

[4] Burke et al. (2022); O'Dell et al. (2019) link wildfire smoke to $PM_{2.5}$. See (Cascio, 2018; Dedoussi, Eastham, Monier, & Barrett, 2020; Heft-Neal, Driscoll, Yang, Shaw, & Burke, 2022; Kochi, Champ, Loomis, & Donovan, 2012; Kondo et al., 2019; Liu et al., 2017; Richardson, Champ, & Loomis, 2012; Wen & Burke, 2022) for examples of the costs of smoke and/or $PM_{2.5}$ exposure.

al., 2017). Much of this literature focuses on documenting unequal exposure to environmental hazards. Our findings complement this EJ thread by showing even when external (outdoor) exposure is ‘equitable’—wildfire smoke covers large areas—adaptation/defensive responses may be unequal. In particular, we show that when wildfire smoke covers an area, historically disadvantaged communities *in that area* are less likely to out-migrate relative to more affluent and more White communities.

We also contribute to a growing literature that documents and measures avoidance behaviors and defensive investments individuals employ against environmental risks. Avoidance strategies in this literature include consumption choices,[5] structural investments,[6] long-term migration,[7] and short-term travel.[8] Short-term travel and equity have received less attention—likely due to historical data limitations.[9][10] Our results help fill this gap—estimating the extent of short-term out-migration and its distribution across socioeconomic groups. Moretti and Neidell (2011) and Deschênes et al. (2017) both find that costs related to avoidance activities and defensive investments are on the same scale as the health costs of exposure—i.e.the health costs mitigated by avoidance may be quite large. Consequently, if avoidance strategies are mainly available to (or employed by) more-advantaged communities, less-advantaged communities may bear substantial and disproportionate shares of the health burden of exposure. Our results suggest that this concern is legitimate.

Finally, we contribute to the literature on the economics of wildfires. Recent work raises equity concerns in the allocation of fire-fighting resources (Anderson, Plantinga, & Wibbenmeyer, 2022; Lennon, 2022; Plantinga, Walsh, & Wibbenmeyer, 2022), the incidence of wildfire hazard (Wibbenmeyer & Robertson, 2022), the incidence of fire suppression costs (Baylis & Boomhower,

[5] *E.g.*, purchases of water bottles (Graff Zivin, Neidell, & Schlenker, 2011), pharmaceuticals (Deschênes, Greenstone, & Shapiro, 2017), masks (Sun, Kahn, & Zheng, 2017; Zhang & Mu, 2018), and air purifiers (Ito & Zhang, 2020)

[6] For example, Baylis and Boomhower (2021).

[7] For examples of long-term migration, see Bayer, Keohane, and Timmins (2009); Boustan, Kahn, and Rhode (2012); S. Chen, Oliva, and Zhang (2022); Freeman, Liang, Song, and Timmins (2019); Hornbeck and Naidu (2014); Khanna, Liang, Mobarak, and Song (2021); Tan Soo (2018); Zheng and Kahn (2008).

[8] For examples of short-term travel-based avoidance, see Burke et al. (2022); S. Chen, Chen, Lei, and Tan-Soo (2021); Graff Zivin and Neidell (2009); Neidell (2009); Richardson et al. (2012).

[9] S. Chen et al. (2022) show younger, more-educated individuals are more likely to permanently migrate due to pollution. Sun et al. (2017) find more affluent individuals are more likely to make defensive investments—masks.

[10] Our (short-run) migration results also relate to the climate-adaption literature—e.g.Barreca, Clay, Deschênes, Greenstone, and Shapiro (2015); Barreca, Clay, Deschenes, Greenstone, and Shapiro (2016); Burke and Emerick (2016); Carleton et al. (2022); Deschênes (2014); Deschênes and Greenstone (2011); Lai, Li, Liu, and Barwick (2022); Massetti and Mendelsohn (2018).

2022), and the burden of wildfire smoke (Borgschulte, Molitor, & Zou, 2022). Limited previous work exists on avoidance behaviors in the setting of wildfires. In one exception, Richardson et al. (2012) provide survey-based evidence of averting actions from sample respondents after the 2009 Station Fire in Los Angeles County, California. Our results merge these two branches of the wildfire literature—equity and avoidance.

Our results on wildfire-smoke avoidance behavior are most similar to Burke et al. (2022). Burke et al. show (1) individuals are aware of smoke exposure (via Google searches), (2) Google searches suggest individuals are seeking protection from smoke (e.g. "air purifier" and "smoke mask"), (3) people are less happy when facing smoke (Twitter sentiment), and (4) smoke-based fine-particulate ($PM_{2.5}$) increases individuals' likelihood of remaining at home for the entire day. This fourth result from Burke et al. (2022) is closest to our central research question—whether out-migration to avoid smoke is equally accessed across socioeconomic groups.

While similar, our approach and results differ from Burke et al. (2022) on several important dimensions—complementing their analysis. Most notably, we document a strong relationship between out-migration and income; Burke et al. find little evidence that income correlates with communities' likelihoods of staying home.

These different conclusions likely stem from several differences in our empirical approaches. First, Burke et al. (2022) focus on counties, while our analysis focuses on Census Block Groups (CBGs). As we discuss below, the substantially higher spatial resolution afforded by CBGs allows a finer match between communities and incomes—resolving error from aggregation/the ecological fallacy (Anderton, Anderson, Oakes, & Fraser, 1994; Banzhaf, Ma, & Timmins, 2019). We also aggregate across days in a week rather than focusing on day-level outcomes as Burke et al.. Day-level timing matches Burke et al.'s goal of mapping daily smoke-induced $PM_{2.5}$ to social outcomes. However, by aggregating across days, we can capture short-term intertemporal substitution of travel[11]—as found in Graff Zivin and Neidell (2009). Intertemporal variation may be more relevant for testing our hypothesis of short-term out-migration. Our exposure variables also differ: We focus on all 'wildfire smoke,' whereas Burke et al. examine $PM_{2.5}$ generated by wildfire smoke. $PM_{2.5}$ undoubtedly represents a significant concern for public health. However, our goal is estimating the effect of smoke *itself*—rather than hazardous particulates caused by smoke. Finally, our outcomes differ. We measure out-migration as the share of a CBG's trips that leave the county or the 75$^{\text{th}}$ percentile of distance traveled

[11] *E.g.*, temporarily delaying visits to the zoo.

by the CBG's residents; Burke et al. focus on individuals that remain home or are absent from their homes for the entirety of the day. These differences in the unit of analysis and definitions of outcomes provide a complementary view of smoke-induced out-migration.

Together, our spatial disaggregation and temporal aggregation allow us to contribute to this literature's understanding of smoke-induced out-migration—finding significant evidence that income, race, and ethnicity correlate with CBGs' tendencies to out-migrate.

2.2 Data and Descriptive Statistics

For our analysis, we construct a panel that measures CBG-level weekly smoke exposure, out-migration, and demographics.

CBGs provide an ideal unit of analysis for several reasons. First, they are the smallest unit at which we can obtain cellphone-based movement data—our out-migration measure—and some socioeconomic data. Second, most CBGs delineate small geographic areas containing small populations (typically 600–3,000 individuals). Smaller areas help reduce aggregation-related errors when we assign smoke exposure, migratory behavior, and socioeconomic measurements to entire CBGs.

To maximize this 'match'—i.e.to keep CBGs' areas small—we focus on urban CBGs—where urban population exceeds rural population. As Figure 1 illustrates, rural CBGs can be quite large (e.g.southeastern Oregon), complicating smoke-exposure assignment. The panel of urban CBGs covers 27,555 in California, Oregon, and Washington through 2018-2021. These urban CBGs represent 91.3% of the US's west-coast CBGs, 93.3% of the population (47.1 million people), and 93.9% of visits (3 billion).[12]

2.2.1 Movement Data. We measure communities' short-term migration patterns using SafeGraph's aggregated and anonymized cellphone-movement data (SafeGraph, 2022b). Specifically, we use SafeGraph's *Weekly Patterns* dataset, which monitors 45 million cellphones' 'visits' to 3.6 million Points-of-Interest (POIs). A POI represents any *visitable* location—e.g.restaurants, schools, parks, doctors' offices. Across the 30,174 west-coast CBGs in our data, we observe 3.2 billion visits to POIs during 2018–2021. SafeGraph uses internal microdata (similar to data in M. K. Chen and Rohla (2018)) to predict a home CBG for each cellphone. The Weekly Patterns dataset contains the number of visits to each POI by the visitors' home CBGs—during each week. From these counts, we calculate our main measure of out-migration: the percentage of a CBG's visits (each week) that occurred outside of the CBG's county.

[12] Tables A1 and A2 summarize the three datasets described below for *urban* west-coast CBGs and *all* west-coast CBGs—by CBG (Panel A) and by CBG-week (Panel B).

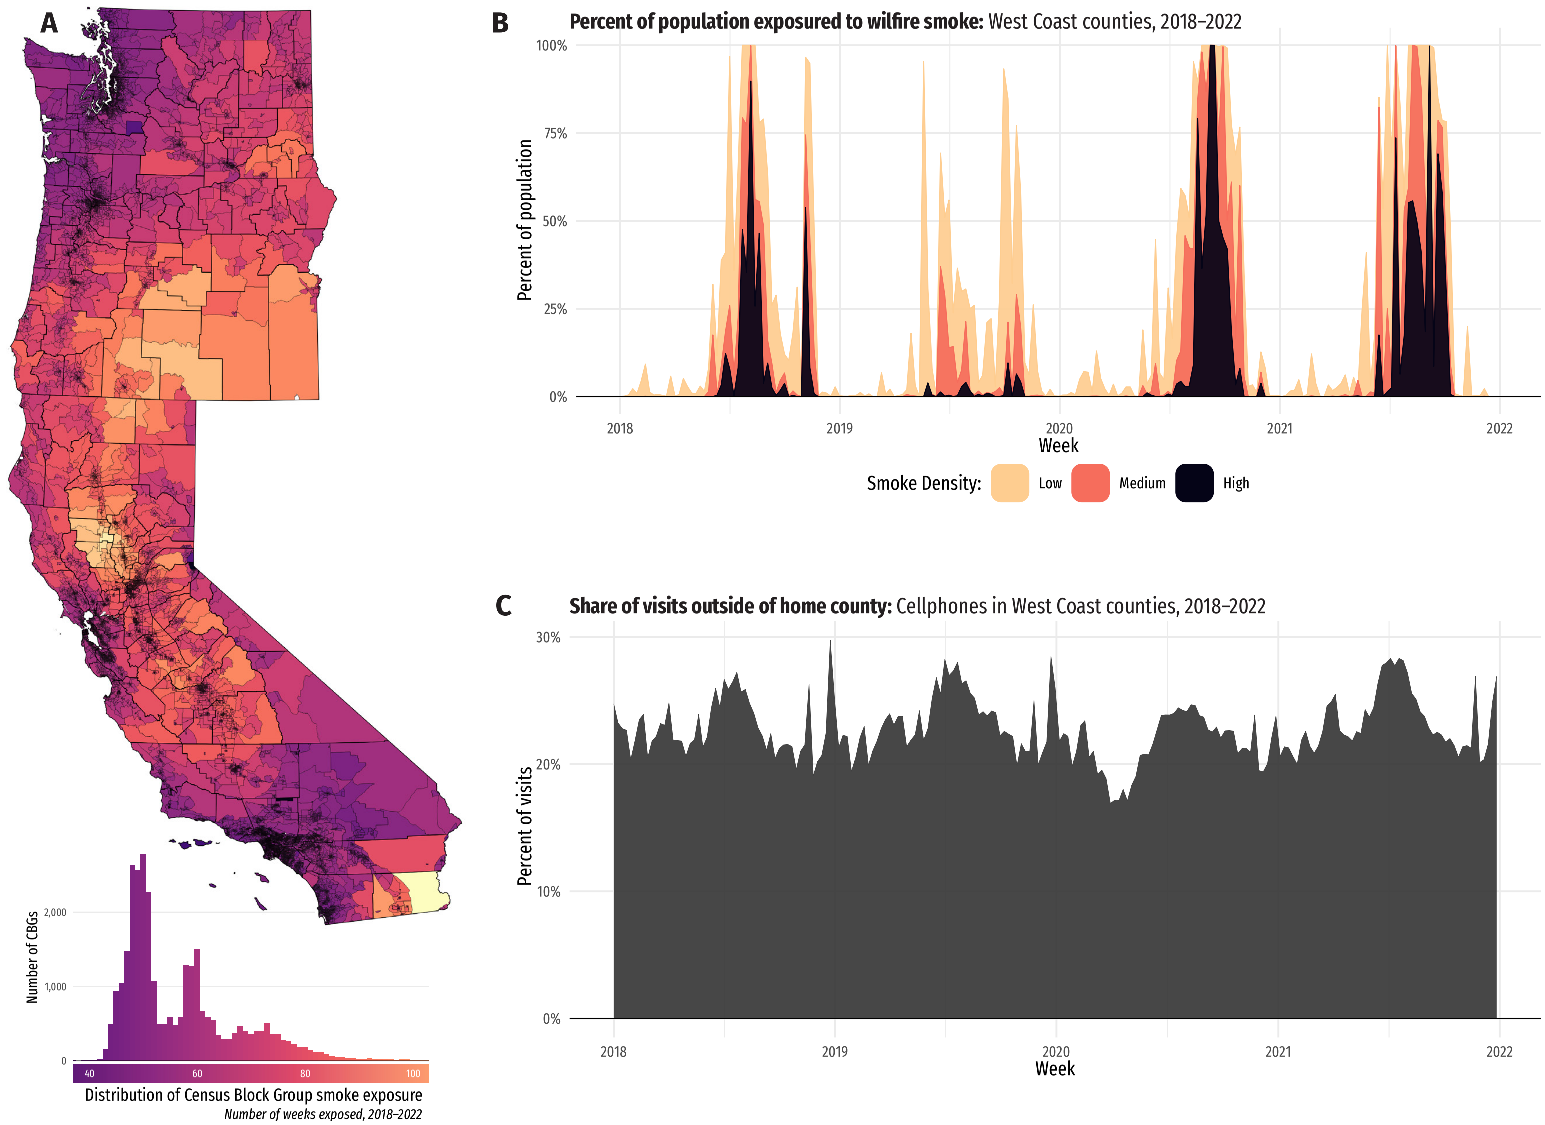

Figure 1. Smoke Exposure and Mobility

Panel **A** illustrates the study area—the West Coast of the United States. The smallest features in the map in **A** are shaded in proportion to the number of weeks that Census Block Groups (CBGs) encountered wildfire smoke 2018–2021. Panel **B** depicts the share of the Western US population exposed to three smoke densities by week. Panel **C** shows each week's share of cellphone-based movement that occurred outside of individuals' home counties.

Our second out-migration measure uses the 75$^{\text{th}}$ percentile of distances traveled by a CBG's residents—measured between CBG and POI centroids—to proxy for the distance traveled away from home by residents of each CBG each week. Panel C of Figure 1 (and Table A1) shows approximately 22% of POI visits occur outside individuals' home counties. Thus, the 75$^{\text{th}}$-percentile distance measure allows us to measure how much farther people are traveling due to smoke—focusing on the part of the distance distribution likely to be affected. Together, these data and calculations provide a unique, spatially resolved view of communities' weekly travel behaviors across four years.

Our cellphone-based movement data offer several strengths for our analysis relative to more traditional datasets. First, movement data provide insights into human behaviors that are largely unavailable—particularly at the scale (10% of the smartphone market) and frequency (all day, every day) of the SafeGraph data. Second, their scale and frequency generate sufficient statistical power to estimate unequal/heterogeneous responses to infrequent events. Answering equity questions about wildfire-smoke exposure requires this power. Finally, the data are *revealed* behaviors—likely suffering less from recall or dishonesty. The strengths of these data are evidenced by the volume of recently published studies that use them.[13]

Cellphone-based movement data are not without concerns and limitations. Some issues relate to the strength of the data—many authors raise ethical and practical concerns for cellphone-based data (Valentino-Devries, Singer, Keller, & Krolik, n.d.).[14] To address some privacy concerns, SafeGraph does not distribute microdata and applies differential-privacy techniques to many features of their aggregated data.[15] Another common concern is external validation—how representative is this sample of 45 million cellphone users? SafeGraph's internal calculations suggest the sample of phones is reasonably balanced at the CBG for income, race, and ethnicity (Squire, 2019a, 2019b). More conservatively: the sample is internally valid for the 45 million users in the dataset—a sizable share of the adult population in the US. In our context,

[13]Recent publications include Allcott et al. (2020); Burke et al. (2022); M. K. Chen, Haggag, Pope, and Rohla (2022); M. K. Chen and Rohla (2018); Dave, Friedson, Matsuzawa, and Sabia (2020); Fajgelbaum, Khandelwal, Kim, Mantovani, and Schaal (2021); Glaeser, Gorback, and Redding (2022); Goolsbee and Syverson (2021); "https://doi.org/10.1086/713787" (2021); Ilin et al. (2021); Jay et al. (2020); Long, Chen, and Rohla (2020); Parolin and Lee (2021); Tai, Mehra, and Blumenstock (2022); Weill, Stigler, Deschenes, and Springborn (2020); Zhao, Holtz, and Aral (2021).

[14]Some prominent early critics of cellphone-based data later published work using cellphone-based data (see S. A. Thompson and Warzel (n.d.) and later Y. Thompson Stuart A. Serkez and Kelley (n.d.)).

[15]Appendix-section Privacy, Noise, and Censoring in SafeGraph Weekly Patterns Data describes SafeGraph's differential-privacy approach and why it is unlikely to affect our results substantially.

the costs associated with aggregated and already-available movement data seem small; with these data, the study is possible.

2.2.2 Smoke Exposure Data. We calculate CBGs' weekly smoke exposures using smoke-plume shapefiles from the US National Oceanic and Atmospheric Administration (NOAA) Hazard Mapping System (HMS) Fire and Smoke Product (NOAA Office of Satellite and Product Operations, 2022). These publicly available data provide daily records of smoke-plume boundaries across North America throughout the sample period.[16] We consider CBG i to be exposed to smoke in week w if any smoke plumes from week w intersect with a i's boundaries. While coarse (Wen & Burke, 2021), numerous studies find communities' and individuals' HMS-based smoke exposure significantly correlate with human responses—e.g.Borgschulte et al. (2022); Burke et al. (2022); Heft-Neal et al. (2022); Miller, Molitor, and Zou (2017).[17]

Smoke exposure varies substantially throughout the sample period—in both the locations and levels of exposure. Panel B of Figure 1 depicts the percent of the west-coast population exposed to smoke in each week of 2018–2021 (colored by the intensity of smoke). This population-smoke-exposure time series includes substantial variation—ranging from weeks with nearly 0% exposure to several weeks with full-population exposure. Panel A of Figure 1 maps the number of weeks of smoke exposure (2018–2021) for each west-coast CBG. Most CBGs faced smoke during 40–80 weeks throughout the sample—though CBGs in central California and southern/eastern Oregon faced smoke during more than 100 weeks. Together, these figures illustrate the spatiotemporal variation in smoke exposure we use to identify communities' responses to smoke.

2.2.3 Demographic Data. Our data on CBGs' racial, ethnic, and income compositions come from the American Community Survey (ACS) 5-year estimates from 2019.[18] Specifically, for each CBG, we use the population counts of Black, Hispanic, and White individuals and each CBG's median household income. As discussed above, our main analyses focus on 'urban' CBGs, where the urban population exceeds the rural population. Data on rural and urban populations come from 2010 decennial census data from NHGIS (Manson, Schroeder, Van Riper, Kugler, & Ruggles, 2022).

[16] The data originate as satellite imagery from NOAA's/NASA's Geostationary Operational Environmental Satellite System (GOES); NOAA analysts then hand-draw plume boundaries—categorizing smoke density as low, medium, or high.

[17] We use historical data from the Wildland Fire Interagency Geospatial Services (Wildland Fire Interagency Geospatial Service, 2022) to determine CBGs located near past wildfires.

[18] The appendix section Five-Year American Community Survey elaborates on this dataset.

2.3 Results

2.3.1 Empirical Approach. Our goal is to estimate (1) the effect of smoke on short-term migration and (2) how this smoke-induced migration varies by a community's racial, ethnic, or income composition.

Toward this goal, we estimate the model

$$\text{Migration}_{iw} = \beta\,\text{Smoke}_{iw} + \delta\,\text{Smoke}_{iw} \times \text{Percentile}_i + \alpha_i + \gamma_w + \varepsilon_{iw} \tag{2.1}$$

where Migration_{iw} measures the intensity of out-migration among residents of CBG i in week w. As described above, we measure out-migration in two ways. First, we use the percentage of POI visits by residents of CBG i in week w that occur outside of i's county. Our second measure of out-migration is the 75^{th} percentile of distances traveled by CBG i's residents in week w. Together, these two measures illustrate how the composition and shape of CBGs' travel distributions change when their residents face wildfire smoke.[19] Smoke_{iw} represents an indicator variable for whether CBG i encountered *any* smoke during week w. Percentile_i refers to CBG i's rank-based percentile along one of several measurements of CBG i's socioeconomic composition (median household income; or population-share Black, Hispanic, or White). We integrate these percentiles with two alternative specifications. The first specification defines Percentile_i as numeric—i.e.imposing linearity in percentile. Chetty et al. find intergenerational mobility is linear in individuals' income percentiles; several of our results also suggest approximate linearity. However, we also relax this linearity assumption. Specifically, we apply a semi-parametric specification where Percentile_i represents a set of indicators for each mutually exclusive two-percentile group (i.e.indicators that identify the bins [0%, 1%), [1%, 2%), etc.). We present the results for these two approaches in sequence.

CBG-specific fixed effects (FEs) (α_i) absorb time-invariant differences in out-migration across CBGs. Week-of-sample FEs (γ_w) account for out-migration shocks and seasonality common to western CBGs. Our results are robust to various fixed-effects specifications—e.g.replacing α_i with a FE for *CBG by month-of-year*. Finally, ε_{iw} is the error term. Both Smoke_{iw} and other determinants of out-migration (in ε_{iw}) may correlate across weeks within a CBG and across CBGs in a given week. To account for this correlation in our inference, we estimate cluster-robust standard errors that allow for correlation within county and within calendar months (e.g.July).

[19] We weight observations (CBG i in week w) by CBGs' populations. Because CBG populations are not uniform, this weighting enables us to draw inferences on the population of individuals—rather than the population of CBGs.

The parameters β and δ directly map to our central empirical questions. If $\beta > 0$, then smoke increases the out-migration. If $\delta \neq 0$, communities differ in their out-migration behavior along the dimension given by Percentile$_i$: larger δs imply greater tendencies to out-migrate in the presence of smoke.

A causal interpretation of β—the effect of smoke on out-migration—requires that CBG i's smoke exposure in week w is independent of other determinants of i's movement in week w, conditional on the fixed effects. This requirement is plausible for several reasons. First, the FEs remove issues from seasonality (e.g.summer vacation) and cross-sectional differences (e.g.affluent, fire-prone areas), ruling out many potential confounds. Second, the sources of the smoke plumes (wildfires) are highly unpredictable in time and space (otherwise, they typically would not occur). Further, the smoke from these wildfires is carried toward or away from communities by erratic meteorological variation. Thus, for our identifying requirement to be violated, there must be a latent factor that (1) only appears when communities are downwind of wildfires—the result of unpredictable and semi-random processes—and (2) induces increased out-migration (but cannot be caused by smoke itself). Finally, event studies centered on communities' first smoke exposures of each calendar year (pooled across events and years) show sharp increases in out-migration beginning in the first week of smoke exposure (Figure B1). The event studies suggest increases in (A) the share of out-of-county visits, (B) the 75^{th} percentile of distance traveled, and (C)the number of visits outside of individuals' home counties. Based upon this evidence and reasoning, we believe this identifying assumption is reasonably plausible.

2.3.2 Population-wide Response to Smoke. We first estimate Equation 2.1 without interactions/heterogeneity—identifying the average response to smoke exposure, pooled across all urban CBGs.

Column (1) of Table 1 demonstrates that, on average, communities increase out-migration when they face wildfire smoke. The two panels of the table separate results for the two dependent variables. In Panel A, the dependent variable is the percentage (0–100) of POI visits from a CBG's residents that occur outside the CBG's county. Panel B's outcome is the 75^{th} percentile of distance traveled to POIs. Each column in each panel results from a separate regression with the same fixed-effect specification—CBG and week-of-sample.[20] The standard-error estimator allows clustering within county and month-of-year (e.g.January).

[20] Table A3 reproduces the estimates in Table 1 but uses state by week-of-sample fixed effects. Table A4 drops CBGs affected by wildfires between 2018-2021. Results across the three tables are very similar.

Table 1. Regression Results: Out-migration Responses to Smoke and Distribution

		Percentile-based heterogeneity			
	(1)	(2) HH Income	(3) % Black	(4) % Hispanic	(5) % White
Panel A *Dependent variable: Percent of POIs visits outside of home county*					
Any smoke	0.28**	-0.41	0.54***	0.89**	-0.10
	(0.10)	(0.30)	(0.15)	(0.39)	(0.21)
Any smoke × Het. percentile		1.4**	-0.49*	-1.2	0.80*
		(0.53)	(0.24)	(0.70)	(0.38)
N obs. (millions)	5.63	5.63	5.63	5.63	5.63
R^2	0.73	0.73	0.73	0.73	0.73
Panel B *Dependent variable: 75^{th} percentile of distance traveled to POIs (km)*					
Any smoke	1.7**	-6.9	6.0***	11.5*	-3.4
	(0.55)	(4.0)	(1.9)	(5.9)	(2.3)
Any smoke × Het. percentile		17.2**	-8.1*	-18.8	10.7*
		(7.6)	(3.7)	(11.4)	(4.9)
N obs. (millions)	5.63	5.63	5.63	5.63	5.63
R^2	0.079	0.079	0.079	0.079	0.079
Fixed effects					
CBG	✓	✓	✓	✓	✓
Week of sample	✓	✓	✓	✓	✓

Notes: Panel **A** estimates the effect of smoke exposure on the percent (0–100) of POI (SafeGraph place of interest) visits that occur within visitors' home counties; Panel **B** estimates the effect of smoke exposure on the 75^{th} percentile of distance traveled to POIs. Columns (2–5) estimate heterogeneity by CBGs' percentile (0–1) of household income, % Black, % Hispanic, and % White. Each column in each panel represents a separate regression—using the same fixed-effect specification of CBG and week-of-sample. Observations are weighted by CBG population. Table A3 reproduces the current table with *state by week-of-sample* fixed effects. Standard errors allow clustering within county and month-of-year (e.g.January). Significance codes: ***: 0.01, **: 0.05, *: 0.1.

In Column (1) of Panel A (Table 1), we estimate that smoke significantly increases a CBG's share of out-of-county POI visits by 0.28 percentage points. On average, approximately 22% of POI visits occur outside of residents' home counties (Table A1, Panel B). Thus, this smoke-induced increase in out-migration represents a 1.3-percent increase relative to the sample-average out-migration rate. This result suggests a small—yet significant—subset of the population consistently travels away from their home counties when smoke plumes cover their homes. In the next section, we ask whether privilege-related demographics predict this behavior.

Column (1) of Panel B also documents statistically significant evidence of smoke-induced out-migration—specifically, the 75$^{\text{th}}$ percentile of a CBG's distance traveled significantly increases when CBGs face smoke. We estimate that the 75$^{\text{th}}$ percentile increases by 1.7 kilometers when smoke plumes intersect with the CBG. This 1.7-kilometer increase is relative to a sample average of 48.4 kilometers (Panel B of Table A1), implying a sizable (3.5-percent) increase in the 75$^{\text{th}}$ percentile of travel in weeks with smoke. Put differently: Wildfire smoke pulls out the right tail of the distance-traveled distribution for urban, west-coast CBGs.

Both panels of Table 1 offer statistically significant evidence that out-migration increases when communities face wildfire smoke. However, these estimators pool behavior across heterogeneous communities—potentially glossing over important differences in individuals' responses to wildfire smoke exposure. The following sections examine how out-migration behavior correlates with income, race, and ethnicity.

2.3.3 Income and Smoke Migration. We now turn to the results of income-based heterogeneity in smoke-induced out-migration. Column (2) of Table 1 repeats the regressions of the previous section but allows heterogeneity by CBGs' income. Specifically, we estimate Equation 2.1 with Percentile$_i$ defined as CBG i's percentile (between 0 and 1) in the West Coast's median-household income distribution. For example, a CBG with median household income of \$74,000 is in the 50$^{\text{th}}$ percentile and would have Percentile$_i = 0.50$.

Both outcomes (panels) in Table 1 reveal sizable and statistically significant relationships between communities' smoke-based out-migration behavior and income.

The direction of this heterogeneity follows a pattern similar to many relationships documented in the environmental- and social-justice literatures: more privileged communities display heightened avoidance behavior—here, out-migration—in the presence of wildfire smoke. In fact, for the lowest-income communities, there is no statistically significant evidence of out-migration. If anything, these communities reduce out-migration: the point estimates for percent out-of-county visits and distance traveled are both negative but do not differ significantly from zero. The level of out-migration only differs significantly from zero for communities above the 47$^{\text{th}}$ percentile of income for out-of-county travel (panel A) and the 49$^{\text{th}}$ percentile for distance traveled (panel B).

More affluent communities out-migrate substantially—and significantly—more than lower-income communities. The interaction coefficients in Column (2) indicate that smoke increases the share of other-county visits for the top percentile by 1.4 percentage points and increases their 75$^{\text{th}}$ percentile of travel by 17.2 kilometers. Notably, the effects for the most affluent communities are 3–6 times larger than the pooled effects presented in Column (1). While there is little evidence that low-income communities travel to avoid smoke, Table 1 provides clear evidence that affluent communities exercise this strategy.

The empirical specification above imposes linearity in the relationship between CBGs' income percentiles and heterogeneous smoke-induced out-migration: an increase in one percentile increases out-migration by $\delta/100$. We relax this restriction by specifying Percentile$_i$ as 50 mutually exclusive indicator variables. These indicators group neighboring percentiles together—e.g.the first and second percentile are in the same indicator bin. We also drop the main effect (Smoke$_{iw}$ in Equation 2.1)—rather than dropping one of the individual indicator variables—so that we can directly compare percentiles' tendencies to out-migrate.

Panel A of Figure 2 illustrates the results of this semi-parametric specification for income-based heterogeneity. Subfigure i shows the point estimates and their 95% confidence intervals for each of the 50 income-percentile bins' tendencies to out-migrate in response to smoke. Subfigure ii depicts bins' general tendencies to travel beyond their home counties (regardless of smoke). Finally, Subfigure iii maps each bins' median income.[21]

The results from this semi-parametric specification largely mirror those of the simpler regression in Table 1: A community's degree of smoke-based out-migration correlates strongly with the community's median income. Communities below the 50$^{\text{th}}$ percentile do not, on average, significantly out-migrate when facing smoke. While the increase in out-migration appears quite linear in communities' income percentile (the x axis), the tendency to out-migrate appears to increase sharply above the 90$^{\text{th}}$ percentile—approximately the same point at which the income distribution sharply increases. Finally, the point estimates of this semi-parametric estimation

[21] Appendix Figure B2A reproduces Figure 2A with distance traveled as the outcome.

(depicted Figure 2Ai) suggest an even higher rate of smoke-induced out-migration (~1.4 percent) relative to the results from the linear-specification results in Table 1.

Subfigure Aii highlights notable differences in other-county travel throughout the income distribution. The same communities that are more likely to out-migrate in the presence of smoke are already traveling more. These insights are also corroborated by our distance-traveled measure (see Appendix Figure B2A).

Both specifications of income-based heterogeneity—and both measures of out-migration—produce the same conclusion: In the presence of wildfire smoke, wealthier communities are significantly more likely to out-migrate than poorer communities. There is no significant evidence that smoke induces any out-migration in communities below the 30^{th} percentile of income.

2.3.4 Race, Ethnicity, and Smoke Migration. Disparities in smoke-induced out-migration extend to race and ethnicity.

Columns (3–5) of Table 1 estimate heterogeneity in smoke-induced out-migration as a function of CBGs' racial- or ethnic-composition percentiles. Similar to our specification of income-based heterogeneity, we estimate Equation 2.1 with Percentile_i defined as CBG i's percentile (between 0 and 1) in the West Coast's distribution of the share of the population that is Black (Column 3), Hispanic (Column 4), or White (Column 5).

Table 1 provides more evidence that historical privilege[22] correlates with communities' tendencies to out-migrate when facing wildfire smoke. Column (3) estimates that smoke significantly increases the share of out-of-county travel in the West Coast's least-Black communities by 0.54 percentage points—and increases their distance traveled by 6 kilometers. However, the most-Black communities show no significant evidence of smoke-induced out-migration: neither out-of-county travel nor distance traveled. Communities that are at least five percent Black (above 65^{th} percentile in the West Coast's distribution) show no significant evidence of smoke-base out-migration (i.e.their confidence intervals include zero).

The story is similar for Hispanic communities: Column (4) documents that the least-Hispanic communities significantly respond to smoke in their share of out-of-county trips (increasing by 0.89 percentage points) and in the distance traveled (increasing the 75^{th} percentile by 11.5 kilometers). As we found with the most-Black communities, the West Coast's most-Hispanic communities show no significant evidence of smoke-induced out-migration. Communities

[22] *I.e.*, populations that are less Black, less Hispanic, and/or more White.

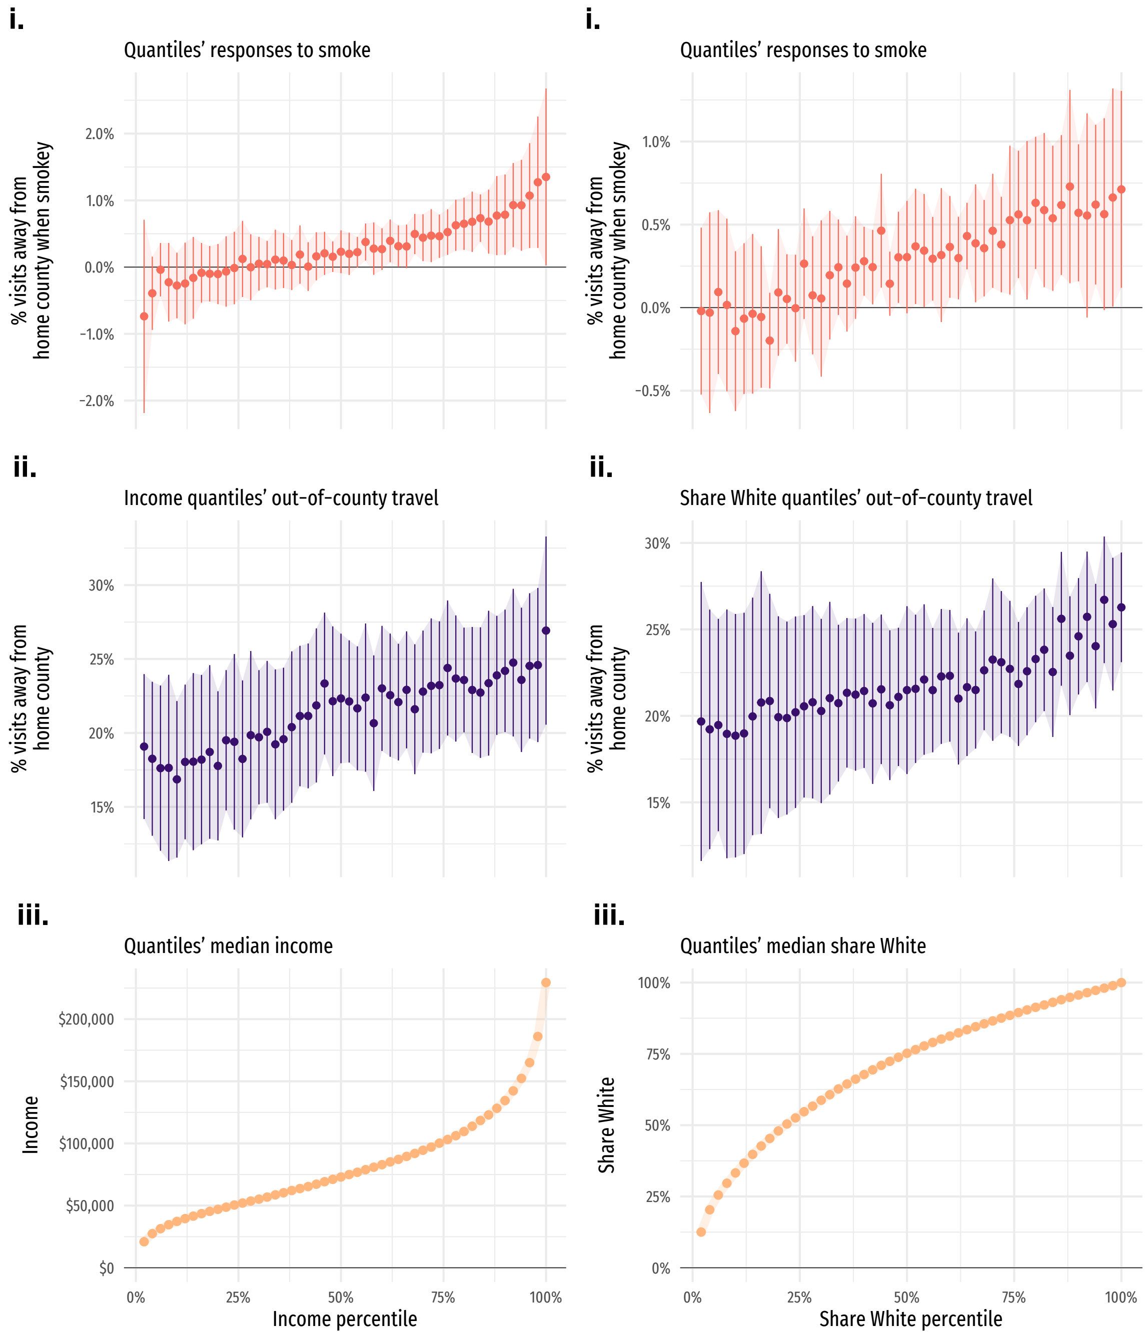

(a) **Income quantiles** (b) **Percent White quantiles**

Figure 2. Inequality in Smoke-induced Out-migration: Income and Percent White

whose population is at least 24-percent Hispanic (the 54$^{\text{th}}$ percentile) show no significant evidence of smoke-based out-migration.

Column (5) of Table 1 also bears evidence of disparities in out-migration that correlate with historical privilege: communities with larger shares of White individuals out-migrate more than less-White communities. The least-White communities (less than 25% White) show no statistically significant evidence of out-migration in terms of out-of-county travel or distance traveled—both estimates are negative and do not differ significantly from zero. The most-White urban communities on the West Coast (∼100% White) out-migrate significantly more than less-White communities. Communities above the 41$^{\text{st}}$ percentile of share White (communities that are at least 69% White) all show statistically significant evidence of smoke-induced out-migration.

As with income-based heterogeneity, we estimate a less-parametric model of heterogeneous out-migration for CBGs' percentiles of their population shares of Black, Hispanic, or White individuals. Panel 3A and Panel 3B display the results of these semi-parametric specifications for share Black and share Hispanic. Panel 2B illustrates the results for share White.[23]

The semi-parametric specification yields similar conclusions to the linear specification: out-migration behavior strongly correlates with less-Black, less-Hispanic, and more-White populations. At approximately the same point in the distribution that communities become majority White (the 25$^{\text{th}}$ percentile, see Panel 2B), they also significantly out-migrate when faced with smoke. As these majority-White CBGs increase in their percentage of White inhabitants, their smoke-induced out-migration continues to increase. Conversely, when the share of Hispanic individuals in a CBG crosses fifty percent (i.e.majority-Hispanic, ∼75$^{\text{th}}$ percentile, see Panel 3B), out-migration is approximately zero—and drops below zero as the share Hispanic increases.

Very few CBGs on the West Coast have a majority-Black population—80 percent of urban west-coast CBGs are less than 10 percent Black. Even with this low representation, Panel 3A still illustrates a clear trend: communities with larger Black-population shares out-migrate significantly less when facing smoke, relative to less-Black CBGs.

In addition to the significant disparities in smoke-induced out-migration that we discuss above, the middle subfigures (labeled *ii*) of Figures 2 and 3 depict striking differences in general (non-smoke-related) travel patterns. As communities become more Black, more Hispanic, less White, or less affluent, they travel out-of-county substantially less. The trend in communities'

[23] Appendix Figures B3A, B3B, and B2B repeat these figures for the 75$^{\text{th}}$ percentile of distance traveled.

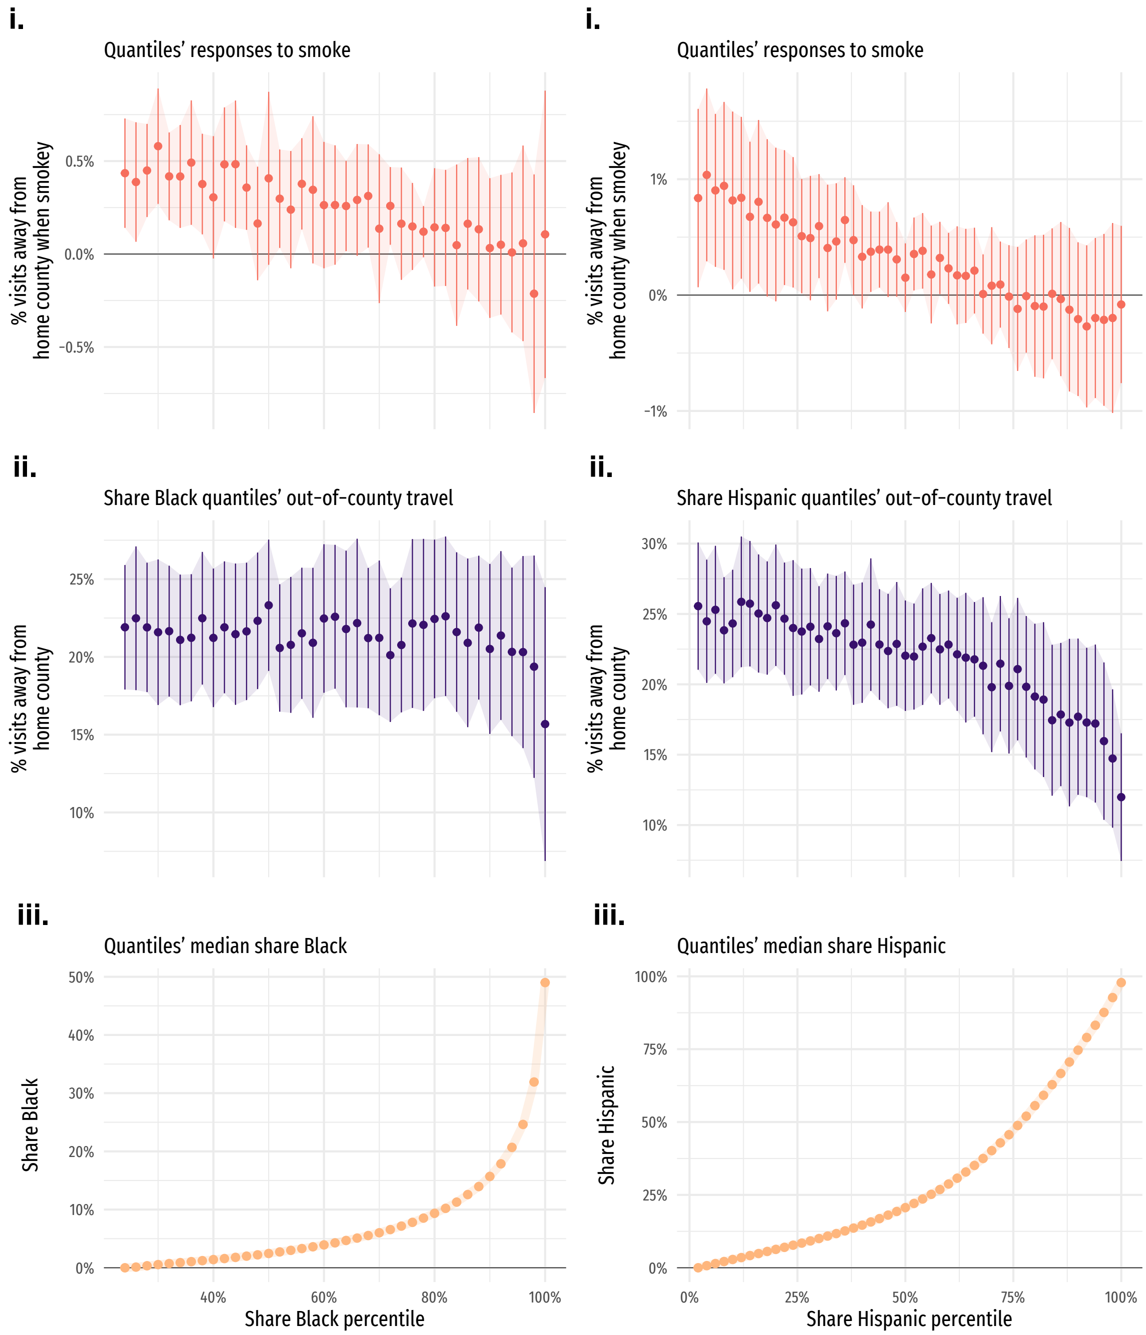

(a) **Percent Black quantiles** (b) **Percent Hispanic quantiles**

Figure 3. Inequality in Smoke-induced Out-migration: Percent Black and Percent Hispanic

Hispanic population is particularly evident. Twenty-five percent of POI visits in the least-Hispanic communities occur outside individuals' home counties; this number drops below twelve percent for the most-Hispanic communities.

2.3.5 Interpreting Out-migration Behavior. Above, we observe evidence that the share of non-home-county visits increases when affluent and historically advantaged communities face wildfire smoke. This outcome—the percentage of visits beyond residents' home counties—is the ratio of (a) the number of POI visits to other counties to (b) the total number of POI visits. Thus, an increase in the denominator (the total number of visits) could cause this ratio to increase—even when the number of other-county visits did not increase. However, an event study for the number of other-county visits (Appendix Figure B1C) suggests that the numerator—the *number* of other-county visits—increases due to smoke exposure. Further, when we estimate percentile-level effects of smoke exposure on the number of visits to non-home counties (rather than the share), the point estimates are positive for communities above the median income (Figure B4). Smoke exposure increases the *level* of out-migration in affluent communities. Consequently, even if the total number of trips (the denominator) declines,[24] the level of out-migration appears to increase in above-median-income communities in response to smoke. The results in the previous subsections can be interpreted as describing the effects of wildfire smoke on communities' levels and shares of out-migration.

2.4 Discussion and Conclusion

Notably, the dimensions along which we document heterogeneous out-migration are descriptive. Our empirical strategy does not provide identifying variation in income, race, or ethnicity. However, showing inequality in avoidance correlates with contemporary and historical privilege is critical for social equity, public policy, and future research.

Figure 4 documents how one of our dimensions—income percentile—correlates with nine socioeconomic variables from the ACS (for urban, west-coast CBGs). Solid dots denote variables' medians for the given income-percentile bin. Darker shading denotes bins' interquartile range (25^{th}–75^{th} percentiles); lighter shading shows bins' 10^{th}–90^{th} percentiles. Across many different dimensions, Figure 4 reinforces how strongly income correlates with many variables important for equity, policy, and understanding potential mechanisms.

Regarding social equity: Our results demonstrate yet another dimension multiplying inequity in disadvantaged communities. Communities less likely to travel away from smoke

[24] Burke et al. (2022)'s find that smoke induces more households to remain at home. Figure B4 suggests a small reduction in total visits.

are also more likely to inhabit homes easily penetrated by smoke and pollution (e.g.rentals and mobile homes), live in polluted areas, face significant health issues, and lack health insurance. In other words, communities that are more likely to stay home amidst smoke are likely facing more smoke *inside* their homes and starting with worse health. Together, these factors may explain why poorer and historically marginalized communities appear to be more susceptible to smoke exposure.[25] This situation compounds inequality.

Compounding disadvantages may suggest high-return areas for policy. For example, policies that improve houses' *seals* (i.e.weatherization and energy efficiency programs that prevent outside smoke and pollution from entering the home) may be especially beneficial in communities less able/likely to travel away from smoke. The environmental-justice literature demonstrates that these households are also more likely to inhabit more polluted areas (Colmer et al., 2020) and face serious health challenges (e.g.higher rates of disability, as in Figure 4). Beyond avoiding smoke exposure, improved seals could also reduce the burden caused by unequal outdoor-pollution exposure. Such improvements would likely improve these homes' energy efficiency, reducing these households' utility bills. Additional research can better direct such policies.

Finally, our results highlight fruitful topics for future research—especially in understanding the mechanisms that cause unequal levels of out-migration across income, race, and ethnicity. Wealth is one obvious possible mechanism. Differing job types (e.g.hourly vs. salaried) or job benefits (income and college degrees correlate strongly in Figure 4) may also explain unequal avoidance; unplanned travel requires a degree of professional flexibility. Access to transportation offers another possible mechanism: Figure 4 shows that low-income households are substantially less likely to have access to a vehicle. Information may also play a role—particularly for subscription-based sources like internet and newspapers. Of course, there are many other potential mechanisms—e.g.liquidity or credit access. Future research can pursue these paths.

While populations often face common hazards, individuals do not equally employ avoidance strategies. We find robust and significant evidence that *some* communities out-migrate in the presence of smoke; other communities show little evidence of out-migration. Our results show avoidance strongly correlates with income, race, and ethnicity—salient dimensions for contemporary and historical disadvantage. Without additional policies or interventions, unequal access to avoidance may exacerbate existing inequality—a potentially important insight as humanity faces a rapidly changing climate.

[25] For examples, see Borgschulte et al. (2022); Miller et al. (2017).

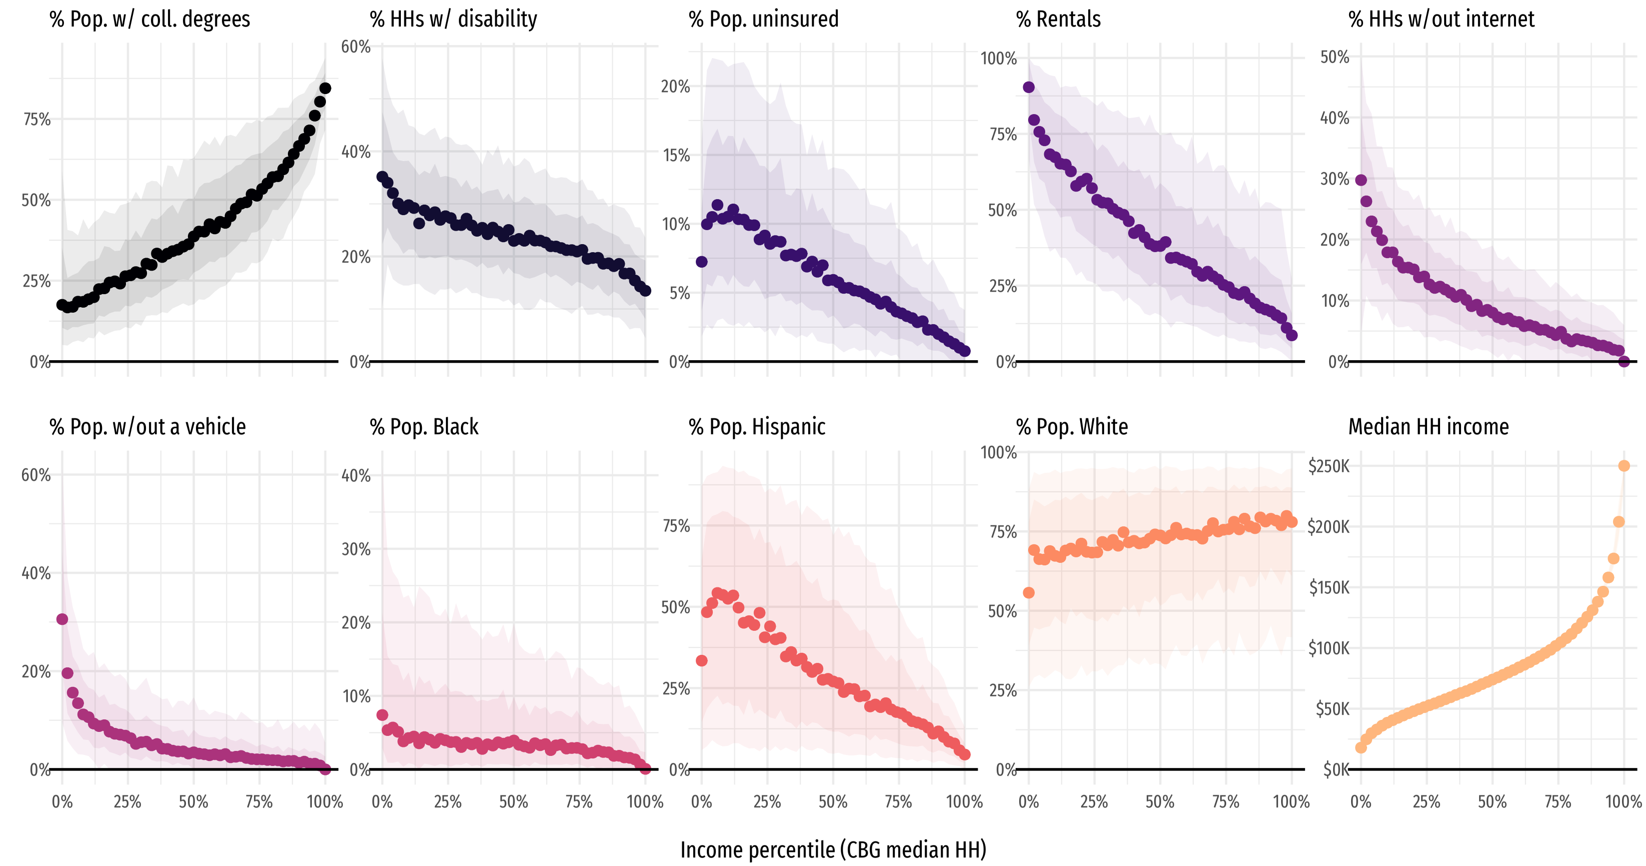

Figure 4. Urban CBG Income Percentiles and Other Socioeconomic Dimensions
Each subfigure summarizes urban, west-coast CBGs levels of the given variable by two-percentile bins (e.g.combining the first and second percentiles). Solid dots the median. The first, darker band gives the interquartile range (25^{th}–75^{th} percentiles) for the given bin. The lighter band marks the 10^{th}–90^{th} percentiles. Note that y-axis scales change across subfigures to match variables' differing variation. Color denotes different variables.

CHAPTER III

TIME-VARYING TIME PREFERENCES

This chapter is co-authored with Connor Thomas Wiegand, doctoral candidate at the University of Oregon in the Department of Economics, and Associate Professor Michael A. Kuhn. I developed the intuition of the theory behind both time-varying and stationary discount functions, as well as developed the first form of the nested exponential function presented. I also created the surveys used and oversaw the experiment week to week, generated results, figures, and tables in the empirical sections and wrote the first draft of the paper. Connor contributed by formalizing proofs, generating the tables and figures, and collaborating in writing the theory section. He also contributed in design of the elicitation device used in the surveys and overall experimental design. Professor Kuhn provided theoretical, empirical, and experimental guidance, funded and acted as administrator for the experiment, and edited and revised after the first draft. IRB approval for this research was obtained from the University of Oregon, Study ID: STUDY00000768

3.1 Introduction

A substantial body of evidence documents that economic preferences can be unstable.[1] Beyond shocks to preferences, decision-makers also learn about their preferences (either through feedback (Charness, Chemaya, & Trujano-Ochoa, 2023) or the realization of uncertainty (Carrera, Royer, Stehr, Sydnor, & Taubinsky, 2022) or "choose" preferences to better suit their environment (Bernheim, Braghieri, Martínez-Marquina, & Zuckerman, 2021). Preference instability, shocks, and drift are especially problematic for researchers looking to estimate time preferences. Consider the framework of Read and van Leeuwen (1998): an individual selects a healthy snack for their future self, but revises their choice to an unhealthy snack when the time for consumption arrives. Is this preference reversal best explained by a fixed feature of that individual's discount function (e.g. hyperbolic discounting), learning about preferences, or some other unobserved environmental factor?

[1]For example, traumatic events like the Great Depression (Malmendier & Nagel, 2011), exposure to violence (Callen, Isaqzadeh, Long, & Sprenger, 2014; Voors et al., 2012), and natural disasters (Bchir & Willinger, 2013; Beine, Charness, Dupuy, & Joxhe, 2020; Bourdeau-Brien & Kryzanowski, 2020; Cameron & Shah, 2013; Cassar, Healy, & Kessler, 2017; Eckel, El-Gamal, & Wilson, 2009; Hanaoka, Shigeoka, & Watanabe, 2018; Kettlewell et al., 2018; Kibris & Uler, 2021; Kuroishi & Sawada, 2019; Page, Savage, & Torgler, 2014; Reynaud & Aubert, 2019), short-term influences like happiness (Ifcher & Zarghamee, 2011), loss of control (A. Gneezy & Imas, 2014), and hunger (Ashton, 2015; Kuhn, Kuhn, & Villeval, 2017), as well as environmental factors like rainfall (Jaramillo, LaFave, & Novak, 2023) and political instability (Sepahvand, Shahbazian, & Swain, 2019) all have demonstrated causal impacts on risk preferences.

Halevy (2015) defines three properties of time preferences: time-consistency, stationarity, and time-invariance. Imagine a decision-maker evaluating a paycheck advance –which results in some distant future disutility– to purchase tickets to a concert. All-else equal, his time preferences violate *stationarity* if he takes an advance **today** to see a concert **tonight**, but does not take an advance **today** to see an equivalent concert **tomorrow**. His time preferences violate *time-consistency* if after deciding **yesterday** not to take an advance for **tonight's** concert, he wakes up **today**, changes his mind and takes the advance to see the concert **tonight**. His time preferences violate *time-invariance* if after deciding **yesterday** not to take an advance for **last night's** concert, he wakes up **today** and decides to take an advance for **tonight's** concert. Any two of these properties imply the third; if the behavioral economist theorizes that the decision-maker is not time-consistent, she assumes that the decision-maker is either not time-stationary or not time-invariant (or neither). Common models of discounting in behavioral economics all pair inconsistency with non-stationarity, while leaving time-invariance intact. However, data shows this is only one of many patterns that people exhibit. Table 2 presents the classification of subjects based on experimental data from Halevy (2015). At most 68%, and as little as 44% of subjects exhibited either all three properties or only time-invariance[2]. Thus, between 32% and 56% of the sample exhibited preferences that cannot be explained with the standard modelling toolkit. In another closely-related paper, Janssens et al. (2017) finds that while 43% of subjects in their study violate time-consistency, only 24% violate *both* time-consistency and stationarity. Similarly, Harrison, Lau, and Yoo (2023) finds that time variance explains why such a large fraction of their experimental sample is time-inconsistent even though they display stationary behavior in their first of their two surveys. Time-variance is an important reason why subjects in all three studies exhibit inconsistent behavior.

Our goal in this paper is to provide a comprehensive treatment of time-variance in estimating time preferences. This entails a theoretical and structural discussion of how to model and allow for time invariance, and a 12-week longitudinal experiment that constructs a panel of intertemporal choices across seven elicitations. While the two-period studies in Halevy (2015), Janssens et al. (2017), and Harrison et al. (2023) allow for the impact of time-variance on experimental measures, they do not give the researcher much power to think about the source of time-variance and model it: is it noise, learning, a result of macroeconomic changes, a result of individual-level shocks to wealth, liquidity, stress, or something else entirely? Answering this question determines the scope of the time-variance problem. If time-variance is idiosyncratic noise,

[2] Janssens, Kramer, and Swart (2017) find a similar composition using a convex time budget elicitation.

Table 2. Time Preference Property Classification From Halevy (2015)

Share of sample:	Min. across approaches	Max. across approaches
Properties of choices	(1)	(2)
A) Invariant, consistent, stationary	35.04%	59.32%
B) Invariant only	6.83%	11.86%
A) + B)	43.58%	67.79%
C) Stationary only	16.95%	27.35%
D) Consistent only	5.08%	16.24%
E) None	4.27%	27.12%
C) + D) + E)	32.21%	56.42%

Notes: Data are from Table 2 in Halevy (2015). There are a total of eight classifications presented in the table, representing the 2x2x2 treatment variation of stake size, strict vs. weak preferences, and whether subjects had full knowledge of the future aspects of the study when they made their initial choices.

then what appears to be time-variance or inconsistency may instead arise from decision error. Failing to account for this would lead to individuals being incorrectly classified as certain types of discounters, but would not impact aggregate estimates of how a population discounts (besides decreasing precision). On the other hand, if time-variance is a systematic process wherein subjects learn about their preferences within the context of the novel experimental setting they encounter, or their time preferences react to correlated environmental factors, the problem is more pernicious.

What is the risk associated with mistaking one kind of time-inconsistency for the other? The behavioral literature on "present-bias" and associated policy interventions (e.g. commitment devices) are predicated on time-inconsistency as a pernicious effect of the opportunity for immediate gratification. If choices can be structured to avoid that opportunity, then the decision-maker can implement their stable long-run preference for saving more, eating healthier, etc. Time-inconsistent choices associated with time variance do not have the same interpretation. Consider a decision-maker that exhibits time-varying preferences such that their intertemporal allocations between rewards on two relative dates (e.g. the day of the choice and eight weeks after the choice) are getting less patient as time passes (meaning that the date of the choice and both reward dates move together in time). First, assume her preferences are time-stationary. If a researcher estimates her quasi-hyperbolic (β-δ) discount function via an observed violation of time-consistency, the researcher will conclude that she is present-biased ($\beta < 1$). Instead, assume her preferences are time-consistent. If her quasi-hyperbolic discount function is estimated via an observed violation of stationarity, the researchers will conclude that she is future-biased ($\beta > 1$).[3].

[3]We show this more formally in Section 3.2.1

The first part of the paper is about why time variance is a problem for identification in time-invariant models, and then introduces models of intertemporal choice that allow time-varying preferences. We present a new model – the "nested exponential" discount function– that is fully general with respect to the properties defined by Halevy (2015). The nested exponential form can satisfy any of the five possible combinations depending on its parameter values: all three properties can fail, all three can hold, or any single property can hold on its own. While this four-parameter model is (of course) not as tractable as the workhorse β-δ model, we develop the intuition of the model as a discrete extension of fast Weibull discounting (Jamison & Jamison, 2011), one of the less-common behavioral models that permits time-inconsistency and non-stationarity. Thus, the model retains a "present-bias" parameter similar to these models.

The second part of the paper presents our longitudinal experiment. Over the course of 12 weeks subjects take part in seven surveys, each featuring 6-10 incentivized time preference measurements. This design gives us six opportunities to observe time-inconsistency and time variance (for each measurement), and seven opportunities to observe non-stationarity. We first look at our experimental data through the lens of traditional time-invariant discounting, and estimate parameters for the quasi-hyperbolic β-δ discount function. We identify small but statistically significant present bias at the aggregate level, and the majority of subjects are classified as present-biased at the individual level. We then shift to analyses that allow for time-varying preferences. We identify a significant time trend in the aggregate discount factor consistent with *decreasing patience* over the period of the study. This trend stems from subject-level changes across time, not sample composition, suggesting systematic rather than idiosyncratic time variance. With time-variance accounted for, we can measure distinct inconsistent and non-stationary behavior at the subject-level; we find that inconsistency is more common and substantial than non-stationarity. The prevalence of individual violations of all three time preference properties leads us to classify subjects as time-variant at the upper end of the range Halevy (2015) suggests. Overall, time-variance explains 72% of the time-inconsistency in our sample.

Finally, we structurally estimate the nested exponential discount function at both the aggregate and individual levels. The nested exponential model fits the data best, but adjusted for its extra parameters, it only slightly outperforms a version of the model constrained to be time-invariant only. In other words, classic "present-bias" where time-inconsistency and non-stationarity always go together does a decent job at explaining aggregate behavior *so long as researcher uses a discount function with some true hyperbolicity*, rather than a discrete approximation. This model vastly outperforms all other two-parameter models, including the quasi-hyperbolic β-δ model.[4] At the individual level, we find that the full nested exponential model and the time-invariant only restricted model both best-describe about 40% of the sample. The remaining 20% fall into roughly equal-sized groups of exponential, stationary-only, and time-consistent-only discounters. Overall, we assign 54% of the sample to time-varying models of discounting, suggesting that understanding, measuring, and modelling preference shifts over time is crucial for accurately estimating time preferences.

3.2 Theory

We adopt notation from Halevy (2015) to describe the stationarity, time-consistency, and time-invariance properties of time preferences. Consider an agent in calendar time $\tau \geq 0$ evaluating bundles (ε, t) and (ψ, t'), which deliver rewards ε, ψ at time periods t, $t' \geq \tau$. We assume that the agent has complete and transitive preferences over all such bundles, according to a preference ordering that may depend on τ, $\succsim_\tau$.

Under classical consumer theory assumptions, including additive separability of utility across time periods, this preference relationship is represented by a *discount function*, $D(\tau, t)$, such that

$$(\varepsilon, t) \succsim_\tau (\psi, t') \iff D(\tau, t) \cdot \varepsilon \;\geq\; D(\tau, t') \cdot \psi \,.^{5} \tag{3.1}$$

$D(\tau, t)$ describes the agent's discount "factor" when evaluating, in period τ, a payoff in period t.

Definition 3.2.1. An agent's preferences satisfy **stationarity** if

$$D(\tau, \tau + \Delta_1)\, \varepsilon \;\geq\; D(\tau, \tau + \Delta_2)\, \psi \iff D(\tau, \tau' + \Delta_1)\, \varepsilon \;\geq\; D(\tau, \tau' + \Delta_2)\, \psi$$

An agent whose preferences are *stationary* will not exhibit preference reversals when the payoff time of all bundles is shifted by the same number of periods, all else equal.

Definition 3.2.2. An agent's preferences satisfy **time consistency** if

$$D(\tau, \tau + \Delta_1)\, \varepsilon \;\geq\; D(\tau, \tau + \Delta_2)\, \psi \iff D(\tau', \tau + \Delta_1)\, \varepsilon \;\geq\; D(\tau', \tau + \Delta_2)\, \psi$$

[4] Hyperbolicity is a key feature of our data; among one-parameter models, the true hyperbolic discount function outperforms exponential discounting.

[5] Without loss of generality, we replace $u(\varepsilon)$ and $u(\psi)$ with ε and ψ, under the assumption an instantaneous utility function representation exists.

Agents whose preferences are *time-consistent* will not exhibit preference reversals when the same rewards, in the same payoff periods, are evaluated in different time periods.

Definition 3.2.3. An agent's preferences satisfy **time invariance** if

$$D(\tau,\, \tau + \Delta_1)\, \varepsilon \;\geq\; D(\tau,\, \tau + \Delta_2)\, \psi \iff D(\tau',\, \tau' + \Delta_1)\, \varepsilon \;\geq\; D(\tau',\, \tau' + \Delta_2)\, \psi$$

Agents whose preferences are *time-invariant* will not exhibit preference reversals when the evaluation period and payoff periods for all bundles are shifted by the same amount. We use TICS (time-invariance, consistency, stationarity) to refer joints to these three properties of time preferences.

Halevy (2015) proves the following proposition, linking these three properties together:

Proposition I. *The TICS properties form a binary. That is, if a preference relation satisfies any two of the properties, then it must imply the third.*

Thus, the TICS properties generate five groups that any agent's preferences may fall into: all, time-invariant only, stationary only, time-consistent only, none. The consequences of this binary relationship are not well-appreciated in the literature. Economists who wish to model time-inconsistent agents, for instance, must also depart from at least one of stationarity or time-invariance. The latter approach is typical; the hyperbolic and β-δ discount functions satisfy only time-invariance. There exists only a nascent literature on discount functions that do not satisfy time-invariance (Strulik, 2021), and on observed time-variant behavior (Brownback, Imas, and Kuhn (2023); DeJarnette (2020); Halevy (2015); Harrison et al. (2023); Imas, Kuhn, and Mironova (2022); Janssens et al. (2017)). Our goal is to both offer new structural discount function that is fully general to the TICS properties and evaluate its time-variance properties in a longitudinal study.

Before proceeding, there are two terms that we will define for use in this paper that are convenient short-hands for phenomena we observe, but come with a lot of baggage in the literature: *(im)patience* and *present/future-bias.* In some papers these terms are closely linked (e.g. Prelec (2004), Attema, Bleichrodt, Rohde, and Wakker (2010), or Takeuchi (2011), where present-bias is considered decreasing impatience). However, it will be useful for us to keep them separate in order to use terms like increasing and decreasing patience to describe time variance.

Specifically, a subject in our study exhibits decreasing patience if their discount factor for some reward delayed by $t - \tau$ periods is decreasing in τ: $D(\tau, \tau + \Delta) > D(\tau', \tau' + \Delta)$ for $\tau' > \tau$, $\Delta > 0$.

Present-bias is a term that either casually or formally refers to the phenomenon wherein an agent's relative preference for sooner rewards over later rewards is greatest when those rewards are immediately available (in the evaluation period). This is usually modeled discontinuously by economists, with extra discounting applied to all non-immediate rewards using the workhorse quasi-hyperbolic (or β-δ) discounting model (Laibson, 1997; O'Donoghue & Rabin, 1999), but is a continuous feature of the hyperbolic discounting model widely used outside of economics. Beyond ambiguity over whether the term applies only to discontinuous immediate gratification models, there is also a recent debate over whether the term "bias" is appropriate to describe an empirical phenomena that could stem from a variety of root causes (Bernheim & Taubinsky, 2018). In this paper we will use "present-bias" and "future-bias" to empirically classify non-stationary and time-inconsistent behavior. If shifting the payoff time of each bundle within a binary choice further into the future decreases the relative preference for the sooner bundle, we call this present-bias. Similarly, if shifting the evaluation time of a binary choice closer to its fixed payoff time periods increases the relative preference for the sooner bundle, we call this present-bias as well.

3.2.1 Measuring Discounting Properties. Consider a decision maker (DM) faced with a choice between an immediate reward at $t = \tau = 0$ and a reward on some future date $t = t_1 > 0$. Call a a scalar measure of how much their choice favors the reward at the later date, t_1, i.e. higher values of a result from more patient discounting of that delayed reward. In the same evaluation period, the DM also faces a choice between utility in two future time periods, $t_2 > 0$ and $t_2 + t_1$. Call b the corresponding measures of that choice. After time passes, such that $\tau = t_2$, the DM faces another similar choice: call c the measure of the choice between immediate utility at $t = \tau = t_2$ and future utility at $t_2 + t_1$. Figure 5 depicts the relationship between the three measurements. We call the collection of these three choices a "decision triangle." Using the values associated with any given triangle, $a - c$ is a measure of time-variance; how much a choice over fixed relative dates changes over time. $b - c$ measures inconsistency; how much a choice over fixed calendar dates changes over time. Finally, $b - a$ measures non-stationarity; how much a choice with a fixed delay changes in response to front-end delay. Combining the three calculations, we have that

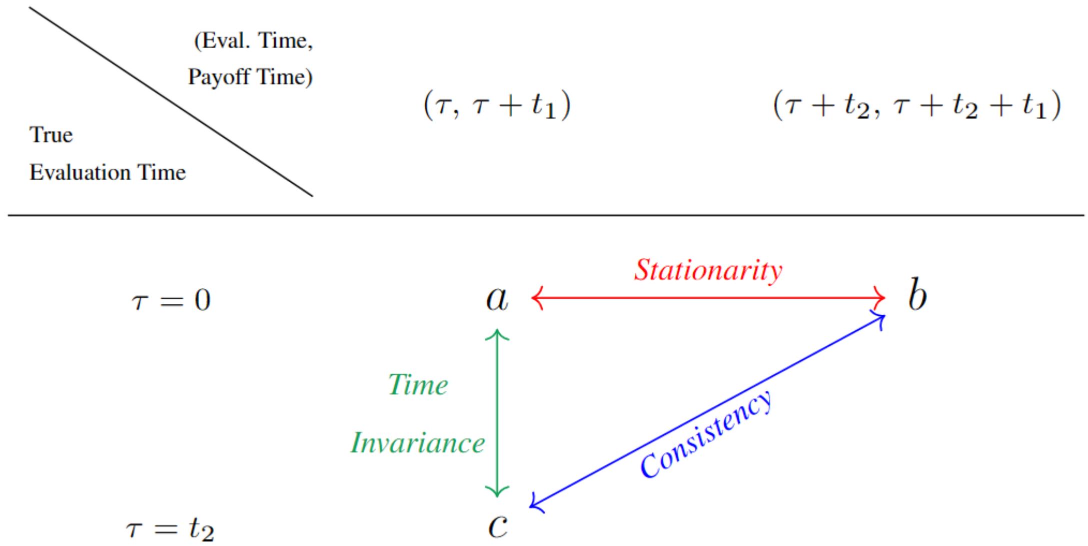

Figure 5. Time Preference Property Decision Triangle

$$\underbrace{a-c}_{\substack{\text{Time}\\\text{Variance}}} = \underbrace{a-b}_{\text{Non-Stationarity}} + \underbrace{b-c}_{\text{Inconsistency}}.$$

With this framework, we can illustrate the pernicious empirical issues that arise when a researcher is not accounting for time-varying preferences. Suppose a researcher wishes to estimate parameters for an agent using the β-δ model, and they set up an experiment to test the agent's time-consistency. Assume the agent exhibits time-varying preferences in the form of decreasing patience, leading to $a-c<0$. Assume also that the agent exhibits stationary preferences, so that $b-a=0$. If this is the case, then the agent must exhibit time-inconsistent preferences such that $b-c=a-c<0$. Thus, the researcher will conclude that the subject is present-biased and estimate $\beta<1$. Now assume instead that the same time-varying preferences are present, but the agent is time-consistent and the researcher sets up an experiment to test the agent's stationarity. Since $b-c=0$, it must be that the agent makes a non-stationary choice so that $a-b=a-c<0$. Thus, the researcher will find the agent is future-biased, and estimate $\beta>1$. Depending on research design, the same unobserved time-variance can lead to different incorrect conclusions about the agent's true behavioral response to an immediate gratification opportunity.

This is more complicated when considering an agent that exhibits violations of all three time preference properties. Consider a decision triangle, defined by the triple (a, b, c) of $(240, 235, 230)$. A stationarity violation exists, as $240 \neq 235$. The researcher would deem this choice future-biased, as an increase in payoff times for both bundles led the agent to be more impatient. A time-consistency violation exists, as $235 \neq 230$. The economist would deem this choice present-biased, as the passing of calendar time until the fixed payoffs led the agent to be more impatient. This confounds the estimation of β and may lead the economist to conclude that this agent is not present- or future-biased. While true on average, this is not a thorough explanation of this agent's preferences.

To account for time-varying preferences, the limited existing literature has used reduced-form approaches to correcting $b - c$ in light of $a \neq c$, or (less-likely given design constraints) correcting $a - b$ in light of $a \neq c$ (e.g. Janssens et al. (2017) or Harrison et al. (2023)). We offer an alternative: a time-varying structural model of intertemporal choice. The following section maps the five groups of time preference properties given by the TICS binary to structural discount functions and their properties.

3.2.2 Classes of Time Preferences. Time-invariant preferences are well represented in the literature. Discount functions that admit these preferences include the standard exponential, as well as the hyperbolic and β-δ models regularly used to model time-inconsistent decisions. These models all share a common property: They are a function of only the distance between payoff time and evaluation time, $t - \tau$. Since these models are only concerned with this distance, the $D(\tau, t)$ function is typically reduced to just $D(t)$, where τ is always 0, without loss of generality. This reliance on distance payoff time characterizes the necessary and sufficient condition for a function to admit time-invariance.

Proposition II. *A discount function admits time-invariant preferences if and only if* $D(\tau, t) = D(t - \tau)$.

See Appendix Section D.1 for the proof.

We are only aware of one discounting functional form in the literature that violates this property. Strulik (2021) introduced a time-consistent hyperbolic discount function which violates the assumption of time-invariance:

$$D(\tau, t) = \left(\frac{1 + \alpha\tau}{1 + \alpha t}\right)^{\beta} \qquad \alpha > 0,\ \beta > 1\ . \tag{3.2}$$

Strulik (2021) designs the function for use in environmental resource models, relying on the increasing patience of the model to avoid extinction steady states. Importantly, commitment to consumption plans in his context requires that the discount function maintains time-consistency. Therefore, this model must be only time-consistent, as it is not time-invariant and so cannot be both time-consistent and stationary. How do we ensure that a discount function maintains time-consistency? Strotz (1956) suggests that the only discount function that can achieve time-consistency is the standard exponential. However, this result relies on the assumption of a discount function that is a function of only $t - \tau$. Strotz's result has since been generalized to prove instead that time-consistency of a discount function is equivalent to multiplicative separability in t and τ (Burness, 1976; Drouhin, 2020).

Proposition III. *A discount function $D(\tau, t)$ admits consistent preferences if and only if it can be written:*

$$D(\tau, t) = f(t) \cdot g(\tau)$$

Further, if the discount function can be written in this way, then $g(x) \equiv \frac{1}{f(x)}$.

See Appendix Section D.1 for the proof.

Including the time-consistent hyperbolic model, three of the five possible time preference groups have been characterized by discount functions in the literature. To our knowledge, discount functions have not been developed to fit the two remaining groups: admitting only stationarity and admitting none of the TICS properties. In order to develop these functions in the next section, we present the necessary and sufficient conditions for a discount function to admit stationarity, which also, to our knowledge, has not been shown in the literature.

Proposition IV. *Let $\alpha \in \mathbb{R}^+$ and f be a C_1 function linear in $t - \tau$. A discount function admits stationary preferences if and only if it takes the form*

$$D(\tau, t) = \alpha^{-(t-\tau)f(\tau)} \ .$$

See Appendix Section D.2 for the proof.[6]

With propositions II, III, and IV, we arrive at the following well-known result:

[6]Note that f may be trivial in τ. If $\alpha > 1$, we require that $f(\tau) < 0$ (equivalent to dropping the minus sign in the exponent).

Corollary. *The discount function $D(\tau,\, t)$ admits preferences exhibiting stationarity, time-consistency, and time invariance if and only if it is the exponential discounting function. That is, $D(\tau,\, t)$ can be written as a function of $t - \tau$ alone,*

$$D(t) = \delta^{\beta(t-\tau)}$$

3.2.3 Nested Exponential Discount Function. With necessary and sufficient conditions for a discount function to satisfy any of the TICS properties, we now have the tools to develop a discount function that can fit all five possible classes of time preferences. Two goals motivate our approach to developing such a function. First, the discount function should be fully general in its admittance of time preference properties. For instance, the β-δ model acts as an extension of the standard exponential, allowing for time-inconsistent and non-stationary preferences by setting $\beta \neq 1$. By extending the standard exponential in a similar manner, we can achieve each of the five time preference classes according to its parameter values. Second, a discount function that allows for time-varying preferences should be capable of exhibiting both increasing or decreasing patience depending on parameter values. The time-consistent hyperbolic model only allows for increasing patience, which is appropriate contextually. Allowing for decreasing patience, as well, however, expands the pool of time preference profiles we can accurately model.

$$D(\tau,\, t) = \delta^{\frac{(t^{\mu} - \tau^{\mu})^{\eta}}{\gamma^{\tau}}} \qquad \eta > 0,\, \gamma > 0,\, \mu > 0$$

This function, which we call the *nested exponential discount function*, can satisfy any of the five possible classes of time preference properties. To see this, suppose $\eta = \gamma = \mu = 1$. Then the nested exponential is identical to the standard exponential discount function. Now suppose $\eta = \gamma = 1$ while $\mu \neq 1$. Then the function is no longer a function of only $t - \tau$ nor does its exponent satisfy the conditions required to maintain stationarity. However, the function does remain multiplicatively separable in t and τ, granting it time-consistency. The argument follows similarly for when $\eta = \mu = 1$ and $\gamma \neq 1$. The function is not a function of only $t - \tau$ and is no longer be multiplicatively separable in t and τ. Its exponent is however can be expressed as $-(t - \tau)f(\tau)$, granting stationarity.

When $\gamma = \mu = 1$, but $\eta \neq 1$, the nested exponential is a discrete-time reformulation of the Weibull discount function, a less common behavioral model that achieves the same violations of the TICS properties as the hyperbolic and β-δ models. We have that the function is not multiplicatively separable in t and τ, the exponent is not of the form $-(t - \tau)f(\tau)$, but the function is still only a function of $t - \tau$. For the final class, satisfying no TICS properties, since any parameter not equalling one results in the violation of two properties, whenever two or more parameters are not equal to one, none of the properties are satisfied. Therefore, all five groups of time preference properties can be modeled with this one function.

This nested exponential function also allows for increasing or decreasing patience depending on the parameter values of η, γ, and μ. For instance, values of $\gamma < 1$ correspond to decreasing patience[7]. As τ increases, the relative distance between t and τ is amplified by a denominator that is increasingly less than 1, in turn increasing the exponent, and therefore decreasing the discount factor. The opposite holds for values of $\gamma > 1$. Values of $\mu < 1$ correspond to increasing patience. For any fixed $t - \tau = C$, the corresponding $t^\mu - \tau^\mu$ is decreasing in τ when $\mu < 1$, resulting in a discount factor increasing in τ. The opposite again holds for $\mu > 1$. Since η on its own cannot cause increasing or decreasing patience, it only acts to mitigate or amplify the effects of the other parameters when active at the same time. The exact condition that dictates the direction of patience of the nested exponential discount function is given by:

$$D(\tau, \tau + \Delta) > D(\tau', \tau' + \Delta) \iff \gamma - \exp\left[\eta\mu\frac{(\tau + C)^{\mu-1} - (\tau)^{\mu-1}}{(\tau + C)^{\mu} - (\tau)^{\mu}}\right] > 0 \tag{3.3}$$

Figures 6 and 7 demonstrate the effect each parameter has on the discount function in comparison to the standard exponential form (note that we have selected the δ parameter to force each function in Figure 6 to be equal at a common point, and that the scales of the vertical axes differ across figures for visual clarity). In Figure 6, the effect of η heavily influences the curvature of the discount function. Letting $\mu = 1 = \gamma$, when $\eta \neq 1$ the function exhibits present- or future-bias in the same way other invariant-only functions do. $\eta < 1$ results in present-bias, where the discount function decreases at a faster rate than the standard exponential for periods close to evaluation period, and then eventually at a slower rate.

The effect of μ can be intuited from the example of competing measurements of present- and future-bias presented in Section 3.2.1. Start by considering time preferences that are only time-consistent. If decreasing patience is present, it will result in future-bias when measuring discounting within evaluation period (a stationarity violation). Indeed, we see that the function where $\mu > 1$ results in a higher discount factor (relative to $\mu \leq 1$) for any delay length due to

[7] Assuming $\eta = \mu = 1$

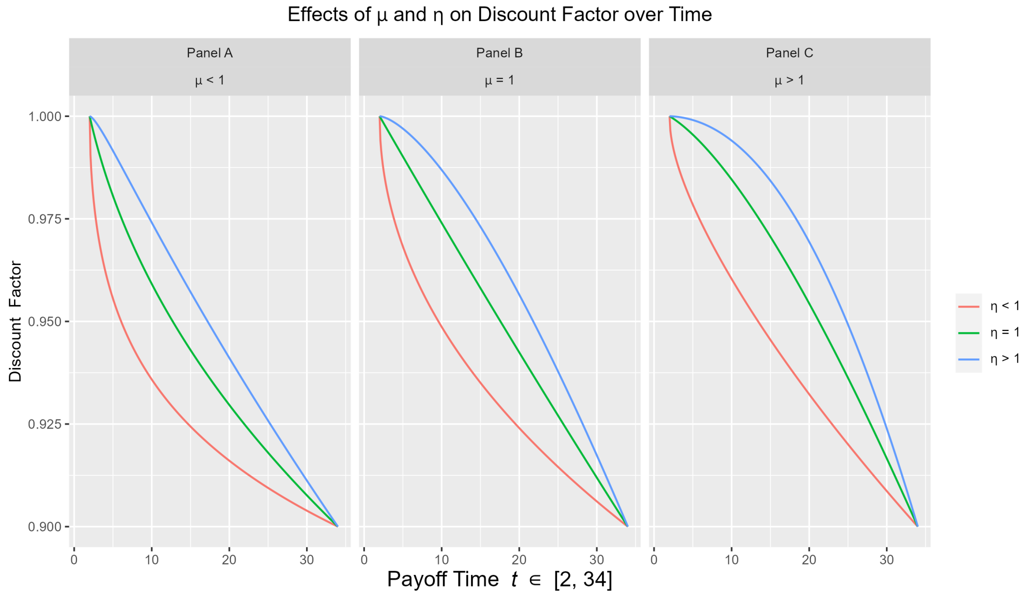

Figure 6. Nested Exponential Discount Function by μ and η Parameter Values

Notes: the value of δ in each function is chosen such that $D(2, 34) = 0.9$. The evaluation time τ is two in each function, and the x-axis shows payoff dates, t, from 2 to 34. We set $\gamma = 1$ for all functions.

a smaller rate of change in response to an increase in delay length. The intensity of this effect is partially driven by evaluation period.

The effect of γ can be described using the same example. If time preferences are stationary, then changing patience (stemming from $\gamma \neq 1$) will result in time-inconsistency. This time-inconsistency occurs because $\gamma \neq 1$ results in higher or lower discount factors for the same delay length across evaluation period. In particular, when $\gamma < 1$, the curve is distorted downwards, so that discount factors are lower for all delay lengths greater than zero as τ increases. This results in a present-biased decision, as the rate of change of the discount rate is not only necessarily more rapid on the curve for the new evaluation period, but the decision is closer to evaluation period, where the rate of change is quickest.

When it comes to estimating time-invariance and the parameters of nested exponential discount function, as opposed to a reduced-form invariance correction, the longitudinal dimension of the data determines the precision and reliability of those estimates. Existing work by Halevy (2015) and Janssens et al. (2017) collects data on a single decision triangle (ignoring different prices or choice delays within a triangle).[8] We observe six triangles per subject over the course of a 12-week study. Time-variance need not be idiosyncratic subject-level noise at different dates;

[8]Harrison et al. (2023) collects longitudinal data and uses structural assumptions to analyze them like a triangle.

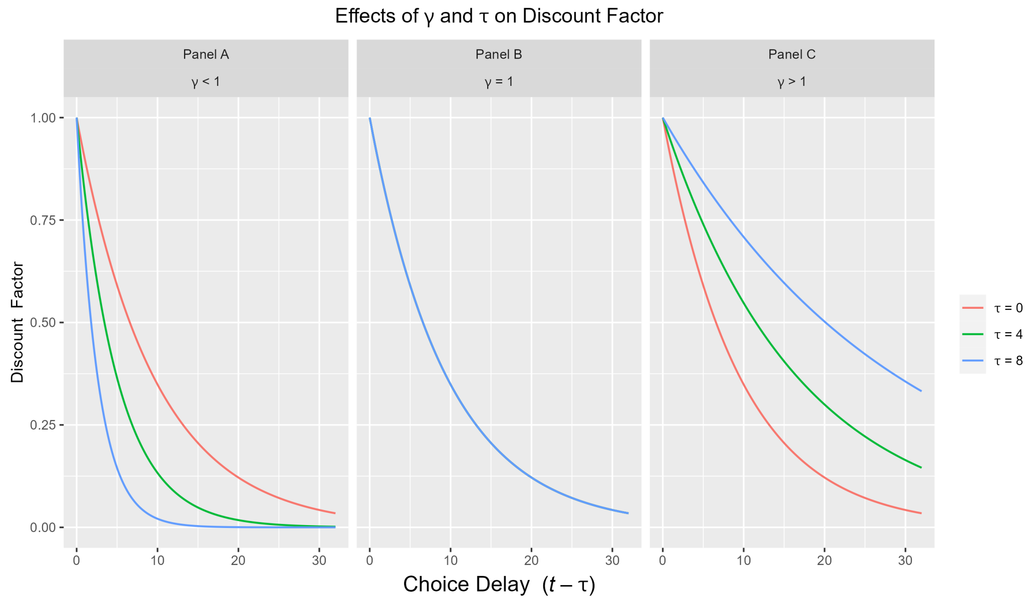

Figure 7. Nested Exponential Discount Function by γ Parameter Value

Notes: for all functions we assume $\mu = \eta = 1$, and $\delta = 0.9$. The x-axis shows choice delay, $t - \tau$.

instead it can be systematic learning (about either the decisions themselves or preferences over them), changes to the macroeconomic environment, the psychological environment, or preference evolution. Indeed, to preview our results, we find evidence of both generally decreasing patience throughout our study, and an uptick in patience during finals week among our student sample.

3.3 Experimental Procedures

The details of our experimental design and analysis were pre-registered on AsPredicted.org, protocol #133180. We recruited subjects from introductory-level economics courses at the University of Oregon. A total of 178 subjects were invited to participate in the experiment, and 153 consented and completed the first survey. After initial recruitment, which occurred both in-person and over email, all communication with subjects was conducted online through the Qualtrics survey platform. While the frequent, short points of contact in our design lend themselves well to an online study, because of the delayed and risky incentives, we decided it was important to run the study in a sample where we could leverage our institutional reputation to establish trust, and where subjects would have in-person access to the study administrators in case of any problems or concerns. Participants were provided with a study email address, as well as the office location and phone number for the faculty study PI.

The study consisted of seven surveys; one every two weeks for a total of twelve weeks. Every subject that completed the first survey earned $5, paid upon completion, and every subject that completed the final survey earned and additional $35 upon completion. All payments were made via instant money transfer on the subjects preferred platform.[9] The first survey featured detailed instructions on our protocol and the last featured an exit questionnaire. Both took subjects roughly 20 minutes to complete. Each intermediate survey featured only prize drawings, price lists, and a very brief questionnaire, each taking roughly 10 minutes to complete.

Within each survey, subjects encountered several 'multiple price lists' (MPLs) that compared two amounts of 'tickets' presented in two columns, where each column represents a prize-drawing for $300 in a specific time period, and each row represents a choice between a (tickets, time period) bundle. Each ticket corresponded to an increment of 1 in 100,000 in odds within a drawing. The left column always offered tickets in a drawing that would take place earlier than the drawing in the right column. For all subjects, in every price list, the left column featured a sequence of 22 rows with 202 to 244 tickets by increments of two, whereas the right column featured a fixed 242 tickets in every row. This design is an adaptation of the multiple-lottery-list elicitation of Belot, Kircher, and Muller (2021). See Appendix Figure B5 for an example price list.

We denote surveys by the number of weeks from the start of the study that they occur (to correspond to the t and τ variables that the discounting models will be based on). In weeks zero and two (the first and second surveys) we offered subjects six choices: three "choice delay" lengths between drawings –two, four, and eight weeks– crossed with two "front-end delays" between the evaluation date and the earlier drawing –zero and two weeks. This design constructs the decision triangles outlined in Section 3.2.1; when a subject responds to sequential surveys, we can observe three decision triangles (one for each choice delay), allowing us to detect violations of any of three time preference properties. Starting in week four, we phased in much longer choice delays of 16 and 32 weeks, at first only with the two-week front-end delay, but then in week six, without the front-end delay. Week 12 was the final week that involved subject choices, and we did not elicit choices with the two week delay, as we would be unable to observe whether they revised those choices in the future. Figure 8 outlines the price lists faced by subjects in each survey.

The goal of any price list elicitation is to observe a single switching point between columns, such that the switch point approximately identifies at which row –as the value of one of the column's rewards changes monotonically– a subject is indifferent between the bundle in the left column and the bundle in the right column. If the subject switches more than once, this

[9] Subjects chose from Venmo, PayPal, CashApp, and Zelle, and we pre-screened potential participants for whether they were comfortable using one of these platforms to receive payment.

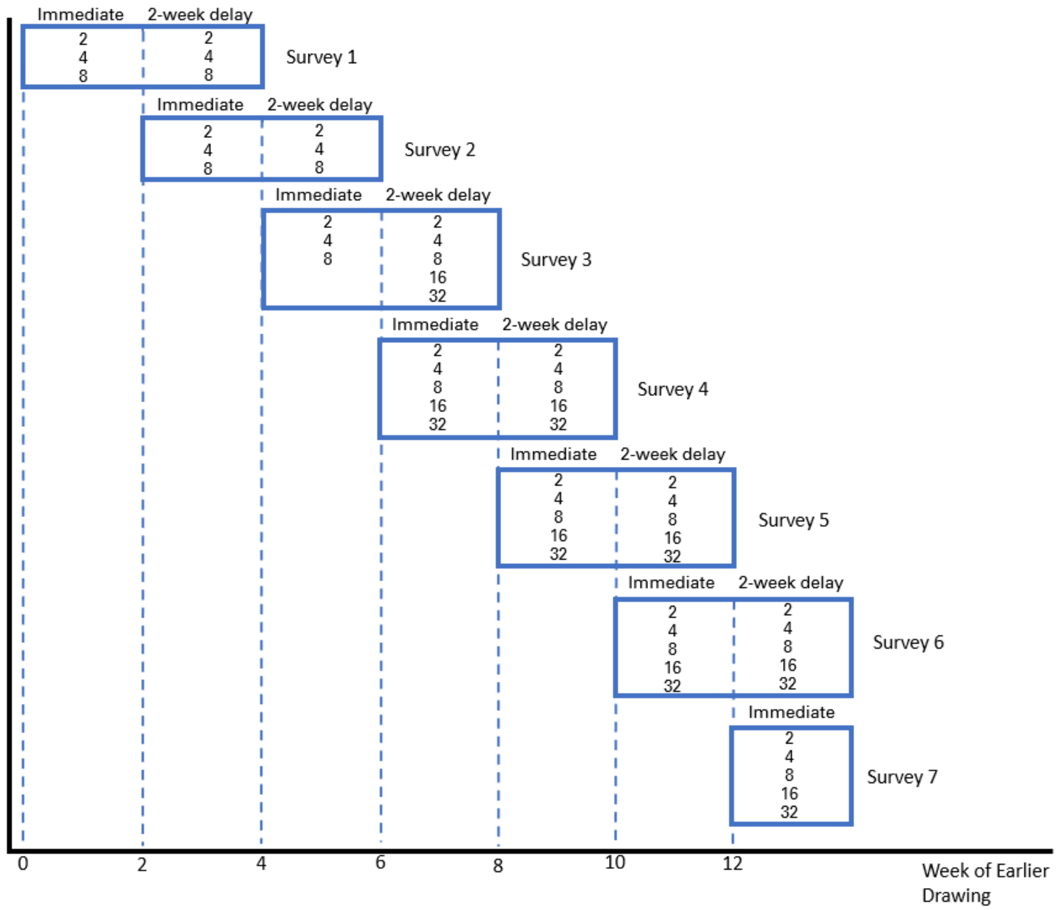

Figure 8. Study Timeline

confounds that identification. A single switching point in tickets was enforced by our elicitation device, a bright yellow sliding bar that separated choices that the subject preferred the earlier drawing from the choices where subjects preferred the later drawing. Subjects were instructed to place the sliding bar *between* the last choice the preferred in the right column and the first choice they preferred in the left column. Specifically, we instructed them to put every row where they prefer the later drawing above the bar, and every row where they prefer the earlier drawing below the bar. The bar itself expressed their switch point in words, which changed when the bar was dragged to a new location. Subjects could leave the bar at the very top –indicating they always prefer the earlier drawing– but they had to actively click and release it to do so.[10] Subjects could not advance through a price list for at least six seconds in order to limit random or thoughtless choices.

[10]We thank Antonia Krefeld-Schwalb for the idea, and the source code to help us build this tool.

All drawings had a positive probability of being realized, but only *one price list per earlier-drawing date* was selected to count. For example, in week zero, with equal probability, one of the price lists with an immediate earlier drawing was randomly selected, and then one of 22 rows was randomly selected from that list, again with equal probability. If the subject picked the immediate drawing in that row, the drawing took place at the end of the week zero survey. If the subject picked the later drawing in that row, the drawing took place *at the beginning* of the survey taking place in that week. The week zero price lists with a two-week front end delay were "saved for later": when subjects returned in week two and made another set of choices featuring an earlier drawing on that date, the price list that counted for that earlier date was randomly selected from the set of week zero and week two choices.[11]

Prize drawings were implemented with the selection of a random number from 1 to 100,000. If the random number was less than or equal to the number of tickets a subject has in a drawing, they would win. Winning automatically triggered an email to a study administrator to enable an immediate transfer of the $300 prize. Three subjects won drawings, including one in the first survey.[12] Many drawings took place after the survey portion of the study ended; in these cases, subjects received emails with a link to participate in the drawing with no data collection. Before participating in incentivized price lists in week zero, subjects had to fill out an example list (however they wanted), correctly use the yellow slider bar to illustrate a specified preference, choose a random drawing number that would allow one hypothetical subject to win and another to lose based on their choices, and answer a question about the independence of choices across price lists. Each survey after week zero included a refresher question that subjects had to correctly answer prior to making their incentivized choices.

Finally, while the data are not the subject of the current paper, each survey concluded with a brief questionnaire on recent life events. Subjects were asked about financial and psychological shocks they may have experienced since the last survey. The same questions were asked every survey, except for one rotating question that assessed risk preferences, cognition, and other measures of time preferences not captured by our main elicitation. The final survey also included an exit questionnaire that asked subjects about their understanding of the survey, demographics and socioeconomic status, expectations about the future, subjective time horizon, patience, and impulsivity.

---

[11] First, we randomly determined with 50-50 chance from which week the choice-that-counted would come from, and then we randomly selected the price list from that week.

[12] As of this draft. The four most-delayed drawings have yet to occur.

The first survey was distributed at 7am on a Wednesday in the middle of the university term. All subsequent surveys were also distributed at 7am on Wednesday, covering the end of that term, finals week, and much of the subsequent term.[13] Each survey was open for a 26 hour window, from 7:00am on Wednesday to 9:00am on Thursday. Subjects were allowed to miss one out of the seven surveys and were barred from further participation after their second missed survey.

3.3.1 Attrition. While a concern in any longitudinal study, the length of time this study covers coupled with seven attempted points of contact, further coupled with the concern that attrition may be non-random with respect to discounting, made attrition a substantial ex-ante concern. For instance, if those that leave the study are less patient than those who remain – because the study features heavily delayed incentives– then the full sample may exhibit increasing-patience time variance in the aggregate despite no individual-level time-variance. Given resource constraints on offering overwhelming incentives and time constraints in hunting down missing individuals to complete a survey within a 26-hour window, we opted for an approach that 1) made completing each survey fast and easy, with multiple reminders and direct, individualized participation links (as opposed to requiring login information), 2) allowed us to carefully test for any selective attrition, and 3) allowed us to analyze time-variance at the subject level.

Of the 153 subjects that completed the first survey, 97 (63%) completed the last survey. 79 (52%) of subjects completed every survey. Throughout the results section of the paper, we treat the data and findings from the "balanced" sample of 79 subjects and 4,345 price lists, but we discuss and reference results from the full sample of 153 subjects throughout, with corresponding tables in the Appendix. None of the main results of the paper depends on which sample we use. In general, the balanced sample results are what we would expect if subjects that attrit make noisier choices and removing them reduces measurement error. We formally test for differential attrition in Section 3.5, and find no significant differences in the discount factor at any point in the study between those who are about to attrit and those who will remain (see Figure 12).

One mitigating factor for concerns about attrition is our focus on subject-level results. The study is designed to allow us to observe the time preference properties exhibited by subjects and estimate subject-level parameters of discount functions. For all subjects that completed at

[13] The term length is 10 weeks, so we were unable to avoid an end of term somewhere in the study. Between having a survey in finals week or spring break, we selected finals week under the assumption that students would at least be on their computers, even if they were busy.

least two subsequent surveys –141 in total, 92% of the initial sample– our experiment provides sufficient data to inform our research question.

3.4 Time-invariant Behavior

We present our results in three sections. Here, we describe discounting behavior in our study from a time-invariant perspective, to establish the comparability of our data with other elicitation efforts in the literature. In the next section, we consider time variance from a reduced form perspective, including a consideration of attrition. Finally, we take a structural approach to time variance, returning to the motivating issue of accurately modeling subjects' axiomatic discounting.

3.4.1 Descriptive Statistics. We first establish sample-wide discounting. Without discounting apparent in the aggregate, it would be unclear whether our elicitation method is inducing subjects to make intertemporal tradeoffs. We measure subject i's implied discount factor in price list j in survey week w, d_{ijw} as the midpoint of their switch interval divided by 242 (the fixed ticket allocation to the later lottery).[14] For example, if a subject prefers 242 tickets in the later draw to 230 tickets in the sooner draw, but 232 tickets in the sooner draw to 242 tickets in the later draw, then $d = \frac{231}{242} = 0.9545$.

Across every survey and price list for subjects that completed every survey[15], the average d is 0.9560 (corresponding to a switch point between 230 and 232 tickets). This is significantly different from one ($p < 0.0001$), consistent with a positive discounting over the tickets in our study.[16] For each week the delay increases, d decreases by 0.07pp ($p = 0.0013$), and when the sooner draw is immediate, d decreases by 0.24pp ($p = 0.0007$), consistent with mild but statistically significant present bias.[17] Assuming exponential discounting, the average *weekly* discount factor across the balanced sample is 0.9920 ,which corresponds to an *annual* discount rate of 51.84%, although we will show in Section 3.4.2 that once adjusted for present bias, the rate is considerably lower. Figure 9 shows the implied discount functions for our data, keeping choices where the sooner lottery is immediate separate from choices where it is delayed by two weeks.

---

[14]When a subject never switches –always preferring the later draw– we assume a switch midpoint of 245. When a subject switches immediately –always preferring the sooner draw– we assume a switch midpoint of 201. In Section 3.4.2 we find that allowing for a censoring process at the endpoints instead has no impact on our estimates.

[15]From this point forward referred to as the balanced sample, and used as the main sample of interest in our analysis

[16]Two-tailed t-test.

[17]OLS regression of d on the delay and an indicator variable for an immediate sooner draw, standard errors clustered at the subject level.

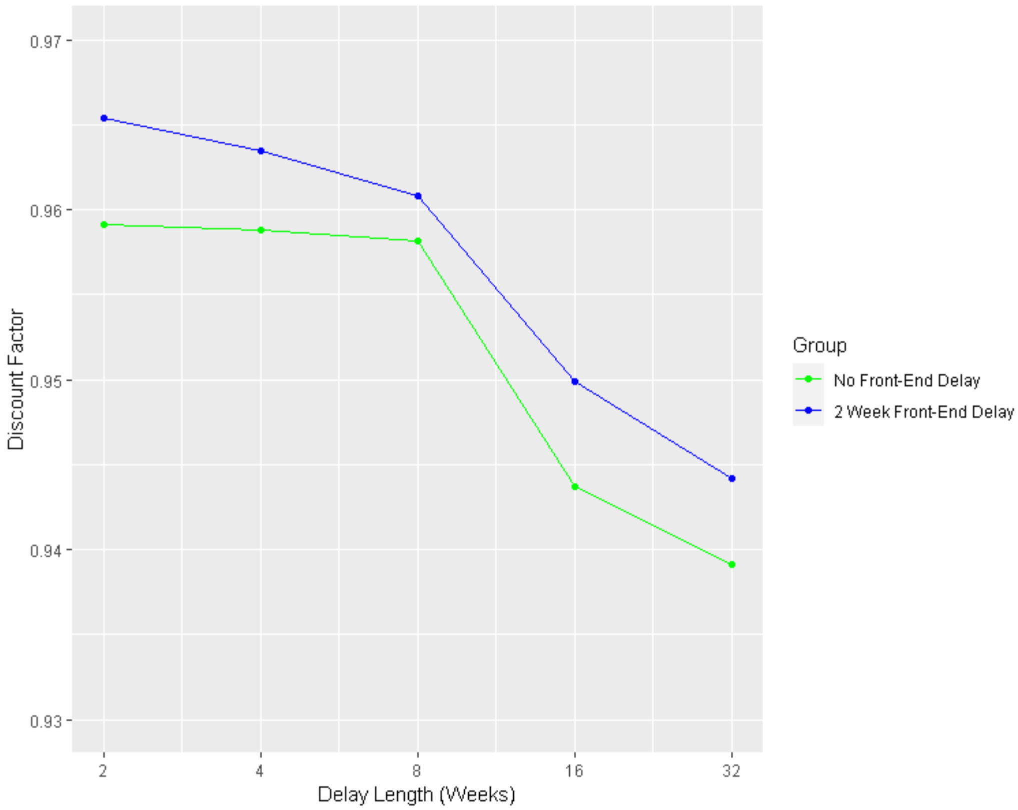

Figure 9. Average Discount Factor by Choice Delay

At every delay length, subjects discount the future more when the sooner draw is immediate, confirming clear present bias in the data. Another notable takeaway from Figure 9 is that when the sooner draw is immediate, subjects treat two-week, four-week, and eight-week delay lengths very similarly. They discount future tickets in all cases, but at a similar rate. When the sooner draw is two weeks in the future, this is not the case, and discounting always responds to delay length.

Given the importance of structural estimation in the literature and this paper, it is important to consider the identification properties of our elicitation. The goal of any price list elicitation is to identify tight bounds on the set containing the indifference point between two options. When subjects never switch (always choose the later draw) or switch immediately (always choose the sooner draw), that set is only bounded below or above. 91% of choices in our study are *interior* switch points, and only one subject failed to deliver a single interior switch point. As such, we prioritize structural estimation strategies that assume a price list switch point identifies a point of *equality* between the sooner and later options. We consider the robustness of our estimates to a set-identification approach in the next section.

3.4.2 Aggregate Discount Parameters. A common application of time preference elicitations in the literature is to structurally estimate discounting parameters –typically an exponential δ parameter, and often a quasi-hyperbolic present-bias β parameter as well. Following our notation from Section 3.2 where the discount factor, $D(\tau, t)$, is a function of payoff period t and evaluation period τ, the quasi-hyperbolic form is $D(\tau, t) = \beta^{\mathbb{1}(t>\tau)} \cdot \delta^{t-\tau}$. For our sooner and later options within each price list, we call t^S and t^L the payoff dates of the respective lotteries.[18] Assuming d identifies the point of indifference between the sooner and later draws in the price list, we have that

$$\beta^{\mathbb{1}(t^S>\tau)} \cdot \delta^{t^S-\tau} \cdot (X-1) = \beta \cdot \delta^{t^L-\tau} \cdot 242 \Rightarrow d = \frac{\beta \cdot \delta^{t^L-\tau}}{\beta^{\mathbb{1}(t^S>\tau)} \cdot \delta^{t^S-\tau}} = \beta^{\mathbb{1}(t^S=\tau)} \cdot \delta^{t^L-t^S}, \tag{3.4}$$

where X is the number of sooner tickets in the switching row (and thus $d = \frac{(X-1)}{242}$). We log-linearize this equality to obtain the regression equation

$$\ln(d_{ijw}) = \mathbb{1}(t_j^S = \tau_j) \cdot \ln(\beta) + (t_j^L - t_j^S) \cdot \ln(\delta) + \varepsilon_{ijw}, \tag{3.5}$$

which we estimate with standard errors clustered at the subject level. Estimates of δ and β from this approach are shown in column (1) of Table 3. Column (2) shows the corresponding estimates from an interval regression (a generalized Tobit model), which uses only the known bounds of each switch interval, including the one-sided bounds as the extremes. Column (3) shows estimate from a non-linear least squares (NLS) regression applied to equation 3.4. NLS will be our preferred technique for estimating the time-variant discount functions, which cannot be neatly linearized, so we include it here as well to show that it produces nearly identical estimates to the other models. When price lists feature larger intervals or do not enforce a single switch point, researchers often take a binary-choice approach to these data at the choice-row level, and use maximum likelihood to estimate the discounting parameters. We take this approach with a logistic choice model, and present estimates in column (4).

Estimates for β and δ are nearly invariant to estimation method. Across the four models, our estimate of annual discount rate ranges from 9% to 11%. This is consistent with the lower end of estimates from the time preference elicitation literature, for example Andersen, Harrison, Lau, and Rutström (2008) and Andersen, Harrison, Lau, and Rutström (2014) find estimates of

[18] Note that t^L is always greater than τ.

Table 3. Aggregate Discount Parameter Estimates From Quasi-hyperbolic Model

Model:	OLS	Interval Regression	NLS	Logistic Choice
	(1)	(2)	(3)	(4)
δ	0.9980	0.9979	0.9981	0.9983
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
β	0.9711	0.9712	0.9723	0.9778
	(0.0032)	(0.0033)	(0.0031)	(0.0031)
r (annual discount rate)	0.1098	0.1146	0.1066	0.0901
	(0.0150)	(0.0174)	(0.0148)	(0.0145)
H_0: $\delta = 1$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$
H_0: $\beta = 1$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$

Notes: standard errors are clustered at the subject level. δ is the weekly exponential discount factor, and $r = \delta^{-52} - 1$ is the annual discount rate. The data consist of 4,345 price lists from 79 subjects.

10% and 9% in a nationally-representative sample of Danes, respectively. Andreoni and Sprenger (2012) estimates an annual rate of 37% using the Convex Time Budget technique with a US undergraduate population, and Kuhn et al. (2017) estimates a rate between 20% and 24% using the same technique among French undergraduates. Our estimates for β are all approximately 0.97, corresponding to present-bias that is roughly equivalent to an increase in delay length of 15 weeks. The estimates are precise; all are statistically significantly different from one with $p < 0.0001$.[19] Using the full sample of subjects, we obtain nearly identical results, shown in Appendix Table A5.

3.4.3 Individual Discount Parameters. When a time preference elicitation is bundled into the design of a larger research study, it is expected to deliver an individual-specific measure of discounting that is either a mediating variable or an outcome of interest. Using the OLS technique from equation 3.5, we successfully estimate individual-specific δ and β parameters for every subject in the sample. Figure 10 shows the distributions of estimates of r, the annual interest rate, and β.

The median r estimate is 9.69%, with 87% of subjects exhibiting positive discounting. The median estimate for β estimates is 0.9822. There is limited evidence of future-bias, with only 19% of subjects assigned $\beta > 1$[20] Overall, we conclude that our time preference elicitation was

[19]The literature is mixed on whether incentivized discounting studies show evidence of present bias, with Andersen et al. (2014); Andreoni, Kuhn, and Sprenger (2017); Andreoni and Sprenger (2012); Augenblick, Niederle, and Sprenger (2015); Kuhn et al. (2017) finding either no or minimal present bias over money in a laboratory setting, Augenblick et al. (2015); Imas et al. (2022) finding present bias over effort, Aycinena, Blazsek, Rentschler, and Sprenger (2022) finding future bias over large stakes in Guatemala, and Belot et al. (2021) finding present bias over lottery tickets.

[20]Results for the full sample can be found in Appendix Figure B7. Medians and proportions are qualitatively similar.

successful, delivering estimated utility parameters for all participants, with both the distributions of individual parameter estimates and aggregate estimates well within the ranges established by previous work. This establishes the platform from which we will examine time-varying properties of discounting.

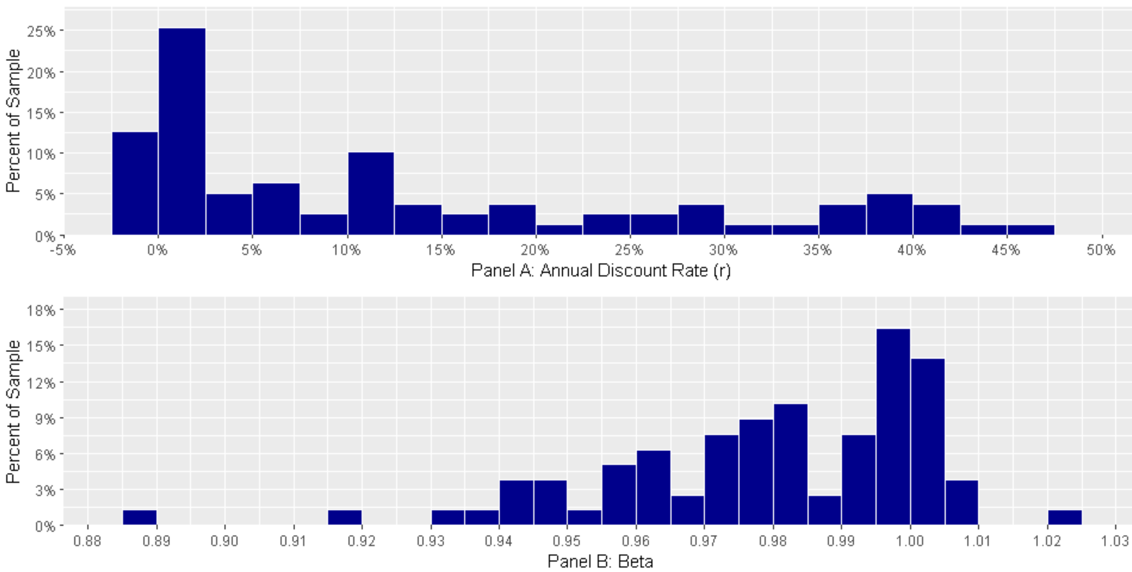

Figure 10. Distributions of Individual Estimates From the β-δ Model

3.5 Time-varying Behavior

3.5.0.1 Aggregate Time Variance. We now exploit the longitudinal nature of our data. In this section, we take a naïve reduced form approach before building up to structural models of time-variance in the next section. We start by augmenting the regressions we use to obtain descriptive statistics of discounting behavior in Section 3.4.1 with survey-week fixed effects. Specifically, we estimate

$$d_{ijw} = \mu_w + \lambda_1 \cdot ((t_j^L - t_j^S) - 8) + \lambda_2 \cdot (1 - \mathbb{1}(t_j^S = \tau_j)) + \eta_{ijw} \tag{3.6}$$

where we use $((t^L - t^S) - 8)$ as the choice delay variable, and $(1 - \mathbb{1}(t^S = \tau))$ as the front-end-delay variable so that the values of μ correspond to average discounting when comparing immediate tickets to tickets in eight weeks. Standard errors (η) are clustered at the subject level. Figure 11 shows these week fixed effects and their 95% confidence intervals.

We highlight three aspects of the data in Figure 11. First, subjects exhibit *decreasing patience* as the study progresses: we see substantially more discounting in week 12 (future utility discounted by 5.6%) compared to week zero (future utility discounted by 3.9%). In terms of an

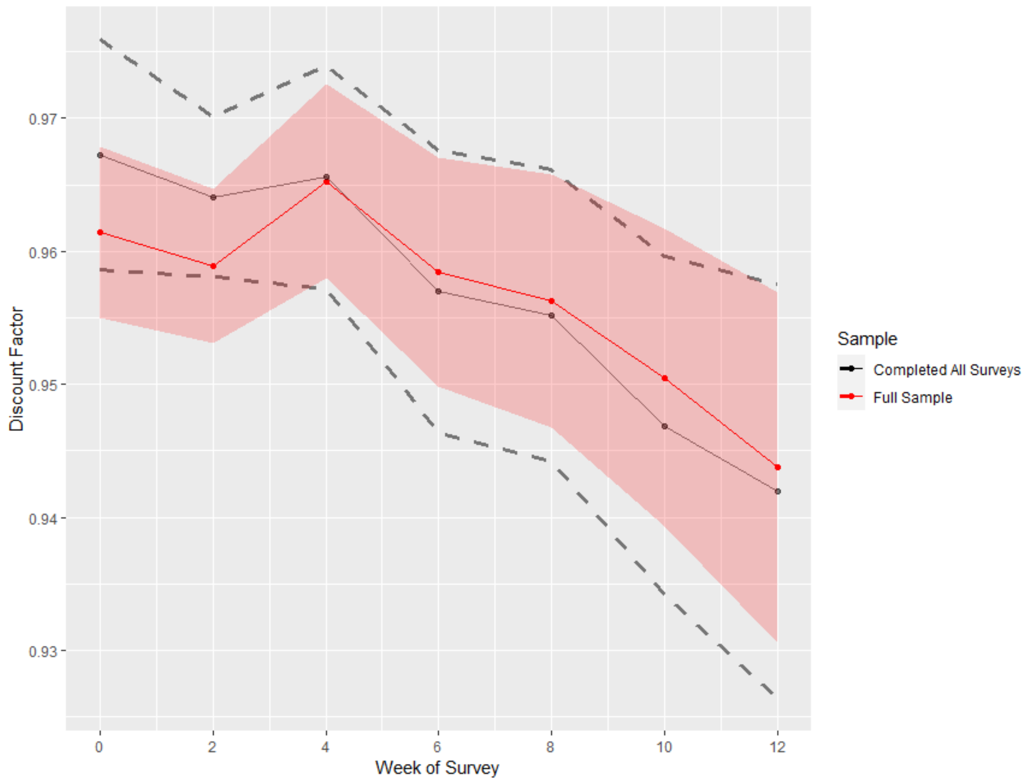

Figure 11. Time Invariance as Measured by Survey-week Fixed Effects

annualized (exponentially) discount rate, this is akin to a shift from 29% to 46% over the course of our study. If we replace μ_w in Equation 3.6 with a linear time trend, $\mu \cdot w$, we estimate that μ is negative (-0.0022) and statistically different than zero ($p = 0.0008$). Second, the time path of discounting is very similar for the full sample and for the sub-sample that completed all seven surveys, especially from week four (the third survey) onward. If anything, using the full sample appears to attenuate the trend of decreasing patience we observe among those who complete all the surveys. If we estimate linear time variance for the full sample we estimate that μ is -0.0013 ($p = 0.0172$). Third, we anticipated a detectable effect of final exam week in our study, and indeed there is a trend-break in survey week four: subjects place a higher value on future tickets in that week. The average discount factor in that survey is higher than in the previous survey, despite the overall decreasing trend ($p = 0.0524$). This is an example of time variance that is likely attributable to the influence of external factors.[21]

[21] This effect is driven not by subjects who fail to complete the finals-week survey, but it is stronger for subjects who will eventually attrit *after* completing it. This is why the trend break is more notable for the full sample.

We estimate additional linear time trends models to check for time-variance that interacts with either the delay length between payoff periods in a choice, or whether the sooner date is immediate. Results are in Table 4[22]. We find robust evidence of across-the-board time variance, at a magnitude of 0.22 percentage points per week. There are no large or statistically significant interactions between the linear time trend and either choice delay or front-end delay in either sample.

Table 4. Linear-time-trend Estimates of Time Variance

	(1)	(2)	(3)	(4)	(5)
Constant	0.9554	0.9698	0.9697	0.9697	0.9697
	(0.0043)	(0.0046)	(0.0047)	(0.0047)	(0.0048)
Choice delay $((t^L - t^S) - 8)$	-0.0007***	-0.0005***	-0.0006	-0.0005***	-0.0006
	(0.0002)	(0.0002)	(0.0003)	(0.0002)	(0.0003)
Front-end delay $(1 - \mathbb{1}(t^S = \tau))$	0.0048***	0.0020	0.0020	0.0022	0.0022
	(0.0014)	(0.0014)	(0.0013)	(0.0029)	(0.0029)
Survey week (w)		-0.0022***	-0.0022***	-0.0022***	-0.0022***
		(0.0006)	(0.0005)	(0.0006)	(0.0006)
$w \cdot ((t^L - t^S) - 8)$			-0.0000		0.0000
			(0.0000)		(0.0000)
$w \cdot (1 - \mathbb{1}(t^S = \tau))$				-0.0000	-0.0000
				(0.0005)	(0.0005)

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$. Coefficients are from linear models with standard errors are clustered at the subject level. The data consist of 4,345 price lists from 79 subjects.

A potential driver of differences in findings between the balanced and full samples is that measurements of time variance can be influenced by attrition. While subjects have a small likelihood of winning a prize each time they take a survey, the guaranteed, substantive payment for participating in the study comes after all surveys have been completed. When it comes time to participate in survey week two, that delayed reward is ten weeks in the future, and it gets two weeks closer in each successive survey. If impatient subjects attrit because they do not value the completion payment, we should expect to see a large initial decrease in discounting from survey weeks zero to two (because week zero is incentivized on its own), and then a stable pattern thereafter.[23] This is not the pattern of time variance we find in Figure 11.

If selective patience-based attrition is affecting our results, it is masking even more decreasing patience than we observe. However, we do not find any relationship between

[22]Corresponding estimates for the full sample are in Appendix Table A6

[23]Purely from a patience perspective, if week two is worth completing for the reward in ten weeks, then week four is worth it for the reward in eight weeks.

discounting and attrition in our study. Figure 12 shows how the discount factor within each survey correlates with whether a subject does not complete the following survey. None of the estimates are statistically significant (p = 0.771, 0.541, 0.354, 0.503, 0.508, and 0.597, for Weeks 0-10 respectively), and the sign is not consistent across surveys.[24]

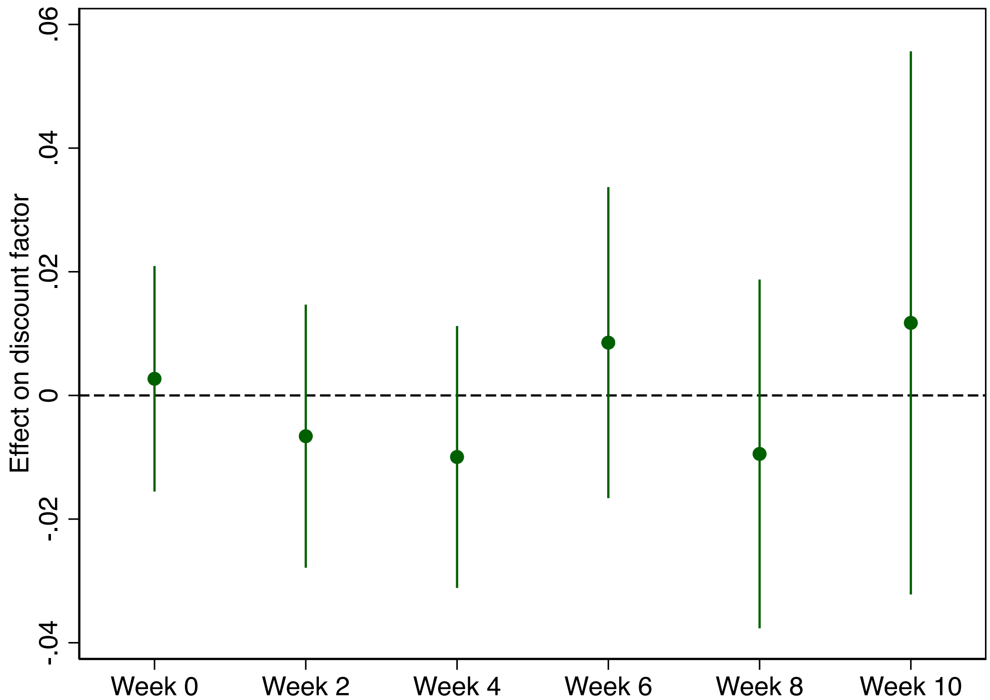

Figure 12. Relationship Between Subsequent Survey Attrition Status and Discounting, Point Estimates and 95% Confidence Intervals
Notes: the plot shows the coefficients and 95% confidence intervals from a regression of the within-survey discount factor on an indicator for whether a subject fails to complete the subsequent survey. We control for the delay-length and front-end delay of each choice so that the discount factor corresponds to a tradeoff between the present and eight weeks in the future.

3.5.0.2 Subject-level Time Variance. A key strength of our data collection is our ability to observe subject-level changes in discounting over time. A reduced-form approach to this is to modify Equation 3.6 to feature subject-week fixed effects, μ_{iw}, rather than pooled week fixed effects. Figure 13 below shows the distribution of these fixed effects by subject across weeks. Across all survey weeks, modal discounting looks very stable, while the right tail expands consistently throughout the study.

While these distributions do not track individual subjects, it suggests a group of time-invariant discounters and a group of decreasing-patience discounters. To confirm this, we measure

[24]Estimates correspond to equation 3.6 estimated at the survey level, with the survey fixed effect replaced with an indicator for whether a subject returns to participate in the next survey.

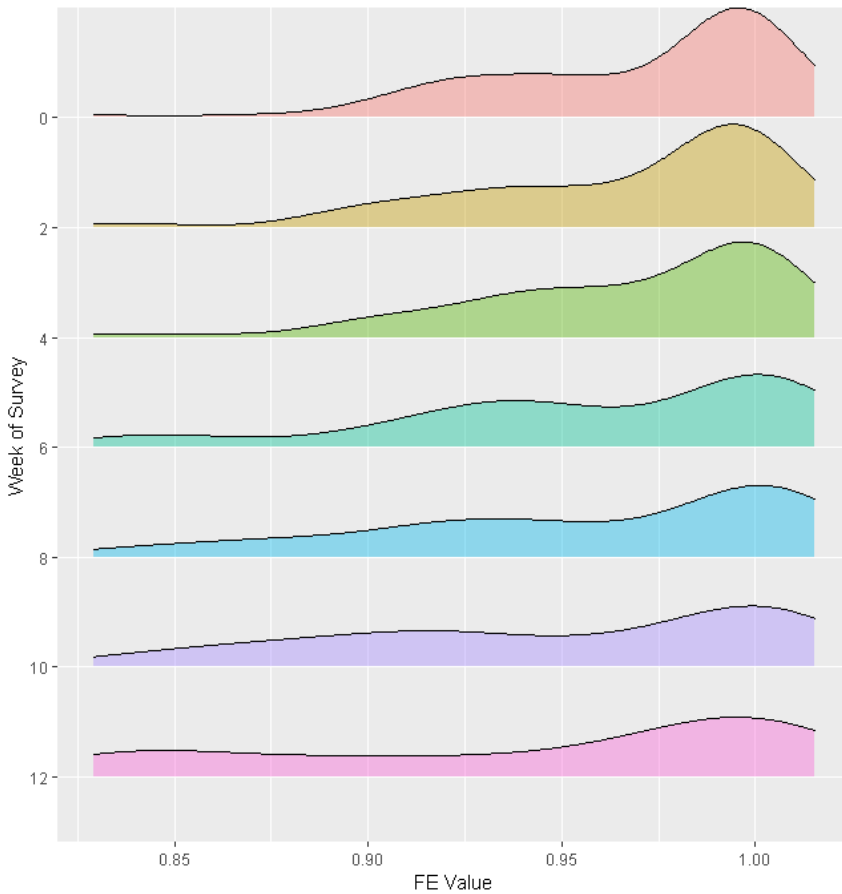

Figure 13. Distributions of Subject-week Fixed Effects, by Survey-week
Notes: the plot shows the distribution of subject-week fixed effects from a regression of the discount factor on those fixed effects and controls for the delay-length and front-end delay of each choice so that the discount factor corresponds to a tradeoff between the present and eight weeks in the future.

individual time trends using linear, individual-specific survey week time trends, replacing μ_w in Equation 3.6 with $\rho_i \cdot w$, while also allowing for fixed level differences across individuals using subject fixed effects, μ_i. Figure 14 plots the ρ_i coefficients along with their standard errors[25]. We can reject time-invariance for 57% of subjects, with 37% and 20% exhibiting decreasing and increasing patience, respectively. Note that conditional on exhibiting time variance, the magnitude of decreasing patience is larger on average than the magnitude of increasing patience, leading to the overall trends shown in Figures 11 and 13.

[25] $N = 72$. 7 subjects had insufficient variation to estimate coefficients

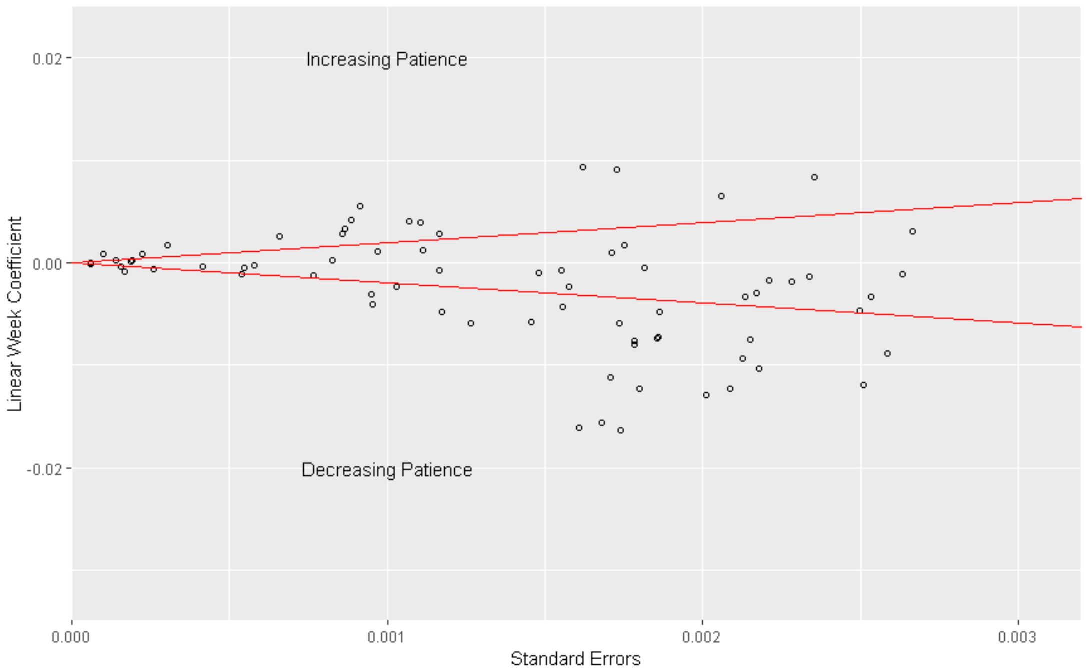

Figure 14. Subject-level Estimates of Time Variance
Notes: the plot shows the coefficients and standard errors from subject-specific regressions of the discount factor on a linear survey week variable. We control for the delay-length and front-end delay of each choice so that the discount factor corresponds to a tradeoff between the present and eight weeks in the future. The red lines plot 1.96 and -1.96 times the standard error. Points lying outside the bounds of the red lines correspond to statistically significant changes in patience over the course of the study.

3.5.1 Inconsistency and Non-Stationarity. In Section 3.2, we illustrated how time-varying behavior implies either non-stationary or inconsistent behavior (or both). We have also empirically established both decreasing patience and inconsistent/non-stationary behavior in the data –we have not yet drawn a distinction between the two because under the assumption of time-invariance, they are constrained to be the same. In this section, we allow them to differ and explore the connection between the two; how much observed inconsistency and non-stationarity is explained by time-variance?

We start by separately estimating non-stationarity and inconsistency. Equation (3.6), models the impact of choice delay (t_j) and front-end delay ($t_j - \tau_j$) with survey-week fixed effects. We adjust these variables such that the constant term describes the discount factor for an eight-week choice delay without a front-end delay. The fixed effects limit the model to within-survey variation in $t_j - \tau_j$, corresponding to a stationarity violation. In columns (1) and (2) of Table 5, we present the survey-week fixed effect models, with and without allowing an interactive effect between choice delay and front-end delay, respectively. We do not estimate a statistically significant relationship between a *within-survey* front-end delay and the discount factor ($p =$

0.1406 and $p = 0.1036$, respectively), although the coefficient is positive, qualitatively consistent with a present-bias. The estimates in column (2) show that for shorter choice delays, we do not observe significantly more stationarity-violating present bias, although we do in the full sample. Results in Appendix Table A7.

Table 5. Separate Estimates of Non-stationarity and Inconsistency

Violation type:	Stationarity		Consistency	
	(1)	(2)	(3)	(4)
Constant	0.9698	0.9698	0.9698	0.9698
	(0.0046)	(0.0046)	(0.0046)	(0.0046)
Choice delay $((t^L - t^S) - 8)$	-0.0005***	-0.0005**	-0.0005***	-0.0005**
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Front-end delay $(1 - \mathbb{1}(t^S = \tau))$	0.0020	0.0064***	0.0065***	0.0063***
	(0.0014)	(0.0013)	(0.0016)	(0.0016)
$((t^L - t^S) - 8) \cdot (1 - \mathbb{1}(t^S = \tau))$		-0.0001		-0.0001
		(0.0001)		(0.0001)
Survey-week FEs	Y	Y	N	N
Drawing-week FEs	N	N	Y	Y

Notes: Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Standard errors are clustered at the subject level. The data consist of 4,345 price lists from 79 subjects.

To estimate inconsistency instead, we replace the week effects with *drawing-week* fixed effects in columns (3) and (4) of Table 5. The drawing week of a decision is the week the *earlier* drawing takes place. For example, decisions in the first survey with a two-week front-end delay have the same drawing week as decisions in the second survey without the front-end delay. We find that making a choice two weeks in advance results in a significantly higher discount factor ($p < 0.0001$ and $p = 0.0001$, respectively). These findings are consistent with Section 3.2, where time-inconsistent present bias is a direct result of decreasing patience under the assumption of stationarity.

Given that inconsistency and non-stationarity differ in their magnitude and frequency, time-variance must play a role in explaining these phenomena. We follow Janssens et al. (2017) in asking how many of the instances of inconsistency and non-stationarity in our study can be attributed to time-variance. Referring to the decision triangle language of Section 3.2, given any level of time invariance $(a - c)$, we can calculate whether any observed inconsistency $(b - c)$ is exactly predicted by, in excess of, or less than what would be predicted assuming stationarity holds $(a = b)$. We can do the same for any observed non-stationarity as well. Specifically, for all decision triangles where inconsistency exist $(b \neq c)$, we calculate $(a - c)/(b - c)$. This is effectively the percentage of time inconsistency that is attributable to time variance. *Time* *variance explains 72% of the inconsistency in our sample.* While inconsistency is typically attributed to a behavioral present-bias (when $c < b$) in decisions, this result suggests that a large part of it could be explained by decreasing patience ($c < a$). For all decision triangles where non-stationarity is present ($a \neq b$), we construct a similar ratio of $(a - c)/(a - b)$. This ratio describes the percentage of non-stationarity that can be attributed to time variance. Time variance explains 56% of non-stationarity in our sample.

With evidence of time variance, inconsistency, and non-stationarity in our sample, we now follow Halevy (2015) and attempt to classify each subject into one of the five possible combinations of time preference properties illustrated in Table 2. We start by classifying each applicable decision within a triangle as either time-invariant or time-varying. This requires that subjects must have attended at least two sequential surveys to be included in this analysis. Out of 153 total subjects, this applies to 141[26]. These 141 subjects account for 6,267 of the 6,367 price lists collected in our study. containing 2,506 decision triangles. Of these 141 subjects, 130 subjects (92%) make a time-varying decision at some point in the study. Of the 2,506 decision triangles, 1,500 (60%) include a time-varying decision[27]. Both the percentage of subjects and percentage of choices that are time-varying exceed the findings of Halevy (2015), that found a maximum of 56.42% of subjects exhibiting time-varying preferences. While the longer longitudinal dimension of our study contributes to that difference at the subject level, this should not be the case at the triangle level. We return to this issue later in this section.

We count the occurrences of inconsistent and non-stationary decisions in a similar manner to the time-varying decisions. While the same 141 subjects and 6,267 price lists generate 2,710 decision triangles where b and c the test for inconsistency is possible and 2,782 decision triangles where the test for non-stationarity is possible, we restrict our analysis to the 2,506 decision triangles that overlap with the preceding time-invariance analysis. We find that 129 subjects (91%) make an inconsistent decision at some point, with 1,448 (58%) of triangles being inconsistent. Similarly, 127 subjects (90%) make a non-stationary decision at some point, with 1,276 (51%) triangles in total being non-stationary. Both proportions exceed those from Halevy (2015) again, where at most 66.33% of subjects exhibited inconsistent preferences and at most 55.22% exhibited non-stationary preferences.

[26]This eliminates ten subjects who only responded to the first survey and two more subjects who only responded to the first and third.

[27]This proportion is very similar if we also consider time-variance in the choices made *with* a front-end delay. In this case, 94% of subjects display time-varying behavior in 59% of triangles.

Table6 shows the classification of subjects by property in our sample in column (1) of Panel A. Only eight subjects (10%) exhibit invariant preferences. Of these eight, seven (9% of the sample) make consistent and stationary decisions. However, these nine subjects display constant discounting, not reacting to differences in choice delay or front-end delay either. Thus, among subjects who reacted to any stimuli, all violated a time preference property. 69 (87%) subjects violate each property at some point, two (2%) maintain only stationarity, and none maintain only consistency[28]

A key caveat to these statistics is that the prevalence of violations in our sample can be, in part, attributed to the number of opportunities we offer to record an error in decision-making within a 22-row price list. As we will show in the next section, true errors in decision-making can be down-weighted by the underlying trend of one's decisions. By contrast, similar studies with only two points of contact have only one comparison to draw conclusions from, meaning what could be an error in decision-making may instead be treated as an out-right violation of one of the properties. Our study design allows for formal modeling and testing of each property at the individual level. However, we first address this issue with two ad-hoc approaches that allow us to continue to compare our results to other studies. First we allow for a margin of error that spans a standard deviation (14 tickets) of the decisions we observe (± 6 tickets due to the two-ticket gap between each choice row). If a subject's choice is within six tickets on either side of their previous choice, we treat the decision as if it did not violate the tested time preference property.[29] Results are in column (2) of Panel A in Table 6. Second, instead of aggregating to the subject-level, we treat each decision triangle as a distinct observation. Results are in column (1) of Panel B, and results using both modifications are in column (2) of Panel B.

Using the margin of error at the subject level, 16 subjects (20%), exhibit time-invariant behavior. All 16 of these subjects are also consistent and stationary. The increase in the size of this group comes mostly from a reduction in size of those that violate every property, now 60 subjects (76%). Three subjects maintain only stationarity (4%) and still none maintain only consistency (1%) using this window of error. Taking the decision-triangle level of analysis approach instead we find that around 43% of decision triangles maintain time-invariance, while 56% do not. This is close to the minimum classification from Halevy (2015), which has 44% of subjects exhibiting time-invariance. Our most conservative estimates of time-varying behavior

[28]These results are qualitatively invariant to the inclusion of the full sample, seen in Appendix Table A8.

[29]One problem with adding a margin of error to this accounting exercise is that it is possible for only one property to be violated, even though theoretically, when two hold, the third is guaranteed. When this is the case, we classify a decision triangle or subject as satisfying all time preference properties.

Table 6. Subject- and Decision-triangle-level Classification of Time Preference Properties

Margin of Error:	None	± 6 tickets
Properties of choices	(1)	(2)
Panel A: Subject level		
A) Invariant, consistent, stationary	8.86%	20.25%
B) Invariant only	1.27%	0%
A) + B)	10.13%	20.25%
C) Stationary only	2.53%	3.80%
D) Consistent only	0%	0%
E) None	87.34%	75.95%
C) + D) + E)	89.87%	79.75%
Panel B: Triangle level		
A) Invariant, consistent, stationary	37.45%	63.40%
B) Invariant only	6.07%	7.01%
A) + B)	43.52%	70.41%
C) Stationary only	15.24%	14.40%
D) Consistent only	7.86%	7.96%
E) None	33.39%	7.23%
C) + D) + E)	56.48%	29.59%

Notes: Data 79 subjects that completed all surveys. Column(1) assumes any deviation from is a time preference property violation. Column(2) assumes any deviation greater than six tickets from an analogous choice is a time preference violation, capturing a range of a standard deviation around a choice ($\sigma/2 \approx 7$, which due to interval of two tickets between choice rows corresponds to equality within six tickets).

come from combining the decision triangle sorting with the margin of error. This method shows 70% of decision triangles maintain time-invariance, which is close to the maximum classification from Halevy (2015), which has 68% of subjects exhibiting time-invariance. These results suggest observing a single decision triangle per individual understates potential for time-varying behavior. Defined strictly, time-invariance is a universal feature of subjects who responded to our stimuli.

3.6 Estimating Time-varying Discount Functions

Beyond establishing the prevalence of time variance and its role in explaining time-inconsistency and non-stationarity in our sample, our goal in this paper is to develop a structural toolkit for estimating time variance as a part of the discount function, and formally testing whether time preference properties hold at the individual level. Recall the "nested exponential" discount function we introduced in Section 3.2:

$$D(\tau, t) = \delta^{\frac{(t^{\mu} - \tau^{\mu})^{\eta}}{\gamma^{\tau}}}. \tag{3.7}$$

Assuming $\delta \neq 1$, when all three (none) of the other parameters differ from one, the function fulfills none (all) of the three time preference properties. If any pair of μ, η, and γ equal one, then only one of three properties holds. In the case where all properties holds the function nests classical exponential discounting, and in the case where only invariance holds ($\mu = \gamma = 1$, $\eta \neq 1$, e.g. the "behavioral" scenario of present bias) the functions nests something akin to Weibull discounting.

We estimate parameters for the nested exponential discount function at the aggregate sample level using the non-linear least squares method, separately allowing for each possible class of time preference properties. Results and descriptive statistics of these estimations are presented in Table 7. Column (1) enforces exponential discounting ($\mu = \eta = \gamma = 1$, thus $D(\tau, t) = \delta^{t-\tau}$). The weekly discount factor of 0.9973 translates into an annual discount rate (r_0) of 15%.

The model in column (2) allows for time-inconsistency and non-stationarity, but enforces time-invariance, our Weibull-like specification ($\mu = \gamma = 1$, thus $D(\tau, t) = \delta^{(t-\tau)^{\eta}}$). This model fits the data much better than the exponential model, even taking its extra parameter into account, as shown by the decrease in the Akaike information criteria (AIC). This method for assessing model fit penalizes the more complex models for the higher degrees of freedom they have in fitting the data. The estimate of $\delta = 0.9779$ (the weekly discount factor for a choice delay of one week) suggests much more rapid discounting of near-term future outcomes, but the "present-bias" parameter η flattens the discount function for longer choice delays. $\eta = 0.3811$ turns a choice delay of two weeks into an effective delay of 1.3 weeks, and a 52-week choice delay into an effective delay of 4.5 weeks, such that the annual discount rate is lower than in the exponential model: 11%.[30] Qualitatively, this means we estimate significant hyperbolicity in discounting, clearly rejecting $\eta = 1$ ($p < 0.0001$). To asses the magnitude of present-bias, we construct measures PB_S and PB_C, which are discount factor ratios that correspond to measurements of the quasi-hyperbolic β through a stationarity violation or a consistency violation, respectively. Specifically,

$$PB_S = \frac{D(0,8)}{D(0,0)} \Big/ \frac{D(0,10)}{D(0,2)}, \text{ and } PB_C = \frac{D(2,10)}{D(2,2)} \Big/ \frac{D(0,10)}{D(0,2)} \tag{3.8}$$

When invariance holds these measures are identical, and we estimate $PB_S = PB_C = 0.9755$, which is significantly different from one ($p < 0.0001$).[31]

[30] Without the hyperbolic "flattening" of the discount factor, a weekly discount factor of 0.9778 exponentially implies an annual discount rate of 221%.

[31] While it lies outside the nested exponential model, it is worth pointing out that the data clearly reject the quasi-hyperbolic model in favor of this Weibull-like restricted model. Hyperbolicity is a key feature of our data; indeed the single parameter true hyperbolic discount function $D(\tau, t) = \frac{1}{1+\alpha \cdot (t-\tau)}$ fits substantially better than the single-parameter classic exponential model.

Table 7. Aggregate Parameter Estimates for the Nested Exponential Model

Restrictions:	$\mu = \eta = \gamma = 1$	$\mu = \gamma = 1$	$\mu = \eta = 1$	$\eta = \gamma = 1$	None
Properties					
Invariant	✓	✓	X	X	X
Stationary	✓	X	✓	X	X
Consistent	✓	X	X	✓	X
	(1)	(2)	(3)	(4)	(5)
δ	0.9973	0.9779	0.9973	0.9787	0.9756
	(0.0003)	(0.0027)	(0.0004)	(0.0045)	(0.0039)
η	1	0.3855	1	1	0.4250
	.	(0.0326)	.	.	(0.0452)
γ	1	1	1.0012	1	0.9703
	.	.	(0.0164)	.	(0.0176)
μ	1	1	1	0.5141	0.7321
	.	.	.	(0.0494)	(0.1380)
Log-likelihood	-18,149.6	-17921.5	-18,149.6	-18,068.5	-17,917.2
AIC	36,301.2	35,846.9	36,303.2	36,141.0	35,842.4
r_0	0.1499	0.1082	0.1513	0.1786	0.0880
	(0.0177)	(0.0130)	(0.0250)	(0.0192)	(0.0138)
r_{12}	$= r_0$	$= r_0$	0.1491	0.1113	0.1178
			(0.0222)	(0.0149)	(0.0174)
PB_S	1	0.9755	1	0.9771	0.9791
	.	(0.0030)	.	(0.0052)	(0.0116)
PB_C	1	$= PB_S$	1.0001	1	0.9744
	.		(0.0007)	.	(0.0032)
H_0: $\delta = 1$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$
H_0: $\eta = 1$	.	$p < 0.0001$	.	.	$p < 0.0001$
H_0: $\gamma = 1$	.	.	$p = 0.9441$	.	$p = 0.0955$
H_0: $\mu = 1$	.	.	.	$p < 0.0001$	$p = 0.0558$
H_0: $\mu = \gamma = 1$	.	.	.	.	$p = 0.1385$
H_0: $\mu = \eta = 1$	.	.	.	.	$p < 0.0001$
H_0: $\eta = \gamma = 1$	.	.	.	.	$p < 0.0001$
H_0: $\mu = \eta = \gamma = 1$	.	.	.	.	$p < 0.0001$
H_0: $r_0 = r_{12}$	.	.	$p = 0.9440$	$p < 0.0001$	$p = 0.1295$
H_0: $PB_S = 1$	.	$p < 0.0001$	.	$p < 0.0001$	$p = 0.0703$
H_0: $PB_C = 1$	.	$p < 0.0001$	$p = 0.9443$	.	$p < 0.0001$
H_0: $PB_S = PB_C$	.	.	.	.	$p = 0.6666$

Notes: standard errors are clustered at the subject level. $r_\tau = D(\tau, \tau + 52)^{-1} - 1$ is a measure of the annual discount rate. PB_S and PB_C are stationarity-based and consistency-based measures of present-bias. The data consist of 4,345 price lists from 79 subjects that completed all seven surveys. All estimates are from non-linear least squares regressions.

The model in column (3) allows for time-variance and inconsistency, but enforces stationarity, the property specification that does not appear in the literature ($\mu = \eta = 1$, thus $D(\tau, t) = \delta^{\frac{t-\tau}{\gamma^{\tau}}}$). While this model represents a modest improvement in fit than exponential discounting in column (1), the fit is much worse than the time-invariant model in column (2). Without access to hyperbolicity, the estimates attempt to fit both short and long choice delays with an initial ($\tau = 0$) weekly discount factor of 0.9973, which translates to an annual rate (r_0) of 15%. Time variance due to $\gamma = 1.0012$ marginally reduces that weekly discount factor as the study progresses. The annual discount rate at the end of our study ($\tau = 12$, termed r_{12} in the table), the annual discount rate falls to just under 15%. In an attempt to best-fit the data, this model estimates statistically insignificant *increasing-patience* time variance ($p = 0.9440$ for the test of $r_0 = r_{12}$, $p = 0.9441$ for the test of $\gamma = 1$), at odds with the reduced form results. A side effect of this is that the model predicts statistically insignificant *future-bias* ($p = 0.9443$ for the test of $PB_C = 1$); this is a direct implication of increasing patience under the assumption of stationarity. These results are more pronounced, and statistically significant for the full sample. Results in Appendix Table A9.

The model in column (4) allows for time-variance and non-stationarity, but enforces consistency, the "consistent hyperbolic" property pattern studied in Strulik (2021), but in a more flexible form ($\eta = \gamma = 1$, thus $D(\tau, t) = \delta^{t^{\mu} - \tau^{\mu}}$). This model lies in between the poor performance of the exponential and stationary-only models in columns (1) and (3), and the improved fit of the invariant-only model in column (2). The initial ($\tau = 0$) weekly discount factor estimate is 0.9787 which captures substantial discounting of near-term future outcomes, with the hyperbolicity parameter μ flattening the discount function for longer choice delays, much like η in the invariant-only model. $\mu = 0.5141$ turns the choice between zero and two weeks into an effective delay of 1.4 weeks, and the choice delay between zero and 52-weeks into an effective delay of 6.7 weeks, yielding an initial ($\tau = 0$) estimate of the annual discount rate of 18%.

Fitting hyperbolicity in a model constrained to be time-consistent forces it to predict *increasing-patience* time variance, just like the stationary-only model, however this model finds the increasing patiences statistically significant. The end-of-study ($\tau = 12$) annual discount rate predicted by this model is 11%, which we can reject is equal to the initial rate ($p < 0.0001$). The μ parameter acts on τ as well as t, thus $\mu < 1$ delivers both hyperbolicity and *increasing hyperbolicity* in τ. While the effective choice delay between zero and 52 weeks is 6.7 weeks, the effective choice delay between 12 and 64 weeks is 4.1 weeks. This model cannot deliver both decreasing patience and hyperbolicity, which explains its poor fit. As a side effect of increasing patience under the assumption of consistency, this model predicts statistically significant present-bias ($p < 0.0001$ for the test of $PB_S = 1$).

The model in column (5) puts no restrictions on the parameter values. It features the best overall fit of the data, but crucially, its AIC is nearly identical to that of the invariant-only model in column (2). We estimate an initial weekly discount factor of 0.9756, and an initial annual discount rate of 9%. This grows to 12% over the course of the study, reflecting the decreasing patience we identified in the reduced-form analysis, but we can only marginally reject that the two are the same ($p = 0.1295$). However, we can marginally reject the assumptions of the invariant-only model. We can reject $\gamma = 1$ ($p = 0.0955$) and that $\mu = 1$ ($p = 0.0558$). We can also nearly reject that $\gamma = \mu = 1$ ($p = 0.1385$). This model predicts a larger degree of consistency-violation present-bias ($PB_C = 0.9744 < 1$, $p = 0.0032$) then stationarity-violation present-bias ($PB_C = 0.9791 < 1$, $p = 0.0032$), although we cannot reject that they are equal in magnitude ($p = 0.6666$). We ultimately find that this fully estimated nested model, is best fitting. The relative likelihood of the invariant-only model to the fully estimated nested model is 0.11, whereas it is essentially zero for the other three models.[32]

Figure 15 shows the discount factors predicted by the estimated functions along with observed discount factors (averages of d from the data) by delay length (for choice delays of two, four, and eight weeks, in Panels A, B, and C, respectively) and week of survey (x-axis). Visualizing the model fit this way helps demonstrate how the nested and invariant-only functions performed better in fitting the data, and what they still fail to capture. The key feature of all three panels is that only the fully estimated model can mimic the decreasing patience shown by subjects. The invariant-only model performs well despite not predicting decreasing patience because –like the nested model– it adapts well to changes in choice delay lengths. The other three models –exponential, stationary-only and consistent-only– not only cannot match the pattern of decreasing patience, but also change very little in response to changes in delay length, resulting in poor fit relative to the invariant-only and nested model.[33]

[32]This measure ranges from zero to one, and assesses the relative explanatory of two models. Using AIC, the relative likelihood is defined as $exp(\frac{AIC_1 - AIC_2}{2})$ where AIC_1 is better-fitting model (lower AIC).

[33]These models fit the level well in the 16-week choice delay, but then badly over-react to the 32-week choice delay.

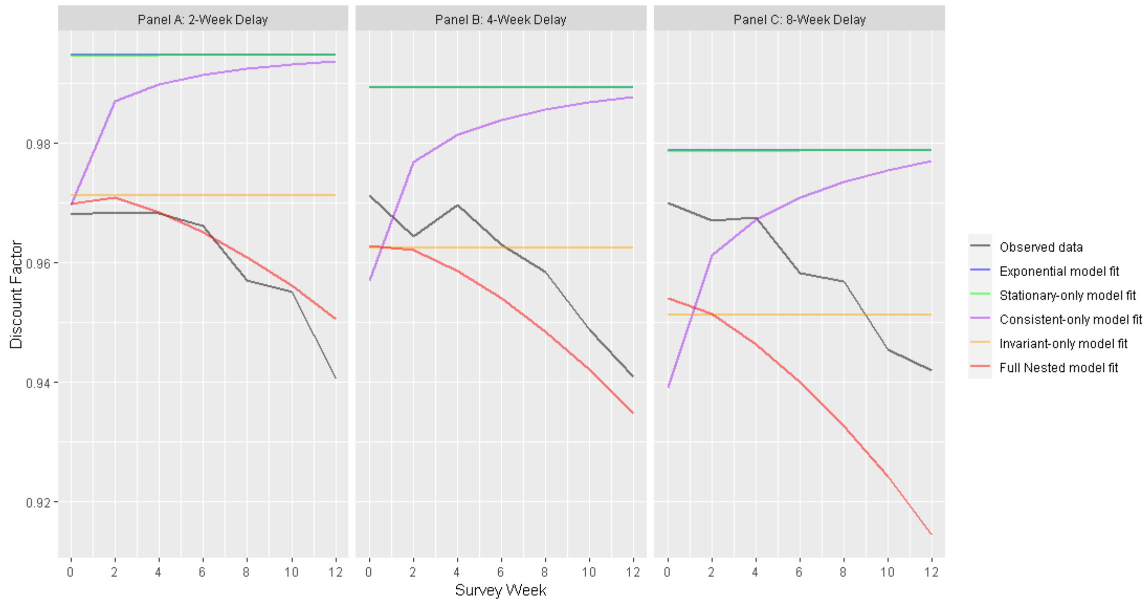

Figure 15. Observed and Estimated Discount Factors by Delay Length

3.6.1 Subject-level Estimates. The aggregate results suggests that marginal time-variance is present, but that when it comes to structurally fitting the preferences of a representative agent, decreasing patience is not as important as other features of the data – notably hyperbolicity– for a model to fit. If a researcher were to opt for an invariant-only discount function to model the preferences of this sample, predictions would be accurate in general, but could underestimate the level of discounting that would occur in future periods. However, even if a time-invariant model may be satisfactory to describe the sample, the literatures on time-preference estimation and the importance of present-bias are focused on individual heterogeneity in preferences. If 20% of people are quasi-hyperbolic discounters, we may not estimate a representative discount function with β significantly less than 1, but we still want to correctly identify that sample. In a time-varying world, this necessitates allowing for heterogeneity in time-variance as well.

We attempt to estimate the parameters of the nested exponential model for each of the 79 subjects that completed every survey, once again using the non-linear least squares method. Given lack of variation in responses for some subjects, we successfully estimate parameters for all models for 50 subjects. As a point of reference to aggregate results, we plot the distribution of parameter estimates for the fully estimated nested model in Figure 16. Again using AIC we determine which model best fits each individual in sample, with the results in Table 8.

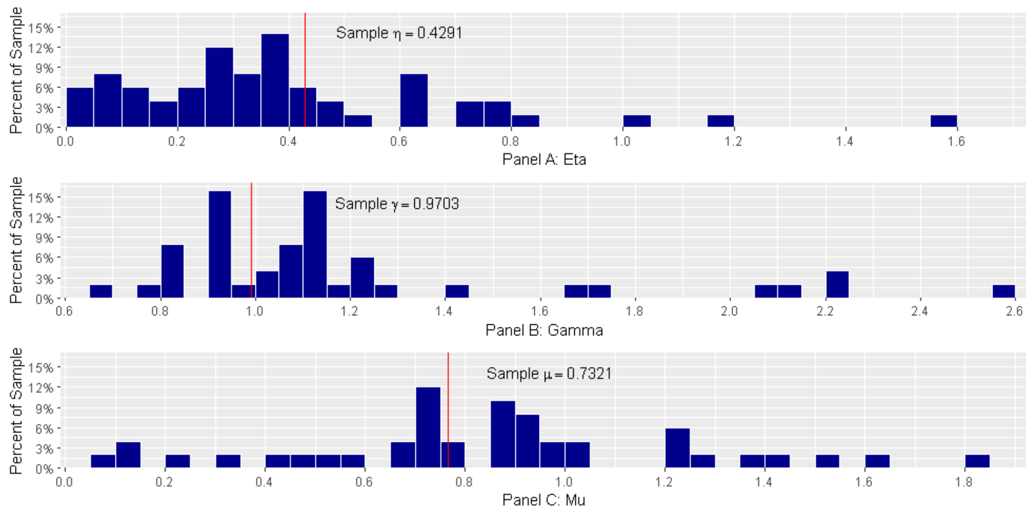

Figure 16. Individual Parameter Estimates for the Unrestricted Nested Exponential Model
Notes: 9 and 8 estimates, respectively, are suppressed from the right tails of the distributions in Panel (B) and Panel (C) for visual clarity.

Table 8. Subject-level Structural Classification of Time Preference Properties

Properties of choices	
A) Invariant, consistent, stationary	0%
B) Invariant only	42%
A) + B)	42%
C) Stationary only	4%
D) Consistent only	6%
E) None	48%
C) + D) + E)	58%

Notes: Data comes from 50 subjects for which we were able to obtain parameter estimates from all models and completed every survey.

Approximately 42% of subjects are best fit by time-invariant models. While this is consistent with the lower end of the range found by Halevy (2015) (44%), the distribution between exponential discounters and invariant-only discounters that we find is very different; no time-invariant discounters in our study are best fit by the standard exponential that also allows for consistency and stationarity. The *lowest* estimate of that proportion in Halevy (2015) is 77%. Similar to the reduced-form methods of sorting used in the previous section, we find few subjects are best-described as being stationary-only or consistent-only. These results hold qualitatively in the full sample, where we successfully estimate parameters for all models for 91 subjects. One key difference is that 14% of time-invariant discounters, six subjects, are best represented by the standard exponential discount function. Analogous estimation and sorting results can be found in Appendix Figure B8 and Appendix Table A10

3.7 Discussion

We use a 12-week longitudinal study to detect violations of three related time preference properties across seven surveys. We find that in all but our most conservative classification approach, the largest single group of subjects by adherence to the TICS properties is "none". A minimum of 30% and a maximum of 90% of subjects exhibit time-variance, a pattern of behavior that confounds the identification of time preferences in traditional time-invariant models. Average discounting exhibits systematically decreasing patience over the course of our study. Without adjusting for time-variance, this would've led us to dramatically overstate the degree of behavioral time-inconsistency in our study (e.g. present-bias). Indeed, 72% of the time-inconsistency we observe in our data is explained by time variance. This is yet another reason that the term "present-bias" to describe decision-making may be ill-founded when not accounting for time-varying behavior.

Beyond highlighting this measurement issue, we developed a new discount function – the nested exponential– that allows researchers to estimate a model that is fully general with respect to the TICS properties, yet tractable enough to deliver subject-level parameter estimates in a simple non-linear least squares estimation. This model is easy to pick up and estimate in a longitudinal research study that features all three choices in a decision triangle, but also it can be used in a calibration to explore the robustness of estimates of time-inconsistency or non-stationarity derived from only two of the three choices.

The presence and acknowledgment of time-varying preferences has potential to impact policy, particularly relating to commitment mechanisms and behavioral nudges. If a policy is meant to reinforce commitment to a consumption plan, take-up of the commitment mechanism may be lower than expected when time-varying preferences exist and are internalized by the policy target. If the target understands that their preferences are changing over time, committing to a previous decision would be welfare-damaging. Similarly, the effect of behavioral nudges may be overstated if the target has time-varying preferences that coincide with the direction the policy is meant to push them.

A deep exploration of the sources of time-varying behavior is left for future work using the data we have collected. Past studies on this topic focus on forces of risk, uncertainty, and liquidity-constraint as likely causes of this behavior. More thorough analysis of these potential factors, as well as the impact of demographics, socioeconomic status, cognition, and macroeconomic changes, will help in determining in what contexts time-varying behavior may be present and to what extent it affects outcomes or can be controlled for.

CHAPTER IV

AN EXPLORATION OF THE MECHANISMS OF TIME-VARYING TIME PREFERENCES

4.1 Introduction

With limited exceptions, the economic literatures on variation in preferences over time and non-classical, "behavioral" preferences have largely proceeded separately. Only recently has a small literature focused on how these two research areas could, in the special case of time preferences, be confounded. Are dynamic preference reversals –a hallmark of non-classical time preferences– manifestations of a *time-invariant* predilection for immediate gratification, or a consequence of preferences that are *time-varying* such that analogous intertemporal tradeoffs are made differently depending on measurement date? Halevy (2015), Janssens et al. (2017), and Holloway, Wiegand, and Kuhn (2024) all empirically show that a majority (but not all) of observed dynamic preference reversals can be explained by time-varying preferences. Given the apparent importance of time-variance in explaining intertemporal choices, in this paper I use the data from Holloway et al. (2024) to explore whether there are types of individuals who are more or less likely to exhibit time-variance and whether time-variance in discounting is attributable to individuals' changing economic or psychological conditions.

After outlining the theory, experimental design, and results from Holloway et al. (2024) in Sections 4.2 and 4.3, I perform heterogeneity analysis on discounting behavior in the study. In addition to a panel of discounting behavior, I collected data on subjects' demographics, socioeconomic status, risk preferences, patience, and other behavioral characteristics that can be treated as static variables in determining 'who' may exhibit time-varying preferences. These data are covered in Section 4.4. Additionally, I collected dynamic data on financial and psychological well-being and food security that is used to help inform 'what' may cause time-varying preferences. These data are covered in section 4.5.

In both sets of data, I routinely find evidence that socioeconomic status is related to subjects' average discounting behavior, but rarely meaningfully informs a pattern of time-varying behavior, especially within-subject over time. I do find, however, that higher perceptions of patience stability, as well as some measures of patience (but not all), are correlated with less impatient decisions over time relative to the average subject, suggesting subjects have some awareness of time-varying behavior they may exhibit. The relative importance of these variables is reinforced by a robustness test employing Bayesian Model Averaging. I also find that discounting is not driven over time by whether a subject's preferences were recently realized (in the form of real money drawings).

These results suggest time-variance may be best described as a preference parameter that is allowed to evolve over time or as an output of a learning process not captured by this study. To explore this possibility, in Section 4.7 I develop an asymptotically time-varying discount model. This model approaches time-invariance as the evaluation period tends to infinity. I find this model performs at least as well as the analogous time-invariant model and outperforms the standard exponential model significantly. This model also estimates that between 33% and 53% of subjects exhibit time-varying behavior, in line with Halevy (2015). While this model does not explicitly inform the mechanisms that cause time-varying behavior, I see it as a useful tool to allow others to seamlessly allow for time-variance in their models, allowing for wider-spread study of the property.

4.2 Theory

We adopt notation from Halevy (2015) to describe the time-invariance property of time preferences. Consider an agent in calendar time $\tau \geq 0$ evaluating bundles (ε, t) and (ψ, t'), which deliver rewards ε, ψ at time periods t, $t' \geq \tau$. We assume that the agent has complete and transitive preferences over all such bundles, according to a preference ordering that may depend on τ, $\succsim_\tau$.

Under classical consumer theory assumptions, including additive separability of utility across time periods, this preference relationship is represented by a *discount function*, $D(\tau, t)$, such that

$$(\varepsilon, t) \succsim_\tau (\psi, t') \iff D(\tau, t) \cdot \varepsilon \ \geq \ D(\tau, t') \cdot \psi \ .^1 \tag{4.1}$$

$D(\tau, t)$ describes the agent's discount "factor" when evaluating, in period τ, a payoff in period t.

Definition 4.2.1. An agent's preferences satisfy **time invariance** if

$$D(\tau, \tau + \Delta_1)\, \varepsilon \ \geq \ D(\tau, \tau + \Delta_2)\, \psi \iff D(\tau', \tau' + \Delta_1)\, \varepsilon \ \geq \ D(\tau', \tau' + \Delta_2)\, \psi$$

Agents whose preferences are *time-invariant* will not exhibit preference reversals when the evaluation period and payoff periods for all bundles are shifted by the same amount.

[1] Without loss of generality, we replace $u(\varepsilon)$ and $u(\psi)$ with ε and ψ, under the assumption an instantaneous utility function representation exists.

Time-invariant preferences are well represented in the literature. Discount functions that admit these preferences include the standard exponential, as well as the hyperbolic and β-δ models regularly used to model time-inconsistent decisions. These models all share a common property: They are a function of only the distance between payoff time and evaluation time, $t - \tau$. Since these models are only concerned with this distance, the $D(\tau, t)$ function is typically reduced to just $D(t)$, where τ is always 0, without loss of generality.

There exists only a nascent literature on discount functions that do not satisfy time-variance (Strulik, 2021), and on observed time-variant behavior (Brownback et al. (2023); DeJarnette (2020); Halevy (2015); Harrison et al. (2023); Holloway et al. (2024); Imas et al. (2022); Janssens et al. (2017). None of the empirical papers in this literature observe time variance using more than two points of observation, and none systematically explore what might cause or predict individual-level time-variance. Holloway et al. (2024) outlines the conditions necessary for a time-varying discount function and introduces the "nested exponential" discount function, which allows for any feasible configuration of time preferences according to the time-invariance, consistency, and stationarity properties depending on its parameter values. This approach treats time-variance as a directly-modeled part of discounting, in the same way that quasi-hyperbolic discounting directly models violations of stationarity and consistency based on the β parameter. In this paper we also explore time-varying discount functions that are defined by allowing a constant parameter of an otherwise time-invariant (and widely-used) function to be dependent on evaluation period. This is a "reduced-form" way of capturing the idea that time-variance could represent subjects learning about their own preferences over time. Furthermore, introducing time-variance through a portable extension of an existing model can reduce the cost of uptake in application for other economists and aids in interpretation of results relative to the baseline model Rabin (2013). We investigate the performance of models employing this method and their interpretation in Section 4.7.

4.2.1 Measuring Time-variance. Consider a decision maker (DM) faced with a choice between an immediate reward at $t = \tau = 0$ and a reward on some future date $t = t_1 > 0$. Call a a scalar measure of how much their choice favors the reward at the later date, t_1, i.e. higher values of a result from more patient discounting of that delayed reward. After time passes, such that $\tau = t_2$, the DM faces another similar choice: call b the measure of the choice between immediate utility at $t = \tau = t_2$ and future utility at $t_2 + t_1$. Then $a - b$ is a measure of time-variance; how much a choice over fixed relative dates changes over time.

Economic experiments concerned with time preferences largely dismiss the potential for time-varying preferences or avoid measurements that may reveal them. Interest in changes in time preferences are generally centered around time-inconsistency of agents. Therefore, measurements of subjects' discount factors are taken with some front-end delay before the decisions realizes and then again when realization of the decision is immediate to determine if a preference reversal over the exact same trade-off between payoffs (in magnitude, delay length, and ultimate payoff periods) is affected by the evaluation period the decision takes place in. This comparison cannot be used to make statements about time-invariance, as the two decisions are referring to payoffs on the same date, only changing evaluation period, as opposed to shifting both payoff and evaluation periods.

Even when measured appropriately, there have been limited attempts to account for time-varying preferences in the literature. Reduced-form approaches have been used to correct observed magnitudes of time-inconsistency when time-variance is also observed (e.g. Janssens et al. (2017) or Harrison et al. (2023)). There have not been prior attempts, however, to rely on the theory outlined above to structurally model time-varying time preferences in the same way hyperbolic and quasi-hyperbolic discount functions were developed to model time-inconsistency. In Section 4.7, we introduce one way of injecting time-varying preferences into existing analyses: a structural model leveraging time-varying parameters of an otherwise time-invariant discount function.

When it comes to estimating time-invariance and the parameters of a potentially time-varying discount function, as opposed to a reduced-form invariance correction, the longitudinal dimension of the data determines the precision and reliability of those estimates. Existing work by Halevy (2015) and Janssens et al. (2017) collects data on only a single comparison of discount factors. This style of measurement cannot distinguish idiosyncratic subject-level noise at different dates from a systematic learning trend(about either the decisions themselves or preferences over them), changes to the macroeconomic environment, the psychological environment, or preference evolution. Further, little work exists on which of these factors actually drive time-varying time preferences [2].

In order to address whether time-variance manifests as a systematic trend or idiosyncratic noise, we observe six comparisons (seven discount factors) per subject per delay length of later payoff over the course of a 12-week study. To then determine what factors drive the trend we find, we collect static data on subjects' risk preferences, cognition, and underlying time preferences

[2]Janssens et al. (2017) do find correlation between time-variance and liquidity constraints

and dynamic data on subjects' financial and psychological well-being. The finer details of our experimental design are discussed in the next section.

4.3 Experimental Procedures and Background

The details of the experimental design and analysis were pre-registered on AsPredicted.org, protocol #133180. Holloway et al. (2024) establishes the prevalence and importance of time-variance in these data. Most instances of time-inconsistency are explained by time-variance, there is an aggregate trend of decreasing discounting over the course of the study, and models that structurally incorporate time variance fit the data best at both the aggregate and individual levels. Given that time variance does appear to be an important feature of the data, the goal of this paper is to understand which subjects are more likely to exhibit time variance and why. I briefly describe the procedures of Holloway et al. (2024), but refer to that paper for the full details.

We recruited subjects from introductory-level economics courses at the University of Oregon. A total of 178 subjects were invited to participate in the experiment, and 153 consented and completed the first survey. After initial recruitment, which occurred both in-person and over email, all communication with subjects was conducted online through the Qualtrics survey platform. The study consisted of seven surveys; one every two weeks for a total of twelve weeks. Every subject that completed the first survey earned \$5, paid upon completion, and every subject that completed the final survey earned an additional \$35 upon completion. All payments were made via instant money transfer on the subject's preferred platform.[3] The first survey featured detailed instructions on the protocol and the last featured an exit questionnaire. Both took subjects roughly 20 minutes to complete. Each intermediate survey featured only prize drawings, price lists, and a very brief questionnaire, each taking roughly 10 minutes to complete.

Within each survey, we elicited subjects' time preferences using an adaptation of the multiple-lottery-list elicitation of Belot et al. (2021). The price lists of this elicitation offer subjects choices between lottery tickets in \$300 prize drawings that take place at different points in time. We offered subjects choices with varying delay lengths between the prize drawings ("choice delay") and with or without a two-week delay between the date of choice and the earlier prize drawing ("front-end delay"). See Appendix Figure B5 for an example price list. We denote surveys by the number of weeks from the start of the study that they occur (to correspond to the t and τ variables that the discounting models will be based on). Week 12 was the final week that

[3] Subjects chose from Venmo, PayPal, CashApp, and Zelle, and we pre-screened potential participants for whether they were comfortable using one of these platforms to receive payment.

involved subject choices, and we did not elicit choices with the two week front-end delay, as we would be unable to observe whether they revised those choices in the future in order to detect time-consistency. Figure 8 in Chapter III outlines the price lists faced by subjects in each survey.

All drawings had a positive probability of being realized, but only *one price list per earlier-drawing date* was selected to count. For example, in week zero, with equal probability, one of the price lists with an immediate earlier drawing was randomly selected, and then one of the price list rows was randomly selected from that list, again with equal probability. If the subject picked the immediate drawing in that row, the drawing took place at the end of the week zero survey. If the subject picked the later drawing in that row, the drawing took place *at the beginning* of the survey taking place in that week. The week zero price lists with a two-week front end delay were "saved for later"; when subjects returned in week two and made another set of choices featuring an earlier drawing on that date, the price list that counted for that earlier date was randomly selected from the set of week zero and week two choices.[4]

Before participating in incentivized price lists in week zero, subjects had to fill out an example list (however they wanted), correctly use the yellow slider bar to illustrate a specified preference, choose a random drawing number that would allow one hypothetical subject to win and another to lose based on their choices, and answer a question about the independence of choices across price lists. Each survey after week zero included a refresher question that subjects had to correctly answer prior to making their incentivized choices.

The first survey was distributed at 7am on a Wednesday in the middle of the university term. All subsequent surveys were also distributed at 7am on Wednesday, covering the end of that term, finals week, and much of the subsequent term.[5] Each survey was open for a 26 hour window, from 7:00am on Wednesday to 9:00am on Thursday. Subjects were allowed to miss one out of the seven surveys and were barred from further participation after their second missed survey.

4.3.1 Questionnaires. Each survey concluded with a brief questionnaire on subjects' experiences since the last survey. Subjects were asked about financial events, their subjective financial stability, food security, stress, and work and school commitments. Exact wording of questions can be found in Table 9. Each of these questions were asked in every survey so that

[4]First, we randomly determined with 50-50 chance from which week the choice-that-counted would come from, and then we randomly selected the price list from that week.

[5]The term length is 10 weeks, so we were unable to avoid an end of term somewhere in the study. Between having a survey in finals week or spring break, we selected finals week under the assumption that students would at least be on their computers, even if they were busy.

repeat observations of time preferences correspond to repeat observations of these variables. I refer to these as “dynamic” variables.

We also have “static” variables –one observation per subject– from a variety of sources. Included at the end of each questionnaire was a rotating question that changed every survey. These questions assessed risk preferences, cognition, and other measures of time preferences not captured by the main elicitation. All questions asked as part of these surveys can be seen in Table 10. The final survey also included an exit questionnaire that included questions about subjects' expectations of the future, their overall experience, background information such as demographics and socioeconomic status, and a battery of questions relating to subjective measures of patience, time horizons, and impulsivity. Wording for these questions are found in Table 11. Additionally, an initial screening survey was run in the process of recruiting subjects for the study. This survey asked students for their major, their gender, and whether they used student loans, family funds, their job, scholarships, or some combination of the four to pay for their schooling. I use gender and student loan dependence in the static analysis, noting that this data is available for all 153 subjects, while other variables are constrained to attendance in the week collected.

4.3.2 Attrition. While a concern in any longitudinal study, the length of time this study covers coupled with seven attempted points of contact, further coupled with the concern that attrition may be non-random with respect to discounting and measures collected in the exit survey, made attrition a substantial ex-ante concern. Given resource constraints on offering overwhelming incentives and time constraints in hunting down missing individuals to complete a survey within a 26-hour window, we opted for an approach that 1) made completing each survey fast and easy, with multiple reminders and direct, individualized participation links (as opposed to requiring login information), 2) allowed us to carefully test for any selective attrition, and 3) allowed us to analyze time-variance at the subject level.

Of the 153 subjects that completed the first survey, 97 (63%) completed the last survey. 79 (52%) of subjects completed every survey. Summary statistics for the Week 0 and Week 12 samples can be seen in Table 12. The samples do not differ significantly on any observable dimensions. It should be noted, however, that any results relying on rotating supplemental questions or exit survey questions are conditional on subjects having reached that far in the survey.

Table 9. Dynamic Survey Questions

Variable	Description
Financial Events	In the last two weeks, have any of the following events occurred? Select all that apply. **(A)** I received a paycheck for work **(B)** I received a sizeable money transfer from some other source (i.e. paycheck-sized or similar, but from another source) **(C)** I learned about a new source of future income **(D)** I made a large, expected payment using my own funds (rent, tuition, car payment, large purchase, etc.) **(E)** I had a large, unexpected expense come up that I paid using my own funds **(F)** I had a large, unexpected expense come up that I have not yet been able to pay **(G)** I learned about a new source of future expenses **(H)** None of the above
Financial Stability	On a scale from 1 to 7, where 1 represents "very concerning" and 7 represents "very solid," in the last two weeks, how would you describe your financial situation? *
Food Security: Type/Amount*	On a scale from 1 to 7, where 1 represents "very concerning" and 7 represents "very solid," in the last two weeks, how would you describe your financial situation? **(A)** Enough of the kind of food you want to eat **(B)** Enough, but not always the kind of food you want to eat **(C)** Sometimes not enough to eat **(D)** Often not enough to eat
Food Security Stress*	In the last two weeks, did you ever worry about whether your food would run out before you got money to buy more? **(A)** Yes **(B)** No
$500 Expense	In the last two weeks, if an unexpected expense of $500 had come up, would you be able to (or were you able to) pay it? **(A)** Yes **(B)** Maybe **(C)** No
Stress	On a scale from 1 to 7, where 1 represents "never" and 7 represents "all the time," in the last two weeks, how often did you find yourself stressed or overwhelmed with your responsibilities (e.g. work, school, family, etc.)?
Work/School Hours	In the last two weeks, how many hours do you think you spent on work and/or school?

Notes: *Questions adopted from USDA (2012)

Table 10. Rotating Survey Questions

Variable	Description
Week 0: Triple Return Investment (TRI)	Imagine that you are given $300, and an opportunity to invest some of it in a project that will succeed with 50% chance and fail with 50% chance. You get to keep whatever you don't invest. If the project succeeds, your investment is tripled, if it fails, you lose your investment. How much would you invest? Your answer should be an integer (whole number) from 0 to 300. From U. Gneezy and Potters (1997).
Week 4: Wait to Redeem (WTR)	Imagine you have won a prize, but the longer you wait to redeem that prize, the bigger it gets. Which of the following options would you most prefer?*
Week 6a: Inflation Perception	What do you think is the current level of inflation in the U.S., as defined by the most recent annualized Consumer Price Index (CPI)? Answer in terms of percent (e.g. 5 = 5% annual inflation). Don't look anything up: please give your best guess. If you think there is deflation, use a negative number. If you think there are no price changes happening, use zero. Decimals are allowed.
Week 6b: Inflation Expectation	What do you expect the rate of inflation in the U.S. to be over the next 12 months? Answer in terms of percent (e.g. 5 = 5% inflation). Don't look anything up: please give your best guess. If you expect there to be deflation, use a negative number. If you expect no price changes, use zero. Decimals are allowed.
Week 8a: Allais Paradox Part 1	Which of the two prize draws below would you prefer? **PRIZE DRAW A** A random number from 1-100 is drawn : If it is less than or equal to 33: you win $2500. If it is between 33 and 100: you win $2400. If it is exactly 100: you do not win. **PRIZE DRAW B** A random number from 1-100 is drawn: If it is less than or equal to 100: you win $2400
Week 8b: Allais Paradox Part 2	Which of the two prize draws below would you prefer? **PRIZE DRAW C** A random number from 1-100 is drawn : If it is less than or equal to 33: you win $2500. If it is greater than 33: you do not win. **PRIZE DRAW D** A random number from 1-100 is drawn: If it is less than or equal to 34: you win $2400. If it is greater than 34: yo do not win.
Week 10: Interest Rate Cognition	Imagine that you are given $300, and it is invested in an account that pays 20% interest per year. How much money would be in the account after 2 years?
Week 12a: Triple Return Investment (TRI)	Same as Week 0
Week 12b: Wait to Redeem (WTR)	Same as Week 4 but with delays reduced by 8 weeks each*
Week 12c: Cognitive Reflection Part 1	A bat and a ball cost $1.10 in total. The bat costs $1.00 more than the ball. How much does the ball cost (in $)? Enter only a number (decimals OK) into the box.
Week 12d: Cognitive Reflection Part 2	If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets (in minutes)? Enter only a number (no decimals) into the box.
Week 12e: Cognitive Reflection Part 3	In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake? Enter only a number (no decimals) into the box.

Notes: *Multiple choice options for these questions can be found in Appendix Figures B9 and B10

Table 11. Exit Survey

Variable	Description
Expectations: Income	Over the rest of 2023, do you expect any substantial changes in your regular monthly income (or spending money, if it comes from a non-income source)? **(A)** No **(B)** Yes, I expect my income will increase **(C)** Yes, I expect my income will decrease
Expectations: Expenses	Over the rest of 2023, do you expect any substantial changes in your regular monthly expenses? **(A)** No **(B)** Yes, I expect my expenses will increase **(C)** Yes, I expect my expenses will decrease
Expectations: Inflow/Outflow	Over the rest of 2023, do you expect any large, one-time, expenses or inflows of money? **(A)** No **(B)** Yes, I expect one or more large inflows of money **(C)** Yes, I expect one or more large expenses
Experience: Understanding	On a scale of 1 to 7, where 1 is did not understand at all, and 7 is understood completely, how well did you feel you understood the prize draws?
Experience: Choice Factors	Can you very briefly explain the key factors that determined your prize draw choices (e.g., timing, number of tickets, etc.)?
Experience: Known Winners	As of the last survey (two weeks ago), three participants have won $300 prize draws. Do you know anyone who won a prize drawing in this study? **(A)** Yes **(B)** No
Experience: Heard Winners	(If answered No in) Did you hear about anyone winning a prize drawing in this study? **(A)** Yes **(B)** No
Experience: Known Participants	Do you know any other participants in this study? **(A)** Yes **(B)** No
Background: Cohort	What year are you in at UO? **(A)** Freshman **(B)** Sophomore **(C)** Junior **(D)** Senior **(E)** 5th year (or greater) undergraduate **(F)** Graduate Student **(G)** Alumni **(H)** Other (please specify)
Background: Parent's Edu.	What is the highest level of education one of your parent's has completed? **(A)** High school/GED **(B)** Associate's Degree **(C)** Bachelor's Degree **(D)** Master's Degree **(E)** Doctorate **(F)** None of the above
Background: Race	What race/ethnicity do you most identify with? **(A)** American Indian or Alaskan Native **(B)** Asian **(C)** Black or African American **(D)** Hispanic or Latin American **(E)** Native Hawaiian or Pacific Islander **(F)** White or Caucasian **(G)** None of the above
Patience	On a scale of 1 to 7, where 1 is not at all patient, and 7 is extremely patient, how patient do you consider yourself to be?
Patience: Stability	Do you think your level of patience is stable over time, or does it depends on what's going on in your life? **(A)** Stable **(B)** Somewhat Stable **(C)** Somewhat Unstable **(D)** Very Unstable
Patience: Falling Through	How often do you fail to follow through on plans you made for yourself in the past? **(A)** Never **(B)** Sometimes **(C)** Almost Always
Patience: Time Horizons	For each of the time horizons below, do you consider events in your life that you expect at that time to be in the "present", the "near future" or the "distant future"?*
Patience: Procrastination	On a scale of 1 to 7, where 1 is not at all, and 7 is very frequently, how often do you procrastinate?
Exercise	On average, how many hours per week do you spend exercising?
Exercise: Regret	Do you exercise as much as you'd like? **(A)** Yes, I am satisfied with how much I exercise **(B)** No, I wish I exercised more **(C)** No, I wish I exercised less
Nicotine	How often do you use nicotine products? **(A)** Daily **(B)** Weekly **(C)** Monthly **(D)** On rare occasion **(E)** Never
Nicotine: Regret	(If answered Nicotine with anything other than 'Never') Do you wish you used nicotine products less often? **(A)** Yes, I wish I used them less often **(B)** No, I have no problem with my nicotine use
Streaming	When thinking of the streaming services that you pay for, which of the following would say is most accurate?**(A)** All of the streaming services I pay for are worth the cost of a subscription **(B)** Most of the streaming services I pay for are worth the cost of a subscription **(C)** Some of the streaming services I pay for are worth the cost of a subscription **(D)** Few of the streaming services I pay for are worth the cost of a subscription occasion **(E)** None of the streaming services I pay for are worth the cost of a subscription **(F)** I don't pay for any streaming services

Notes: *Question formatting can be seen in Appendix Figure B11

Table 12. Summary Statistics

Week 0 Statistic	Week 0 Sample	Week 12 Sample	$H_0 : \mu_0 = \mu_{12}$
Subject Characteristics			
Female (%)	58.2	57.7	$p = 0.9729$
Student Loan Dependence (%)	39.22	32.99	$p = 0.3218$
Fraction Invested (Risk Tolerance)	0.4744	0.4608	$p = 0.6898$
Time Preferences			
Tickets Chosen: 2 week delay, Immediate	233.86	234.87	$p = 0.4667$
Tickets Chosen: 2 week delay, In advance	234.39	234.78	$p = 0.8111$
Tickets Chosen: 8 week delay, Immediate	234.05	234.72	$p = 0.6614$
Tickets Chosen: 8 week delay In advance	233.88	235.51	$p = 0.3081$
Observations	153	97	

Notes: Student Loan Dependence (%) describes the percentage of the sample that stated they in-part or fully rely on student loans to pay for college. Fraction Invested (Risk Tolerance) refers to subjects' response to the Week 0 Rotating Question adopted from U. Gneezy and Potters (1997) (Table 10). Time preference entries refer to the average response in week 0 of the survey to the appropriate multiple price list.

One mitigating factor for concerns about attrition relating to questions asked in every survey is my focus on subject-level results. The study is designed to allow us to observe the time preference properties exhibited by subjects and estimate subject-level parameters of discount functions. For all subjects that completed at least two subsequent surveys –141 in total, 92% of the initial sample– the experiment provides sufficient data to inform my research question.

4.3.3 Established Discounting. Holloway et al. (2024) establishes the presence of both overall discounting behavior and time-varying discounting behavior in the data. I summarize the findings here before addressing the heterogeneity within these results. Holloway et al. (2024) focuses on the balanced sample - the set of subjects that completed every survey. My summary reports the full sample results with this balanced sample in parentheses. Define d_{ijw} as subject i's implied discount factor in price list j in survey week w. This is equal to the midpoint of a subject's switch interval divided by 242 (the fixed ticket allocation to the later lottery).[6] For example, if a subject prefers 242 tickets in the later draw to 220 tickets in the sooner draw, but 220 tickets in the sooner draw to 242 tickets in the later draw, then $d = \frac{219}{242} = 0.9050$.

Across every survey and price list, the average d is 0.9566 (0.9560) - corresponding to a switch point between 230 and 232 tickets. This is significantly different from one ($p < 0.0001$),

[6] When a subject never switches –always preferring the later draw– a switch midpoint of 245 is assumed. When a subject switches immediately –always preferring the sooner draw– a switch midpoint of 201 is assumed.

consistent with positive discounting over the tickets in the study.[7] For each week the delay increases, d decreases by 0.06pp (0.07pp). This is significantly different from zero ($p = 0.0002$), demonstrating meaningful response in discounting given a longer choice delay. [8] In terms of time-variance, moving a decision forward in time by a week resulted in a reduction in the discount factor of 0.13pp per week (0.22pp per week), once again statistically significant ($p = 0.0005$). Figure 11 in Chapter II presents results for survey-week effects as dummies for both the full and balanced sample. Regardless of sample chosen, a clear downward trend in discount factor is present in the aggregate.

The aggregate trend of decreasing patience obscures the composition of discounting present at the subject level. Holloway et al. (2024) leverages the subject-week variation in the experimental design to speak to the subject level effects of survey week. Time-invariance is rejected for 55% of subjects, with 30% and 25% exhibiting decreasing and increasing patience, respectively. Additionally, conditioning on exhibiting time variance, the magnitude of decreasing patience is larger on average than the magnitude of increasing patience, lending to the overall trends shown in Figure 11.

4.4 Who Has Time-Varying Preferences?

Our questionnaires elicited two types of data: (1) time-invariant, “static” data and (2) time-varying, “dynamic” data. Static data collected at only one point in time in the study. I interpret static data as subject characteristics at the point in time they are collected. Dynamic data are collected in every survey, and so can be interpreted as describing the state that a subject is in at the time of the survey. In this section, I analyze the static data to determine which characteristics of subjects are correlated with time-varying discounting.

4.4.1 Demographics and Socioeconomic Status. As part of the screening process [9] and exit survey (Table 11), we collected information on subjects’ demographics and socioeconomic status which have been shown in past studies to be correlated with time preferences (e.g. (Castillo, Ferraro, Jordan, & Petrie, 2011)). To estimate their associations with discounting as well as time-variance in discounting, I estimate the following model:

$$d_{ijw} = \alpha + \lambda_0 \cdot w + \lambda_1 \cdot X_i + \lambda_2 \cdot X_i \cdot w + \lambda_3 \cdot ((t_j^L - t_j^S) - 8) + \lambda_4 \cdot (1 - \mathbb{1}(t_j^S = \tau_j)) + \eta_{ijw} \quad (4.2)$$

[7] Two-tailed t-test.

[8] OLS regression of d on the delay and an indicator variable for an immediate sooner draw, standard errors clustered at the subject level.

[9] Relevant data from the screening process includes gender and sources of college tuition payment.

where w denotes survey week, $((t_j^L - t_j^S) - 8)$ denotes the choice delay between the sooner and later options in list j (adjusted so that α measures the eight-week discount factor), $\mathbb{1}(t_j^S = \tau)$ is an indicator for whether there is no front-end delay prior to the sooner option in choice j (adjusted so that α measures the discount factor without a front-end delay), and X_i denotes the demographic or socioeconomic characteristic of interest. The λ_1 coefficient is an estimate of the characteristic's association with time-invariant discounting. Figure 17 presents the coefficients and their 95% confidence intervals on each of the characteristics of interest.

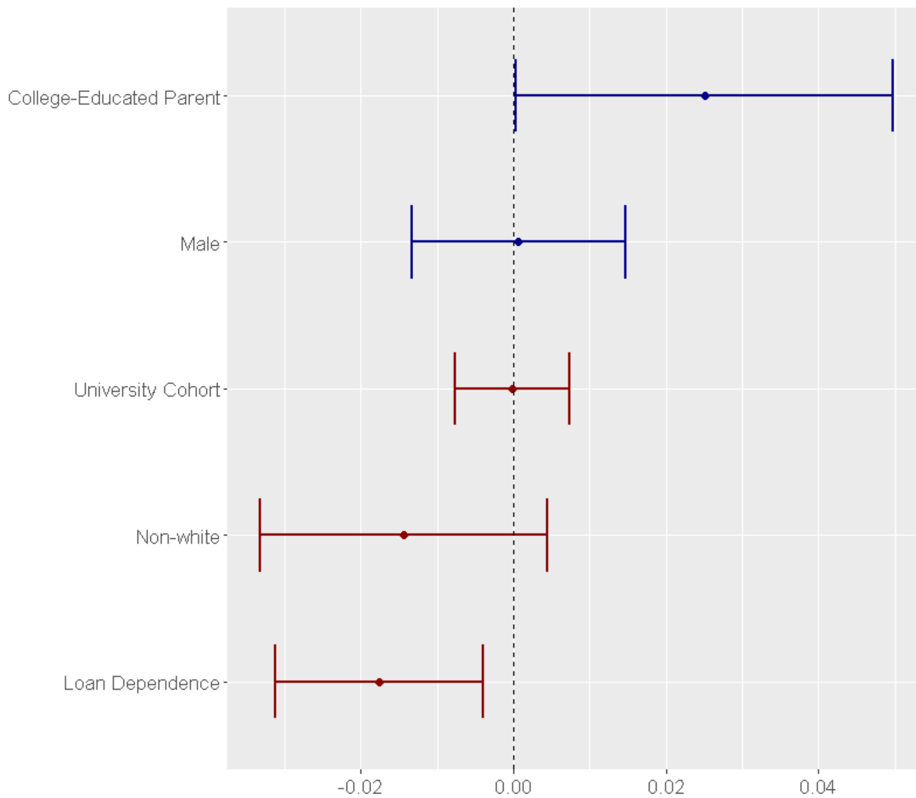

Figure 17. Coefficient Plot: Demographic and Socioeconomic Variables and the Average Discount Factor

Notes: the plot shows the coefficients and 95% confidence intervals of the λ_1 coefficients regression models of Equation 4.2. Standard errors are clustered at the subject level. The data consist of 6,367 price lists from 153 subjects for Male and Loan Dependence and 5,175 price lists from 97 subjects for University Cohort, Non-white, and College-Educated Parent.

We find that, all else equal, subjects whose parents hold any college degree exhibit marginally higher discount factors. The points estimate is 0.0250 ($p = 0.0471$), meaning that subjects with college-educated parents applied a discount factor of 0.9731 (annual discount rate of 19%) to their 8 week delay choice with no front end delay in week 0 on average, compared to 0.9481 (annual discount rate of 41%) for subjects that do not have college educated parents.[10]

[10]We will continue to use this discount factor for subsequent comparisons.

I also find that subjects who rely on student loans to pay for college exhibit significantly lower discount factors than those who do not. With a point estimate of -0.0176 ($p = 0.0117$), those relying on student loans applied a discount factor of 0.9532 on average (annual discount rate of 37%), while those that do not rely on student loans had an average discount factor of 0.9708 (annual discount rate of 21%). I take these results as suggestive evidence of correlation between socioeconomic status and discounting in the sample. To the extent that this informs liquidity-constraints, this result aligns with other research in the literature (Harrison, Lau, and Williams (2002),Matousek, Havranek, and Irsova (2022)).

The λ_2 coefficient on the interaction between the characteristic of interest and survey week gives an estimate of the variable's association with changes in discounting over time. Figure 18 presents the coefficients and their 95% confidence intervals for each of the characteristics of interest. As a point of comparison, the average weekly change in the eight-week discount factor for all subjects is -0.0013 (-0.0026 for the two-week increments between surveys). Non-zero coefficients on interaction terms are estimates of systematic deviations from this average. I find no statistically significant associations of the demographic and socioeconomic characteristics with time variance in discounting.

4.4.2 Exit Survey - Non-demographic. We continue to use Equation 4.2 to estimate the associations between each non-demographic variable of interest and both overall discounting and changes in discounting over time. The exit survey focused on variables relating to self-perception of patience and myopia, as well as behavioral choices linked with patience and myopia. The questions are in Table 11 and Table 13 shows exactly how responses are measured.

Figure 19 presents the coefficients and their 95% confidence intervals for each of the variables of interest. Subjects that perceive their patience as unstable over time exhibit higher discount factors, all else equal. A single point increase in the four-point scale of stability is associated with a reduction of the average eight-week discount factor of -0.0130 ($p = 0.0243$). The standard deviation across all responses is 0.6713, resulting in an reduction in the discount factor of 0.0087 from 0.9923 to 0.9836. Nicotine use is also correlated with a reduction of discount factor of -0.0237 ($p = 0.0141$) from 0.9744 (18% annual discount rate) to 0.9507 (39% annual discount rate).

Figure 20 presents the coefficients and their 95% confidence intervals for each of the variables of interest. The only statistically significant evidence of systematically time-varying preferences comes from those with self-perception of high instability of patience. While I estimate a decrease in patience over time on average, I find that this does not hold for subjects who –

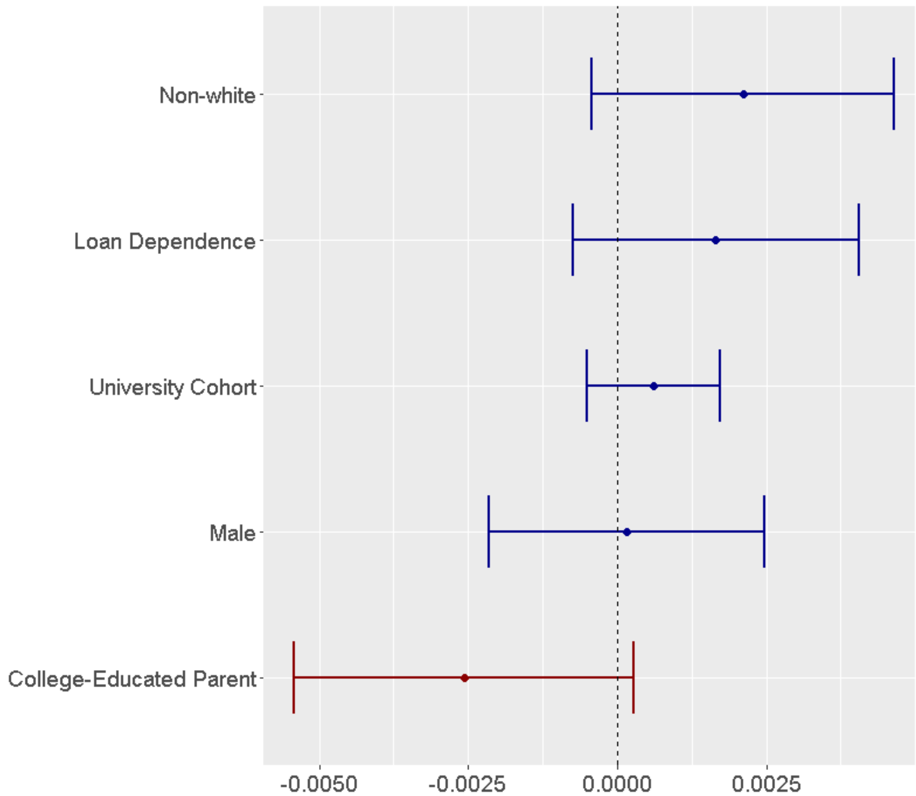

Figure 18. Coefficient Plot: Demographic and Socioeconomic Variables and Discount Factor Time-variance
Notes: the plot shows the coefficients and 95% confidence intervals on the interaction between the variables of interest and survey week from a regression model of Equation 4.2. Standard errors are clustered at the subject level. The data consist of 6,367 price lists from 153 subjects for Male and Loan Dependence and 5,175 price lists from 97 subjects for University Cohort, Non-white, and College-Educated Parent.

accurately– believe their patience is stable over time. The positive λ_2 coefficient of 0.0021 ($p =$ 0.0075), when compared to the effect of an increase in survey week (λ_0) of -0.0058 (p ¡ 0.0001), offsets most of the time variance I observe when moving from one end of the scale to the other.

Apart from a focus on average discounting or time-variance I do find that, consistent with the literature, those that use nicotine products and regret how often they use them do exhibit present-bias. This is demonstrated by an increased discount factor in the presence of a front-end delay (equivalently a lower discount factor in the presence of an immediate decision). In fact, regret of nicotine use becomes the only significant predictor of discount factor within drawing week. While not the focus of this study, this result serves as validation for the measurement methods. Results for the referenced regression in Appendix Table A11.

Table 13. Exit Survey Non-demographic Question Responses

Variable	Description
Patience	Response to Patience
Stability of Patience	Response to Patience: Stability - Scale from 0 to 3, **(D)** being 0 and **(A)** being 3
Fall Through	Response to Patience: Falling Through - Scale from 0 to 2, **(C)** being 0 and **(A)** being 2
Procrastination	Response to Patience: Procrastination
Exercise Hours	Response to Exercise
Exercise Regret	Binary variable - Equal to 1 if answer to Exercise: Regret is **(B)** or **(C)**, 0 otherwise
Nicotine Use	Binary variable - Equal to 0 if answer to Nicotine is **(E)**, 1 otherwise
Nicotine Regret	Binary variable - Equal to 1 if answer to Nicotine: Regret is **(A)**, 0 otherwise
Streaming Regret	Response to Streaming - Scale from 0 to 2 (**(A)** = 0, **(B)** or **(C)** = 1, **(D)** or **(E)** = 2)

4.4.3 Rotating Questions. The last group of static data are the rotating questions from each survey. This data is either used directly as the variable of interest in analogous regressions to the prior two sections, or in some cases are transformed and combined to construct a meaningful variable of interest. Variable definitions can be found in Table 14. Coefficients on the variable of interest and its interaction with survey week can be found in Figures 21 and 22, respectively.

We find some evidence of correlation between these variables addressing risk, cognition, and general time preferences with time-varying discounting or discounting overall [11]. Answering the interest rate cognition question correctly resulted in an increase of average discount factor of 0.0228 ($p = 0.0066$) from 0.9562 to 0.979. Inflation perceptions also have marginal impact on discounting. A one percentage point increase in what a subject believes inflation currently is and what inflation will be in one year result in a decrease of average discount factor of 0.0010 ($p = 0.0190$ and $p = 0.0661$). Standard deviations are 6.927 and 5.678, respectively, totalling decreases of 0.0693 and 0.0568 on average for a one standard deviation increase.

Beliefs about current inflation also marginally impact discounting over time, with a point estimate of 0.0002 ($p = 0.0341$), a one standard deviation in response results in a relative increase

[11] While not indicative of correlation with discount factor, subjects in the sample performed on par with the literature on 'fast versus slow' thinking questions with an average of 1.247 out of 3 correct (Frederick, 2005)

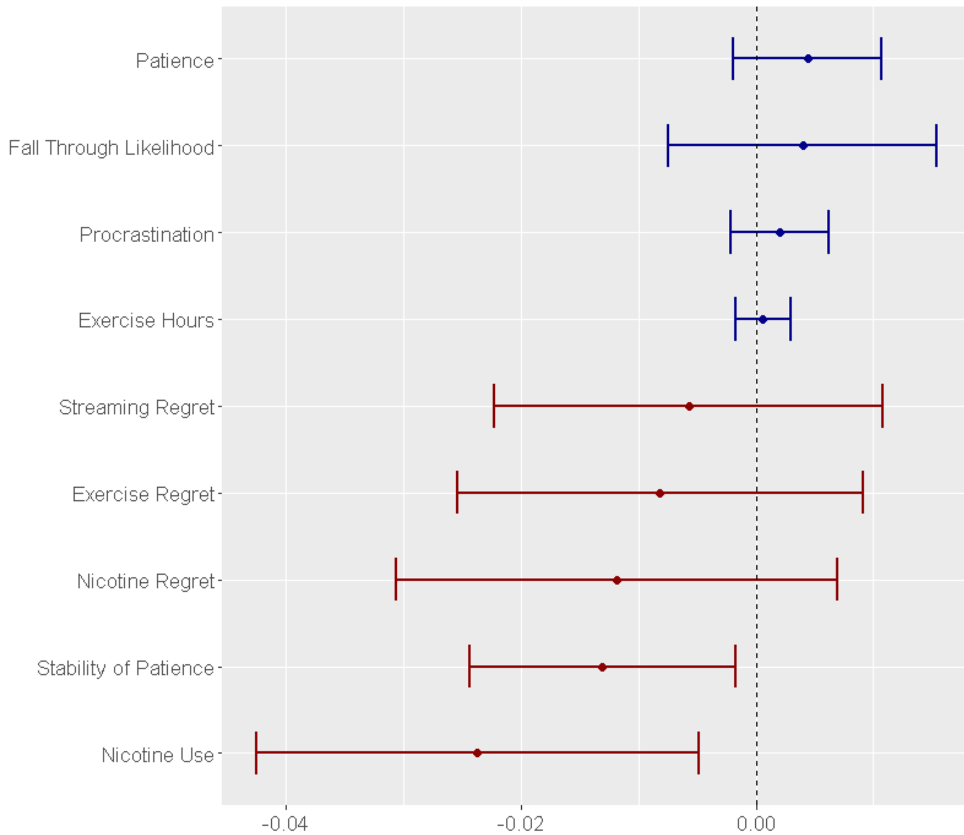

Figure 19. Coefficient Plot: Non-demographics Variables and the Average Discount Factor
Notes: the plot shows the coefficients and 95% confidence intervals on the variables of interest from a regression model of Equation 4.2. Standard errors are clustered at the subject level. The data consist of 5,175 price lists from 97 subjects.

over average of 0.0024 in discount factor by week 12 of the survey. This is once again consistent with less decreasing patience than those with lower perceptions of inflation. This effect is reversed when comparing perception of current inflation to expectation of future inflation. The higher a subject's perception of inflation is today relative to future perceptions, the more impatient they become over the course of the study. The other variable that has impact on discounting over time is the Wait to Redeem question asked in week 4 of the survey. The higher a pay out subjects are willing to wait longer for, the less decreasing patience subjects exhibit relative to the average. An increase of one dollar results in an increase of 0.0003 ($p = 0.0036$), meaning that an increase of one standard deviation (5.226) results an increase of 0.0199 over the average by week 12 of the survey. This result does not hold for the same question shifted forward eight weeks asked in week 12. This is not surprising given the shorter time to the earliest payoff, as the present-biased detailed in Section 4.3 would suggest lower discount factors applied to decisions when evaluation period is pushed forward in time.

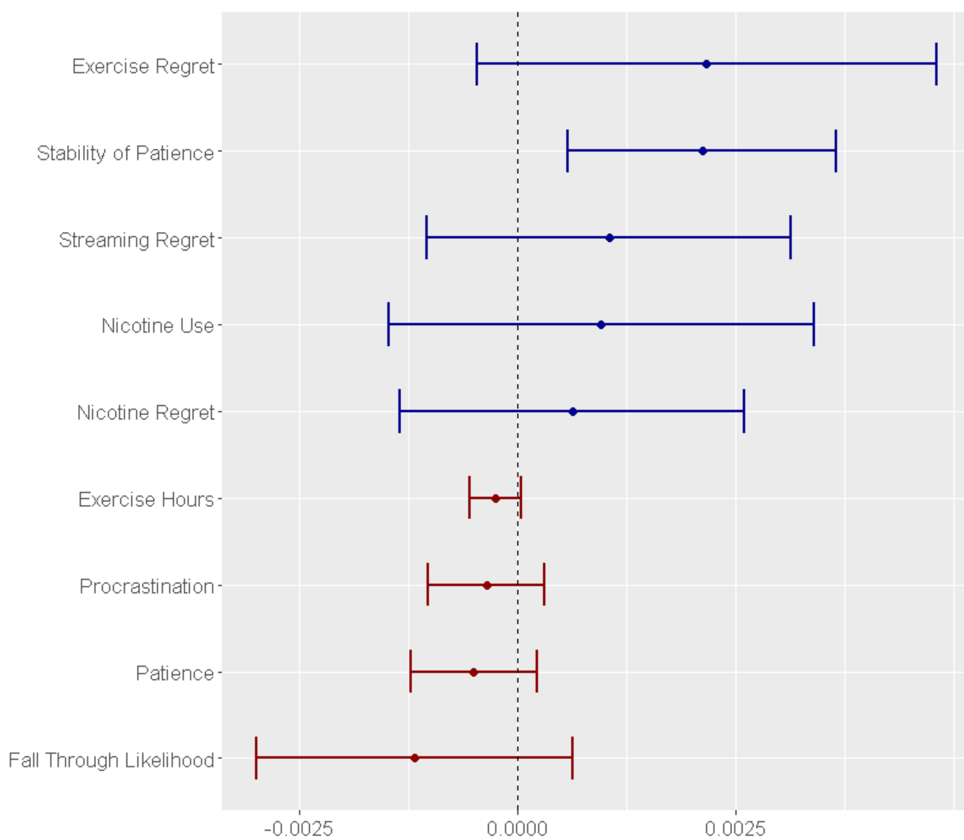

Figure 20. Coefficient Plot: Non-demographics Variables and the Discount Factor Time-variance

Notes: the plot shows the coefficients and 95% confidence intervals on the interaction between the variables of interest and survey week from a regression model of Equation 4.2. Standard errors are clustered at the subject level. The data consist of 5,175 price lists from 97 subjects.

Table 14. Rotating Survey Questions

Variable	Description
TRI: Wk 0	Response to Week 0
TRI: Wk 12	Response to Week 12
TRI: Difference	Change in response to TRI question between Week 0 and Week 12: (TRI: Wk 12) - (TRI: Wk 0)
WTR: Wk 4	Dollar value of response to Week 4 normalized to 50
WTR: Wk 12	Dollar value of response to Week 12 normalized to 50
WTR: Difference	Change in dollar value of response to WTR question between Week 4 and Week 12: (WTR: Wk 12) - (WTR: Wk 0)
Inflation Perception	Response to Week 6a
Inflation Expectation	Response to Week 6b
Inflation Difference	Difference between response to Week 6a and Week 6b: (Week 6b) - (Week 6a)
Expected Utility Violation	Indicator variable equal to 1 if prize draws B and C are chosen in Week 8a and Week 8b, respectively
Interest Rate Cognition	Indicator variable equal to 1 if response to Week 10 is equal to the correct answer when accounting for compounding interest (432)
Interest Rate Miscalculation	Difference between 432 and response to Week 10: 432 - (Week 10)
Cognitive Reflection Score	Sum of correct responses to Cognitive Reflection questions. (0.05, 5, and 47 for parts 1, 2, and 3, respectively)

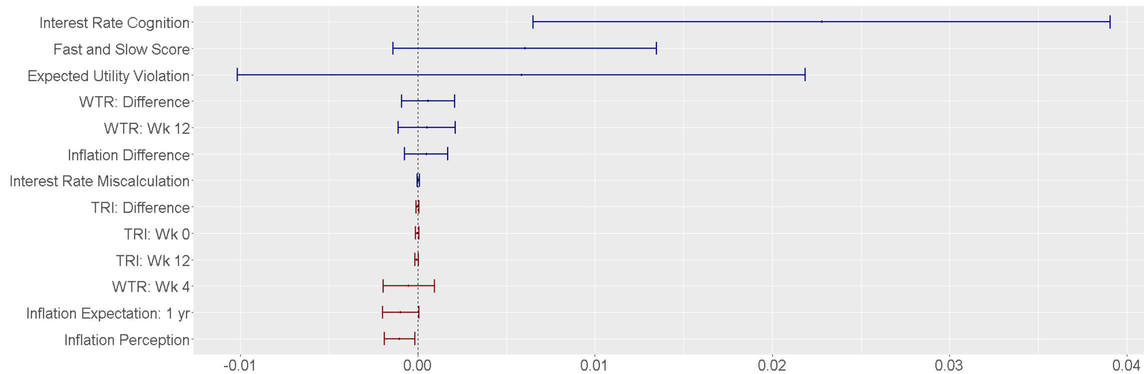

Figure 21. Coefficient Plot: Rotating Covariates' Effects on Average Discount Factor

Notes: the plot shows the coefficients and 95% confidence intervals on the variables of interest from a regression model of Equation 4.2. Standard errors are clustered at the subject level. The data consist of 6,367 price lists from 153 subjects for Week 0 variables, 5,695 price lists from 117 subjects for Week 4 variables, 5,753 price lists from 114 subjects fro Week 6 variables, 5,562 price lists from 106 subjects for Week 8 variables, 5,292 price lists from 99 subjects fro Week 10 variables, and 5,175 price lists from 97 subjects for Week 12 variables.

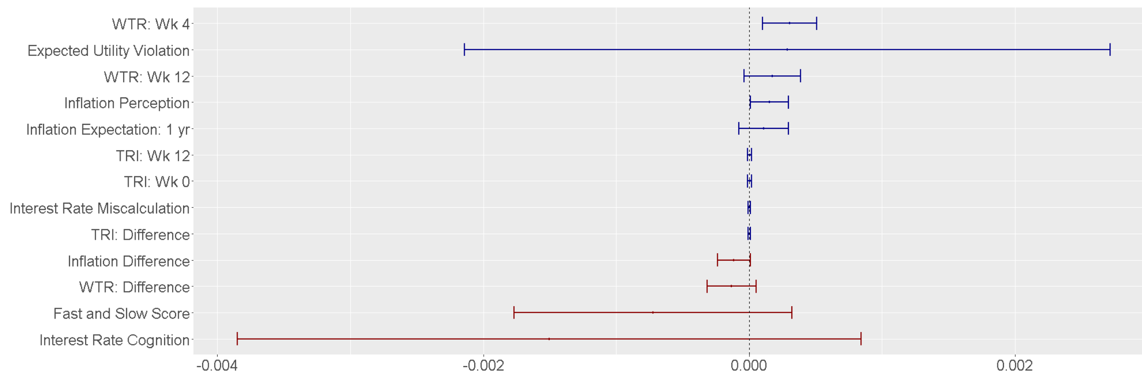

Figure 22. Coefficient Plot: Rotating Covariates' Effects on Discount Factor Over Time

Notes: the plot shows the coefficients and 95% confidence intervals on the interaction between the variables of interest and survey week from a regression model of Equation 4.2. Standard errors are clustered at the subject level. The data consist of 6,367 price lists from 153 subjects for Week 0 variables, 5,695 price lists from 117 subjects for Week 4 variables, 5,753 price lists from 114 subjects fro Week 6 variables, 5,562 price lists from 106 subjects for Week 8 variables, 5,292 price lists from 99 subjects fro Week 10 variables, and 5,175 price lists from 97 subjects for Week 12 variables.

4.5 What Causes Time-Varying Preferences?

In the last section I find limited evidence of innate characteristics driving time-varying behavior. What else might explain the variation I see over time in the data? In this section I look at, instead of characteristics of subjects, the conditions they find themselves in at different points in the study. In particular, the study collected information on the financial and psychological standing of subjects in every survey.

4.5.1 Dynamic Covariates: Aggregate. We take a similar approach to the last section to study heterogeneity on time-varying margins. I regress the discount factor displayed for a given MLL on characteristics of the MLL and a variable of interest, as in Equation 4.2. Relevant variables are described in Table 15. The interaction term with survey week used in Equation 4.2 is omitted. Since dynamic covariates change across time, their interaction with survey week can not be readily interpreted compared to some baseline as can be done for the stationary variables in the previous section. Instead, I first determine the effect these variables have on overall discounting and then determine if these variables themselves have time trends.

Coefficients on the variables of interest are found in Figure 23. I find evidence of positive correlation between one's subjective financial security and one's discount factor, on average. A point estimate of 0.0042 ($p = 0.0141$) results in one standard deviation (1.4743) increasing discount factor by 0.0062. To reinforce this result, I find two other measures of financial well-being positively impact discount factor. Self-perceived ability to cover an unexpected $500 expense is consistent with an increase of discount factor by 0.0094 ($p = 0.0071$) and having addressed a routine cost in the past two weeks is consistent with an increase of 0.0103 ($p = 0.0353$). As previously noted, these results align with findings that liquidity constraints are correlated with more impatient decisions (Harrison et al. (2002),Matousek et al. (2022)). Those that have the ability to cover not only unexpected costs, but also their routine costs have financial security and so do not act as impatiently when faced with a potential payoff. I also find that, while negatively correlated with financial situation, stress levels are not meaningfully correlated with discount factor. This suggests that if negative psychological conditions do have an effect on discounting, it may either be a noisy process or mostly driven by financial security's effect on one's psychological condition.

None of the dynamic variables tested meaningfully change the coefficient on survey week. All regressions run capture -0.0013 in their 95% confidence interval, making them insignificantly different from the effect found in the simple regression. While these variables do not soak up variation from the week coefficient, if the variables I find to be correlated with discount factor are

Table 15. Dynamic Variables

Variable	Description
Paycheck	Binary variable - Equal to 1 if Response to Financial Events included **(A)**
Other Income	Binary variable - Equal to 1 if Response to Financial Events included **(B)**
Other Future Income	Binary variable -Equal to 1 if Response to Financial Events included **(C)**
Routine Costs	Binary variable - Equal to 1 if Response to Financial Events included **(D)**
Unexpected Coverable Cost	Binary variable - Equal to 1 if Response to Financial Events included **(E)**
Unexpected Uncoverable Cost	Binary variable - Equal to 1 if Response to Financial Events included **(F)**
Future Costs	Binary variable - Equal to 1 if response to Financial Events included **(G)**
Financial Stability	Response to Financial Stability
Food Security Scale	Response to Food Security: Type/Amount - Scale from 1 to 4 **(D)** being 1 and **(A)** being 4
Food Secure	Binary variable - Equal to 1 if answer to Food Security Scale is **(A)** or **(B)**, 0 otherwise
Food Security Stress	Binary variable - Equal to 1 if response to Food Security: Binary is **(A)**, 0 otherwise
$500 Expense	Response to $500 Expense - Scale from 1 to 3, **(C)** being 1 and **(A)** being 3
Stress	Response to Stress
Log(Work/School Hours)	Natural log of response to Work/School Hours
70 Work/School Hours	Binary variable - Equal to 1 if response to Work/School Hours is greater than 70, 0 otherwise

also correlated with survey week, then the time-varying behavior I find could still be a result of this time trend. Figure 24 shows the time trends of Financial Security, Ability to Cover a $500 Expense, and Routine Costs for the full sample of the survey. While survey-to-survey changes are small, all three variables change statistically significantly over time. An additional two weeks is consistent with an increase of financial security of 0.0552 ($p = 0.0182$) on a 7-point scale. An additional two weeks is also consistent with an increase in the confidence one can cover a $500 expense of 0.0282 ($p = 0.0126$) on a 3-point scale. The probability of having covered a routine cost in the past two weeks decreases over time, with an additional two weeks reducing the probability by 0.0219 ($p = 0.0010$).

Only the Routine Cost result is consistent with the decreasing patience found in aggregate. Having covered a routine cost in the past two weeks results in more patient decisions,

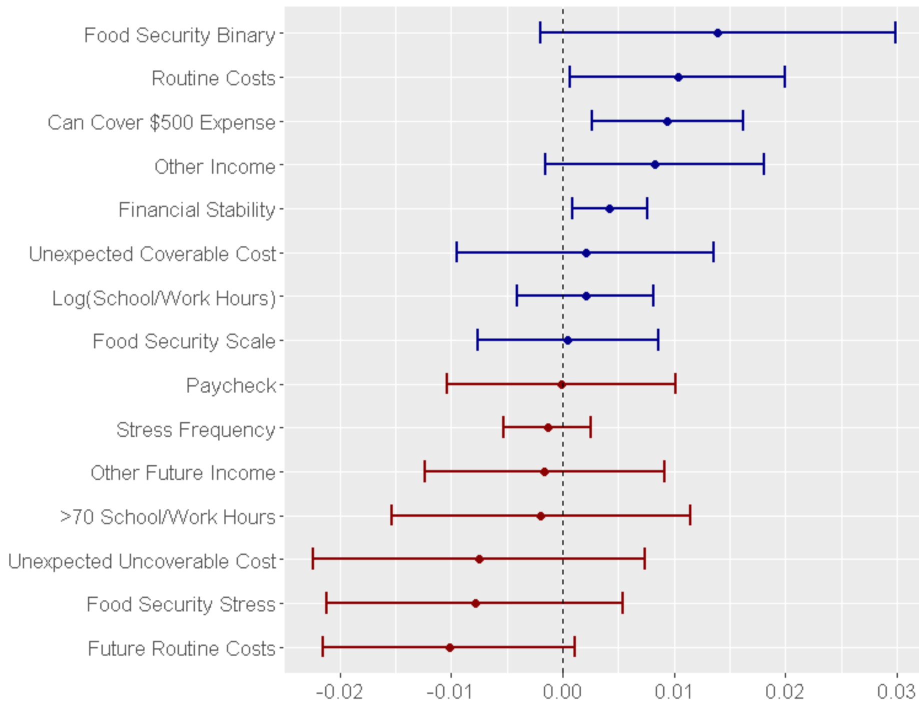

Figure 23. Coefficient Plot: Dynamic Covariates' Effects on Average Discount Factor
Notes: the plot shows the coefficients and 95% confidence intervals on the variables of interest and survey week from a regression model of Equation 4.2. Standard errors are clustered at the subject level. The data consist of 6,367 price lists from 153 subjects.

and it is less likely on average for subjects to have done so in later weeks of the survey. The results for Financial Security and $500 Expense are instead consistent with increasing patience. This suggests that what effect these two variables have on discounting may be overwhelmed by the effect Routine Costs, or some other factor, have on discounting.

The relative impact of these variables is reinforced when looking at trends for the balanced sample. In Section 4.3.3 I show that the trend of decreasing patience is effectively the same between the balanced and full sample. When testing the effect of additional weeks on the variables of interest in the balanced sample, I find no significant change in either Financial Security or $500 Expense. However, I do find a statistically significant effect of -0.0182(p =0.0009) of an additional two weeks on probability of having covered a routine cost[12].

[12] Analog to Figure 24 is Appendix Figure B12

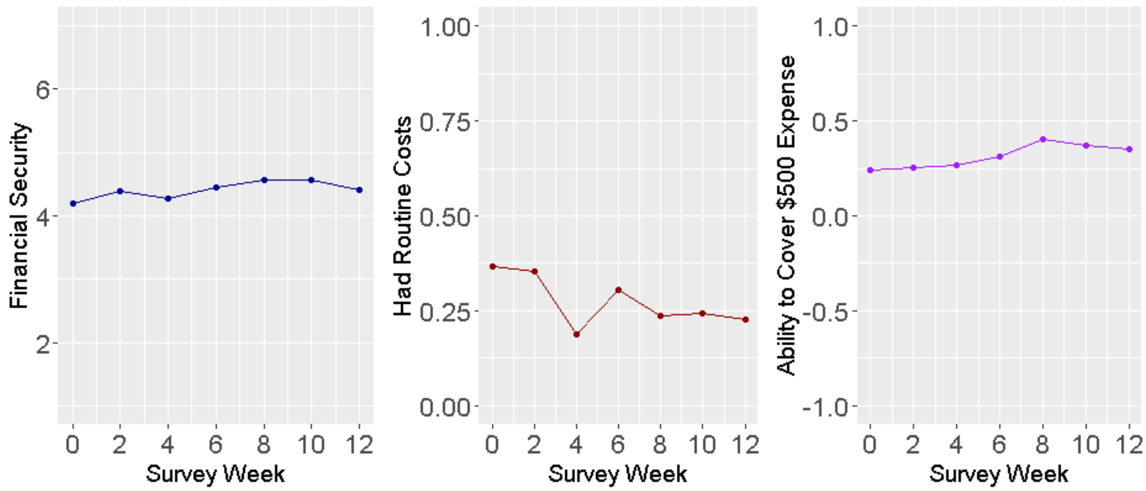

Figure 24. Time Trends of Significant Dynamic Variables

4.5.2 Dynamic Covariates: Subject-level. We now turn to analysis made possible by the novel longitudinal nature of the study. I have access to not only dynamic data on financial and psychological factors within subject, but dynamic data on the discounting behavior within those same subjects. This allows us to analyze the effect of the dynamic covariates within subject. To do this, I modify Equation 4.2 to include subject fixed effects and the variable of interest, with no survey-week interaction. The coefficient on the dynamic variable of interest in this regression is the effect of a change in the variable, relative to the average for that subject, all else equal. These coefficients are plotted in Figure 25.

The results of these regressions are consistently null. I find no evidence of within subject variation dependent on the margins tested. I take this as suggestive evidence that subjects are relatively invariant to changes in their outside circumstances, either due to not considering these changes when making decisions in the context of the experiment, inconsistent directional responses across individual, or insufficient variation in outside circumstances within subjects to generate precise estimates.

To compare to aggregate results, I also test for correlation between survey week and the variables found to be correlated with aggregate discounting while accounting for subject fixed effects. As in the balanced sample, I find null results for both Financial Security and $500 Expense, but a significant decrease in probability of having covered routine costs of about 0.0204 ($p = 0.0010$) per survey on average. This finding cements the downward trend of routine costs faced by the sample, however, given that the effect of Routine Costs is null at the subject-level, does not help explain differential changes of discount factors across subjects.

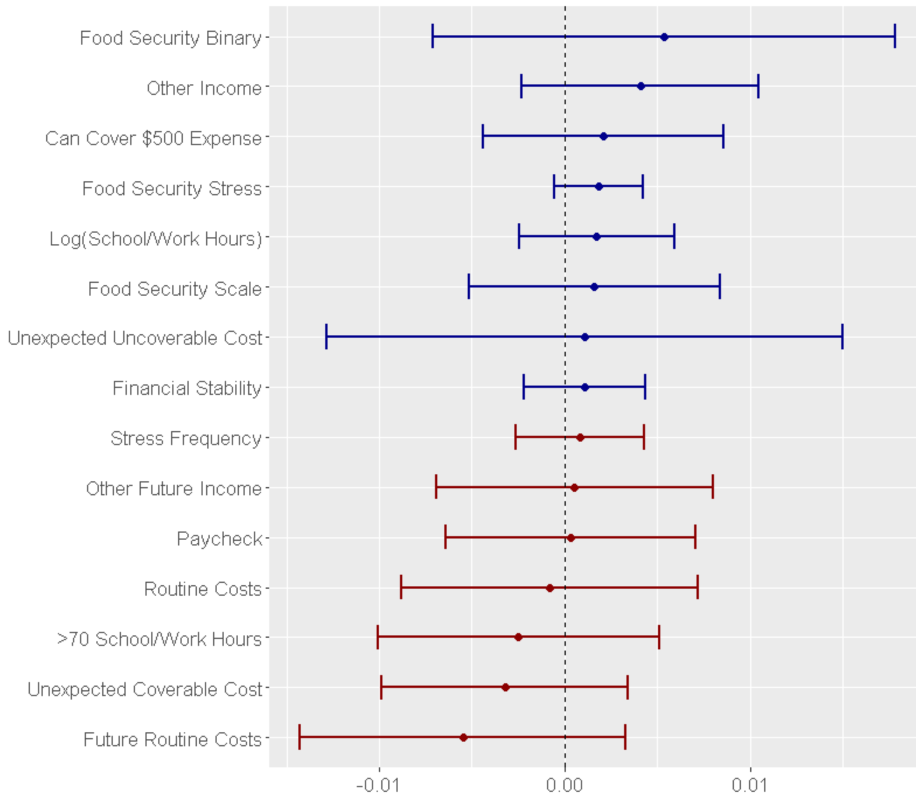

Figure 25. Coefficient Plot: Dynamic Covariates' Effects on Subject Discount Factor
Notes: the plot shows the coefficients and 95% confidence intervals on the variables of interest and survey week from a regression model of Equation 4.2 including subject fixed effects. Standard errors are clustered at the subject level. The data consist of 6,367 price lists from 153 subjects.

4.5.3 Drawing Effects. While subjects do not seem to respond to outside changes in their circumstances, they may however respond to changes in the context of the decisions they make in the study. I now analyze the effect on discounting of subjects receiving a drawing. Recall that each survey taken results in a drawing for the subject. However, this drawing may take place immediately or as far as 32 weeks in the future. This creates variation within week across subjects in terms of whether they took part in a drawing in the last survey if an immediate drawing occurred or prior to decision-making in a survey if a 2-, 4- or 8-week delayed drawing occurred, as well as how many drawings a subject has taken part in.

First, I estimate the effect of having received a post-survey drawing in the most recent survey. A drawing occurs in the most recent survey if, for whatever delay length was randomly chosen, the subject switched towards preferring the earlier, lower odds drawing at a higher number of tickets than was randomly selected. This decision is also dependent on whether the most recent survey's 2-week front-end delay was randomly chosen to override the present survey's decisions. I again regress on delay-length, presence of front-end delay, and week survey along with an indicator for if a post-survey drawing occurred in the most recent survey. I also control for the delay length randomly selected in the most recent survey, as well as the switch point chosen in the multiple lottery list for that delay length. This ensures that the effect of a post-survey drawing across subjects in this regression comes from comparing those that discounted the same in the decision that randomly generated a post-survey drawing for some but not for others. Results can be found in the first column of Table 16. I do not find evidence that experiencing a post-survey drawing impacts decision-making in the next survey. Instead, I find that the choice made in the most recent survey for the analogous decision is very important in explaining this survey's choice. This result is expected, as it suggests that, whether subjects' choices are varying over time or not, the way they have made the choice in the past heavily influences how the decision is made in the present.

Considering the effects of a pre-survey drawing and the number of drawings a subject has experienced requires a different approach than for most recent survey post-survey drawings. A pre-survey drawing occurs if 2, 4, or 8 weeks ago the delay length randomly chosen reflected a later drawing in the present survey week and if the subject's switch point in the multiple lottery list corresponding to the randomly chosen delay length reflects a preference for a later, higher odds lottery than the number of tickets randomly selected. This means that whether a pre-survey drawing occurs is correlated with discounting behavior in past surveys, which is clearly correlated with discounting behavior in the present survey. This correlation results in endogeneity when considering either whether a pre-survey drawing occurred, and by extension how many drawings a subject has experienced.

To combat this endogeneity I adopt a 2-stage least squares approach. The reduced form regression once again includes delay-length, presence of front-end delay, and survey week as regressors, along with either an indicator for presence of a pre-survey drawing or a variable denoting the number of drawings experienced prior to making decisions in the present survey. In separate regressions, I instrument these variables with indicators for whether the delay length randomly selected 2, 4, or 8 weeks ago, respectively, correspond to the present survey. The results of these regressions are in columns 2 and 3 of Table 16.

These instruments are clearly positively correlated with the variables of interest and satisfy the monotonicity assumption. The likelihood of experiencing a pre-survey drawing is strictly positively correlated with whether these delays were chosen in the past. If the delay was not chosen, the probability of receiving a drawing through that channel is zero. If the delay

Table 16. Two-Stage Least Squares: Drawing Participation

Model:	Post-Survey Drawing	Pre-Survey Drawing	# of Drawings
	(1)	(2)	(3)
Post-survey Drawing in $w-2$	-0.0054		
	(0.0059)		
Random Delay in $w-2$	-0.0002		
	(0.0003)		
Discount Factor for Random Delay in $w-2$	0.3969***		
	(0.0647)		
Pre-Survey Drawing in $w^{\dagger}$		-0.0027	
		(0.0048)	
# of Drawings before w responses†			-0.0026
			(0.0050)
First Stage			
2 week delay in $w-2$		0.5020***	0.7356***
		(0.0394)	(0.0621)
4 week delay in $w-4$		0.4881***	0.4093***
		(0.0473)	(0.0641)
8 week delay in $w-8$		0.6148***	0.0663
		(0.0618)	(0.1071)
F-Stat		1765	420
N	4,612	4,457	4,357

Notes: $^{***}p < 0.01$ standard errors are clustered at the subject level. † Variables are instrumented with first stage variables. All regressions restricted to survey weeks 2 through 12. Subjects that won drawings at any point in the survey (3 in total) are omitted.

was chosen, then the probability of receiving a drawing is positive and increasing in switch point for each one of the indicators. These instruments also plausibly satisfy the exclusion restriction. Whether a 2, 4, or 8, week delay was chosen in a previous survey should not affect how a subject discounts in the present, except through the impact of the outcome of a drawing that occurs because that delay was chosen. The strength of these instruments is corroborated by high F-statistics for both pre-survey drawing and number of drawings experienced (1765 and 420, respectively). The instruments are further validated by pre-survey drawing's higher F-statistic, as delays randomly selected 2, 4, or 8 weeks ago directly explain whether a pre-survey drawing occurred in the present, while they only explain how many surveys have been experienced through the channel of whether a pre-survey drawing occurred in the present.

4.6 Robustness

The analysis of static and dynamic variables in the previous sections rely on significance testing of coefficients for 42 variables (and their interaction with week for the static variables). Evaluating this many hypotheses opens up risk of false discovery of significance for any one of the variables found meaningful. One method of addressing multiple hypothesis testing is the use of transformation of the P-value called the Q-value (Storey, 2011). The Q-value effectively multiplies the P-value by the number of hypotheses being tested. Q-values are known to become quite conservative as the hypothesis set grows larger. Indeed, if Q-values were to be used as the measure of significance in this study, effectively all heterogeneity tests would be rendered null due to the number of tests being performed.

Instead of using overly conservative significance thresholds, I employ a Bayesian Model Averaging (BMA) method (Wasserman, 2000) of qualitatively addressing the multiple hypotheses tested in this study. BMA approaches take all variables being tested for significance in an OLS regression and sample all possible model combinations to construct a weighted average over each of them. This results in a Posterior Model Probability (due to Bayes' Theorem) for each model sampled. Summing these Posterior Model Probabilities (PMPs) for each model a variable is included in gives the Posterior Inclusion Probability(PIP), which can be interpreted as a qualitative measure of the variable's importance in explaining variation in the dependent variable. Figure 26 presents the PIPs for all variables regressed on in both static and dynamic tests, including interaction terms with survey week.

144 models account for 100% of possible model probability according to Markov Chain Monte Carlo simulation estimates. The 106 models with the highest PMPs, cumulatively accounting for 90% of model probability, are presented in the figure. Variables with PIP of at least 0.5 are presented. Blue bars indicate the coefficient is positive, while red bars indicate a negative coefficient. The X-axis describes the cumulative probability of the models implied by the bars up to that point, with variables being included subsequently along the Y-axis. Starting from the top of the graph, all variables through patience interacted with survey week (patience#week) are included in 100 percent of the 106 models. For those variables with 1 PIP, the graph does not delineate their relative importance via their ordering.

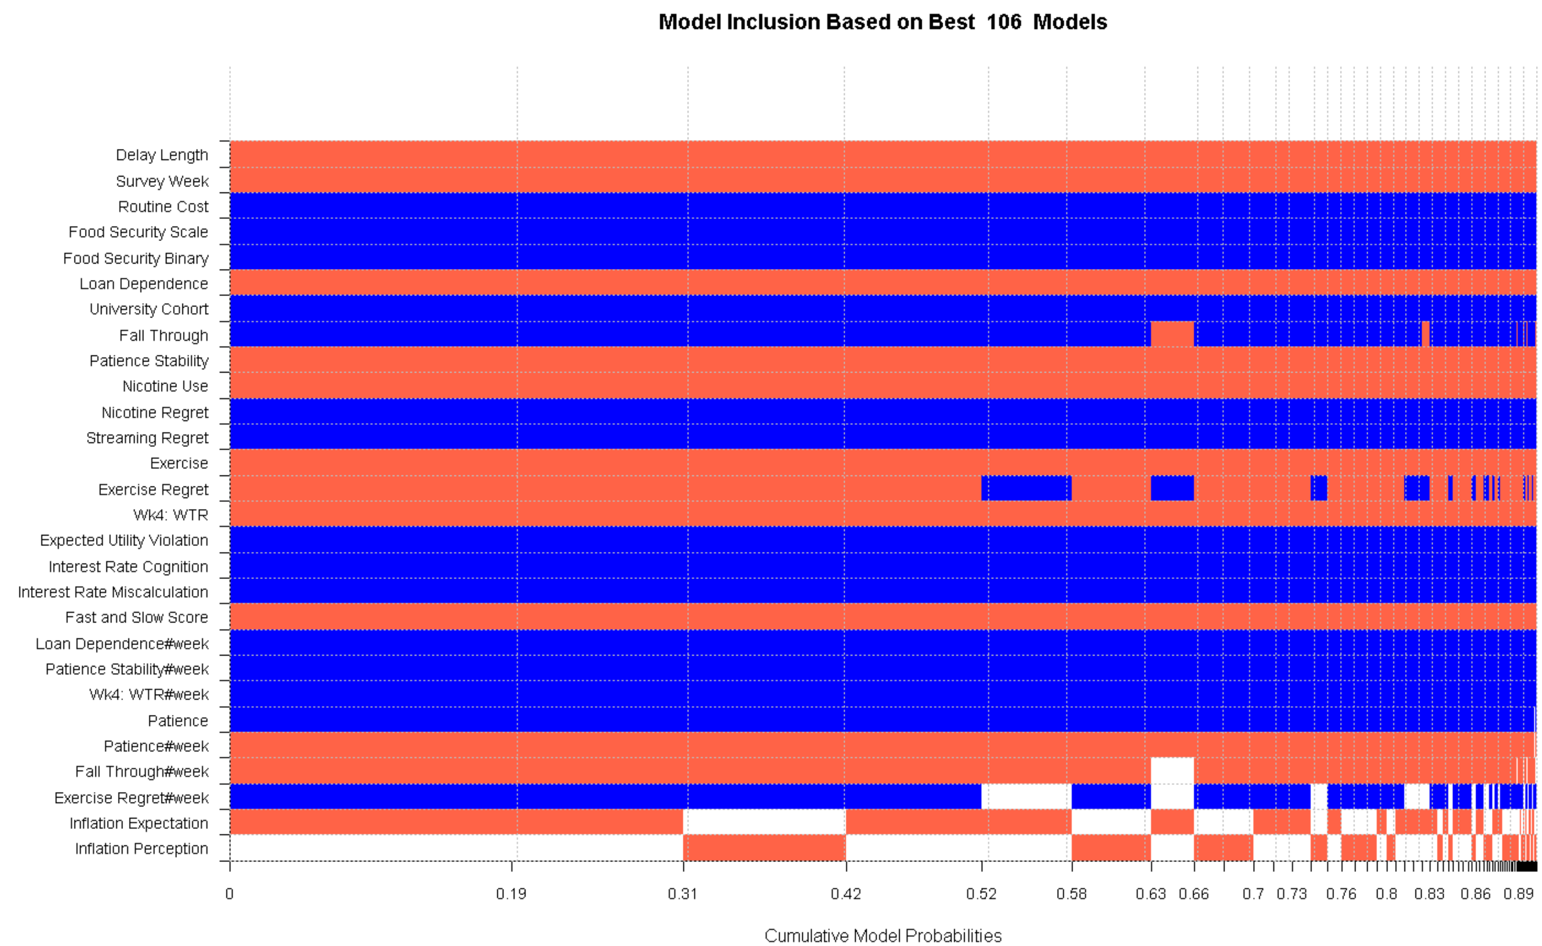

Figure 26. Bayesian Model Averaging Plot

Notes: the plot shows the composition of the 106 models with the highest PMPs that describe 90% of cumulative model probability. The X-axis lists variables with PIPs of at least 0.5 in the 106 models and the Y-axis indicates a model's marginal probability. Blue (red) bars indicate positive (negative) coefficients and white space indicates non-inclusion in a given model. Interaction terms denoted with "#week"

One reassuring result in terms of robustness is the PIP of 1 for survey week. This suggests that given all possible combinations of variables, the importance of survey week on discounting is not superseded by some combination of other variables. The same holds true for back-end delay length. In contrast, front-end delay has a PIP of 0.05, suggesting that the impact of front-end delay[13] *can* generally be explained by other variables without including front-end delay in the model. Another reassuring result is the PIP of 1 for food security variables, Routine Cost, and Loan Dependence. This reinforces the relative importance of measures of socioeconomic status in explaining discounting behavior on average. Notably, measures like $500 Expense, Financial Security, and College-Educated Parent do not have PIPs of at least 0.5 (all < 0.1), suggesting that impact of socioeconomic status can be better summarized through the use of other variables. In a similar vein, the PIP of 1 for specifically Routine Cost and the food security variables underlines the importance of dynamic variables in this analysis. Their inclusion in all models is indicative of external events impacting discounting, even if they do not impart a statistically meaningful trend over time.

Interaction terms' meaningful impact are still relatively rare under this approach. As in the basic OLS approach, Stability of Patience's and Week 4's Wait to Redeem interactions are meaningful and positive in 100% of models. This suggests that the interactions are indeed helpful in explaining variation in discounting on their own, as suggested by the basic OLS approach. However, Loan Dependence and Patience also all have significant interaction terms (PIPs of 1 and 0.9667, respectively). In both cases, the coefficient is negative, meaning the downward trend in discounting could be driven by loan-dependent or self-perceived highly patient individuals.

This same pattern extends to many un-interacted variables that are included in 100% of models, but do not find significance on their own under the basic OLS approach (for instance, Streaming Regret and Fast and Slow Score). This highlights a potential complex interdependence of some variables collected, but does not necessarily detract from the reinforced importance of variables found to be meaningful under the basic OLS approach. Overall, BMA analysis supports the conclusion that survey week's importance cannot be explained by other variables and that heterogeneity in the sample does not aid meaningfully in interpreting the time trend found in the data.

[13] Coefficients on front-end delay in all models relate to *within* survey-week variation. All models displayed include the survey-week variable, resulting in comparisons across front-end delay that hold survey-week fixed. While not the focus of this study, this suggests that non-stationarity of decisions can be explained by other variables, but does not inform whether inconsistency of decisions can be explained by other variables (Holloway et al., 2024).

4.7 Time-Variance as an Asymptotic Process

Reduced-form analysis finds very little correlation between time-variance and observed variables historically found to impact other time discounting characteristics. While the mechanism for what causes time-variance may still be unclear, developing models that can account for the observed behavior is still beneficial in predicting behavior and may also uncover patterns that are suggestive of some untested mechanism. I use the theory outlined in Section 4.2 to develop a discount function that can account for time-variance. First, Consider the following time-invariant discount function:

$$D(\tau,\, t) = \delta^{(t-\tau)^{\eta}} \qquad \eta > 0 \ . \tag{4.3}$$

This discount function is a discrete time reformulation of the Weibull discount function, a behavioral model that satisfies the same time preference properties as the $\beta - \delta$ and hyperbolic discount functions. This function is time-invariant, so long as δ and η are constant across evaluation period.[14] However, suppose instead η can be decomposed such that $\eta = \alpha + \beta * f(\tau)$ where α is constrained to be greater than zero. The resulting discount function is now no longer a function of purely $(t-\tau)$, as τ enters separately as part of η, and so cannot satisfy the definition of time-invariance.

Let $f(\tau) = \frac{1}{1+\tau}$. The resulting time-varying discount function is written:

$$D(\tau,\, t) = \delta^{(t-\tau)^{\alpha + \frac{\beta}{1+\tau}}} \qquad \alpha > 0 \ . \tag{4.4}$$

Note that as τ approaches infinity, $\frac{\beta}{1+\tau}$ approaches zero. Therefore, in the limit, this discount function behaves as a time-invariant function where α can be interpreted as an agent's 'true' η parameter. This model describes a change in discount factor across evaluation period as an agent asymptotically approaching a time-invariant function through some internal process that changes their η parameter over time. β can then be interpreted as a 'variance' parameter. The larger the magnitude of the original β, whether positive or negative, the farther away from true η the agent starts, and the longer it will take for the function to converge to time-invariance. A

[14] Note that the function is equivalent to the standard exponential discount function when $\eta = 1$.

positive β is consistent with increasing patience over time. The η applied to the function in Week 0 will be higher than the η applied to the function in Week 12, and therefore the distance between t and τ will be increasingly down-weighted over time, resulting in increasingly patient decisions. A negative β is conversely consistent with decreasing patience.

We test the aggregate performance of Equation 4.7 against the same model with a constant η, as well as the standard exponential discount function. Estimation of parameters is done using non-linear least squares and results can be seen in Table 17. Results for both the full sample and only those who completed every survey are shown. It is clear based on Akaike Information Criterion (AIC)[15] that not only does the behavioral model with constant η outperform the standard exponential, but that allowing for time-variance further increases fit. For the full sample, the time-varying model estimates a significant, positive β, consistent with increasing patience and the flexibility allowed from the time-varying parameter results in this model achieving the lowest AIC of the three models. Using the balanced sample of those who completed every survey, in contrast, estimates an insignificantly negative β, consistent with marginal decreasing patience that better aligns with reduced-form results. The increased degree of freedom is not as valuable to the fit of the model in the balanced sample, resulting in effectively the same AIC as the Behavioral Invariant.

The conflicting directions of patience estimated by the full and balanced sample is suggestive of two possibilities. One potential mechanism is that those that left the study early may have been more impatient than those remained for the entire study. This would be an intuitive result, however this is not supported by the attrition analysis in Section 4.3.3 nor does it explain all the variation in patience across the study, as I also find within-subject decreasing patience for some subjects. Another potential mechanism is that those that left the study early were more likely to exhibit increasing patience, leaving a relatively higher composition of subjects with marginally decreasing patience by the end of the study. These are not mutually exclusive possibilities, as subjects that exhibited very impatient early decisions may have become more patient over time in the surveys that they did participate in.

To further investigate the differences in the aggregate estimates, I estimate parameters at the subject-level for each of the three functions previously estimated at the aggregate-level. Table 18 sorts subjects (with sufficient variation to be estimated by all three models) by lowest-AIC, and therefore best-fitting, model. Column 1 reports the percentage of the sample best fit

[15] $AIC = -(Log\ \ Likelihood) + 2k$, where k is number of free parameters.

Table 17. Aggregate Parameter Estimates for Invariant and Asymptotic Models

Model:	Standard		Behavioral Invariant		Time-Varying	
	(1)	(2)	(3)	(4)	(5)	(6)
δ	0.9973	0.9973	0.9778	0.9779	0.9784	0.9778
	(0.0002)	(0.0003)	(0.0020)	(0.0027)	(0.0020)	(0.0026)
α	-	-	0.3811	0.3855	0.3704	0.3864
	(-)	(-)	(0.0257)	(0.0326)	(0.0286)	(0.0370)
β	-	-	-	-	0.1123	-0.0094
	(-)	(-)	(-)	(-)	(0.0503)	(0.0754)
Log-likelihood	-26,484.8	-18,149.6	-26,134.2	-17,921.5	-26,127.6	-17,921.4
AIC	52,971.7	36,301.2	52,272.3	35,846.9	52,261.2	35,848.9
r_0	0.1481	0.1499	0.1065	0.1082	0.1581	0.1046
	(0.0140)	(0.0177)	(0.0102)	(0.0130)	(0.0240)	(0.0260)
r_{12}	$= r_0$	$= r_0$	$= r_0$	$= r_0$	0.1024	0.1085
					(0.0110)	(0.0141)
Sample	Full	Balanced	Full	Balanced	Full	Balanced
H_0: $\delta = 1$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$
H_0: $\alpha = 1$	-	-	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$
H_0: $\beta = 0$	-	-	-	-	$p = 0.0827$	$p = 0.9006$

Notes: standard errors are clustered at the subject level. The Full Sample data consist of 6,367 price lists from 153 subjects.The Balanced Sample data consist of 4,345 price lists from 79 subjects. $r_\tau = D(\tau, \tau + 52)^{-1} - 1$ is a measure of the annual discount rate.

by each model, while Column 2 reports the percentage of the sample that exhibits statistical significance of the relevant parameters for distinguishing the models. I find for 52.85% of the sample, the time-varying model is the best fitting model, although only 32.52% of the sample exhibit parameters that actually suggest significant time-variance. Of those best-fit by the time-varying model, about 35.38%(or 18.70% of the sample) exhibit decreasing patience, down to about 25% when correcting for significance. Once again the flexibility afforded by the time-varying model is beneficial to fit, even for those that ultimately exhibit no significant changes in patience across the study. This is seen in the shift in proportion best-fit by the time-varying model, even though parameters ultimately suggest a time-invariant discount function.

One limitation of the time-varying model is the types of time-variance that it can actually describe. Consider a subject who initially makes present-biased decisions so that when no front-end delay to an earlier choice is present, they discount at a higher rate than in the presence of a front-end delay two weeks prior when predicting their future preference. Assuming time-variance of their decisions enters in the form of trends, as opposed to one-off shocks, time-variance could manifest in four different ways. (1) The subject could become weakly decreasingly patient over

Table 18. Subject-level Classification of Best-fitting Model

Model	(1) Best-fitting	(2) Significance-Corrected
Standard	7.32%	8.76%
Behavioral Invariant	39.84%	57.72 %
Time-Varying	52.85% (18.70%)	32.52% (8.13%)

Notes: Data comes from 123 subjects for which we were able to obtain parameter estimates from all models. Column 2 sorts subjects based on the implied model when accounting for 10% statistical significance. If β has $p < 0.1$, the model implied is the learning model. If β has $p > 0.1$, but α has $p < 0.1$, the model implied is the behavioral invariant. Otherwise the model implied is the standard exponential. Numbers in parentheses denote the percentage of the sample with $\beta < 0$ and $\beta < 0$ with significance, in columns (1) and (2) respectively, consistent with decreasing patience.

time as well as weakly decreasingly present biased.[16] This occurs, for instance if the subject meets the decision with no front-end delay more impatiently over time, but at a slower rate than the analogous decision with a two week front-end delay. (2) The subject could become weakly increasingly patient over time as well as weakly increasingly present biased. This occurs, for instance, if the subject meets the decision with no front-end delay more patiently over time, and at a slower rate than the analogous decision with a two week front-end delay.

The first two options are choice profiles that the time-varying function is capable of describing. However, two equally valid choice profiles are (3) the subject becoming weakly decreasingly patient but weakly increasingly present-biased, and (4) the subject becoming weakly increasingly patient but weakly decreasingly present-biased. While my model cannot describe a subset of time-varying choice profiles due to constraints on how the time-varying exponent can evolve, I still find that the model is beneficial in allowing for time-variance that exists in the study that would otherwise not be accounted for at all under traditional discounting models. To investigate the relative importance of the inclusion of different choice profiles, I estimate parameters of a model similar to Equation 4.7:

$$D(\tau, t) = (\delta + \frac{\beta}{1+\tau})^{(t-\tau)} \ . \tag{4.5}$$

[16]Changing present bias in this context is defined by increasing or decreasing percentage difference between the two-week front-end delay choice and the re-evaluated choice with no front-end delay.

Rather than modifying the exponent on time distance, the variance term is added to δ. This function is restricted even further in choice profiles it can describe. If $\beta < 0$, the model will exhibit increasing patience and decreasing future bias. If $\beta > 0$, the model will exhibit decreasing patience and decreasing present bias. Further, the model asymptotically approaches the standard exponential discount function. Results of the estimation presented in Appendix Table A12 demonstrate that the reduction in possible choice profiles does potentially impact aggregate fit. While estimates for both samples reflect the increasing patience seen in the full sample estimation of the behavioral counterpart, the model fits significantly worse, regardless of sample. As noted previously, the standard exponential model that Equation 4.7 approaches is of worse fit than the behavioral-invariant model, as well, suggesting that approaching a relatively more flexible model in the limit may also be beneficial to aggregate fit. After comparing the two models, I view the model that asymptotes to the behavioral-invariant as a step in the right direction of explaining and modelling time-variance of decisions, whether from increasing or decreasing patience, as an evolution of internal discount function parameters.

4.8 Discussion

As part of a 12-week longitudinal study on intertemporal discounting behavior, I collected static data on subjects such as demographics and socioeconomic status, risk preferences, subjective patience, cognition and other characteristics to test their impact on average discounting and time-variance of discounting behavior. I also collected dynamic data relating to financial and psychological of subjects to determine if changes in circumstance may drive time-varying behavior. Socioeconomic factors have the most significant associations with discounting; having a college educated parent and not having to rely on student loans for college payment both predict higher average discount factors, as does higher self perception of financial security, the ability to cover a \$500 expense, and having recently had to pay routine expenses. All of these variables suggest that financial liquidity could be drivers of discounting behavior, which aligns with previous findings in the literature (Harrison et al. (2002),Matousek et al. (2022)). However, these same factors are not strongly related to time-variance. Financial security and the ability to cover a \$500 expense trend upwards in the subject over the course of the study,while the likelihood of having recently covered routine costs trends downwards. Only the latter result, combined with its effect on discount factor is consistent with the time-variance observed in the aggregate, and the magnitude of this effect is small relative to sample average time trend. No dynamic variable accounts for significant within-subject variation over time in discounting.

I also find that subjects perceptions of their own time-invariance are at least partially reflected in their responses. Those that viewed themselves as having relatively stable levels of patience over time applied lower discount factors to decisions on average, but these discount factors changed less over time than those with self-admittedly less stable patience. This does not extend to the level of discounting; those with higher stated levels of patience do not exhibit significantly more patient decisions.

Given the limited ability of the dynamic variables collected to predict changes in time preference, I instead consider whether there is evidence to support time variance as a "native" feature of the experiment. I leverage variation in number of drawings experienced to determine if the realization of a loss drives discounting behavior systematically. I find no evidence that time-varying behavior is driven by the realization of drawings. Participation in both pre-survey drawings immediately prior to a discounting elicitation and post-survey drawings after the prior survey's discounting elicitation have no effect on subsequent discounting behavior.

Characteristics of subjects, their changing circumstances, and internal experience in the study do little to explain the time-varying behavior I observe. Finally, I consider whether models that consider time-variance as an asymptotic learning process nested within a time-invariant discount function explain the data well. This model fits the data better than traditional models and matches the qualitative pattern of decreasing patience in the full sample, but does not for the balanced sample. While not successful in describing the aggregate trend of time-varying behavior, this model does better explain discounting behavior for between 33 and 53% of subjects in the study. Even though the model does not impart information about the specific mechanism for time-variance in the study, the simplicity of the model allows for the inclusion of time-varying decision makers in a wide array of dynamic models for further research and can aid in increased accuracy of prediction of economic behavior.

Further exploration of time-varying behavior is necessary to determine the driving mechanisms and the most efficient model to describe time-variance. The relative importance of financial variables in the study suggest variation in prize size, studying a more socioeconomically diverse population, or designing a study that leverages financial liquidity, may aid in increasing understanding of their impact. The performance of the asymptotic model indicates an experiment designed to create exogenous variation in number of opportunities to learn about one's discounting behavior could also help to determine if the asymptotic pattern remains present and whether the pattern is driven by experience or some other internal process.

A TABLES

Table A1. Summary Statistics for Urban CBGs

	N obs.	Mean	Stnd. Dev.	Min.	Median	Max.
Panel A: CBG-level summaries						
POI visits						
Total	27,555	109,174.2	100,220.5	9,807	85,730	2,770,827
N within home county	27,555	84,966.2	78,177.5	2,725	67,239	2,159,970
N within home CBG	27,555	4,538.8	10,349.6	0	2,021	483,216
Travel dist. (km, 75^{th} *pctl.*)	27,555	20.4	16.6	3	17.1	1,046.7
Smoke						
Weeks of smoke	27,555	55.8	9.7	42	51	114
Population counts						
Total	27,555	1,709.8	1,063.6	0	1,492	38,754
Black	27,555	116.2	194.0	0	45	3,821
Hispanic	27,555	585.3	674.4	0	345	11,073
White	27,555	1,141.2	772.5	0	984	30,573
Rural	27,555	22.5	114.3	0	0	5,675
Urban	27,555	1,581	855.8	1	1,403	31,777
Population shares						
Black	27,545	6.7%	10.5%	0%	3%	100%
Hispanic	27,545	32.5%	27.8%	0%	23.2%	100%
White	27,545	68.4%	22.9%	0%	73.4%	100%
Rural	27,555	1.3%	5.9%	0%	0%	49.9%
Urban	27,555	98.7%	5.9%	50.1%	100%	100%
Income						
Med. HH income	26,929	$83,627.5	$43,194	$2,499	$75,017	$250,001
Panel B: CBG-by-week summaries						
POI visit counts						
All	5,758,995	522.4	523.9	4	399	23,395
Within home county	5,758,995	406.5	407.8	0	312	18,114
Within home CBG	5,758,995	21.7	59.6	0	8	17,490
Travel dist. (km, 75^{th} *pctl.*)	5,758,995	48.4	224.3	0	17.6	5,844.2
Smoke						
Any smoke	5,758,995	26.7%		0%		100%
Any ‘low’ smoke	5,758,995	26.6%		0%		100%
Any ‘medium’ smoke	5,758,995	13.6%		0%		100%
Any ‘high’ smoke	5,758,995	7.4%		0%		100%

Notes: This table summarizes west-coast urban CBGs, our main area of study. Table A2 summarize all west-coast CBGs (including rural CBGs). Panel A here summarizes CBG-level data; Panel B summarizes CBG-by-week data—i.e.the level of analysis. We define *urban* CBGs as communities where the urban population exceeds the rural population. We omit socioeconomic data from Panel B because our demographic data (population counts/shares and income) do not vary with time. The Data and Descriptive Statistics section describes variables and sources. The *Smoke* variables in Panel B summarize indicators, so we omit the percentile summaries.

Table A2. Summary Statistics for All CBGs

	N obs.	Mean	Stnd. Dev.	Min.	Median	Max.
Panel A: CBG-level summaries						
POI visit counts						
All	30,174	106,196.6	97,906	7,295	83,509.5	2,770,827
Within home county	30,174	81,820.3	76,608.7	577	64,819	2,159,970
Within home CBG	30,174	4,374.8	10,007.7	0	1,934	483,216
Travel dist. (km, *75th pctl.*)	30,174	22.9	21	3	17.8	1,046.7
Smoke						
Weeks of smoke	30,174	56.7	10.6	42	52	130
Population counts						
Total	30,174	1,673.1	1,044.7	0	1,460	38,754
Black	30,174	108.0	188.4	0	38	3,821
Hispanic	30,174	553.8	660.5	0	312	11,073
White	30,174	1,141.6	759.9	0	988	30,573
Rural	30,172	118.6	365	0	0	5,675
Urban	30,172	1,453.5	918.8	0	1,330	31,777
Population shares						
Black	30,159	6.3%	10.1%	0%	2.6%	100%
Hispanic	30,159	31.1%	27.6%	0%	21.4%	100%
White	30,159	70.3%	23.1%	0%	75.9%	100%
Rural	30,165	9.1%	26.4%	0%	0%	100%
Urban	30,165	90.9%	26.4%	0%	100%	100%
Income						
Med. HH income	29,443	$82,607.2	$42,328.3	$2,499	$73,984	$250,001
Panel B: CBG-by-week summaries						
POI visit counts						
All	6,306,366	508.1	512.2	4	388	26,695
Within home county	6,306,366	391.5	399	0	299	18,114
Within home CBG	6,306,366	20.9	57.5	0	8	17,490
Travel dist. (km, *75th pctl.*)	6,306,366	53	232.4	0	18.7	5,844.2
Smoke						
Any smoke	6,306,366	27.1		0%		100%
Any 'low' smoke	6,306,366	27.1		0%		100%
Any 'medium' smoke	6,306,366	13.9		0%		100%
Any 'high' smoke	6,306,366	7.6		0%		100%

Notes: This table exands the summaries of Table A1 to all CBGs (rather than restricting to urban CBGs).

Table A3. Robustness of Regression Results: Adding State Interaction of Fixed Effects

		Percentile-based heterogeneity			
	(1)	(2) HH Income	(3) % Black	(4) % Hispanic	(5) % White
Panel A *Dependent variable: Percent of POIs visits outside of home county*					
Any smoke	0.30***	-0.41	0.54***	0.82*	0.05
	(0.09)	(0.30)	(0.14)	(0.38)	(0.20)
Any smoke × Het. percentile		1.4**	-0.46*	-0.93	0.55
		(0.53)	(0.24)	(0.67)	(0.35)
N obs. (millions)	5.63	5.63	5.63	5.63	5.63
R^2	0.73	0.73	0.73	0.73	0.73
Panel B *Dependent variable: 75^{th} percentile of distance traveled to POIs (km)*					
Any smoke	1.4**	-7.1*	5.8**	12.9*	-3.6*
	(0.53)	(3.9)	(2.1)	(6.8)	(1.8)
Any smoke × Het. percentile		17.1**	-8.3*	-20.3	11.1**
		(7.6)	(3.8)	(11.8)	(4.7)
N obs. (millions)	5.63	5.63	5.63	5.63	5.63
R^2	0.081	0.081	0.081	0.081	0.081
Fixed effects					
CBG	✓	✓	✓	✓	✓
State × Week of sample	✓	✓	✓	✓	✓

Notes: This table re-estimates the results in Table 1 but with state by week-of-sample fixed effects.

Table A4. Robustness of Regression Results: Dropping CBGs Directly Affected by Wildfires

		Percentile-based heterogeneity			
	(1)	(2) HH Income	(3) % Black	(4) % Hispanic	(5) % White
Panel A *Dependent variable: Percent of POIs visits outside of home county*					
Any smoke	0.28**	-0.41	0.55***	0.90*	-0.10
	(0.10)	(0.30)	(0.15)	(0.39)	(0.22)
Any smoke × Het. percentile		1.4**	-0.50*	-1.20	0.80*
		(0.53)	(0.24)	(0.70)	(0.38)
N obs. (millions)	5.54	5.54	5.54	5.54	5.54
R^2	0.73	0.73	0.73	0.73	0.73
Panel B *Dependent variable: 75^{th} percentile of distance traveled to POIs (km)*					
Any smoke	1.7***	-6.9*	6.0***	11.7*	-3.4
	(0.54)	(4.1)	(2.1)	(6.0)	(2.3)
Any smoke × Het. percentile		17.2**	-8.1*	-19.0	10.8*
		(7.7)	(3.7)	(11.6)	(4.9)
N obs. (millions)	5.54	5.54	5.54	5.54	5.54
R^2	0.078	0.079	0.079	0.079	0.079
Fixed effects					
CBG	✓	✓	✓	✓	✓
Week of sample	✓	✓	✓	✓	✓

Notes: This table re-estimates the results in Table 1 but without CBGs that were ever affected by wildfires between 2018–2021 (i.e.CBG boundaries that intersect with wildfire perimeters).

Table A5. Aggregate Discount Parameter Estimates From Quasi-hyperbolic Model, Full Sample

Model:	OLS	Interval Regression	NLS	Logistic Choice
	(1)	(2)	(3)	(4)
δ	0.9980	0.9980	0.9981	0.9984
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
β	0.9712	0.9716	0.9724	0.9778
	(0.0024)	(0.0025)	(0.0024)	(0.0024)
r (annual discount rate)	0.1074	0.1104	0.1047	0.0877
	(0.0120)	(0.0136)	(0.0118)	(0.0115)
H_0: $\delta = 1$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$
H_0: $\beta = 1$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$

Notes: standard errors are clustered at the subject level. δ is the weekly exponential discount factor, and $r = \delta^{-52} - 1$ is the annual discount rate. The data consist of 6,367 price lists from 153 subjects.

Table A6. Linear-time-trend Estimates of Time Variance, Full Sample

	(1)	(2)	(3)	(4)	(5)
Constant	0.9561	0.9640	0.9643	0.9643	0.9646
	(0.0033)	(0.0034)	(0.0036)	(0.0035)	(0.0036)
Choice delay $((t^L - t^S) - 8)$	-0.0006***	-0.0005***	-0.0003	-0.0005***	-0.0003
	(0.0002)	(0.0002)	(0.0003)	(0.0002)	(0.0003)
Front-end delay $(1 - \mathbb{1}(t^S = \tau))$	0.0034***	0.0019	0.0018	0.0011	0.0010
	(0.0013)	(0.0012)	(0.0012)	(0.0021)	(0.0021)
Survey week (w)		-0.0013**	-0.0013**	-0.0014**	-0.0014**
		(0.0005)	(0.0005)	(0.0005)	(0.0005)
$w \cdot ((t^L - t^S) - 8)$			-0.0000		-0.0000
			(0.0000)		(0.0000)
$w \cdot (1 - \mathbb{1}(t^S = \tau))$				0.0002	0.0001
				(0.0004)	(0.0004)

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$. Coefficients are from linear models with standard errors are clustered at the subject level. The data consist of 6,367 price lists from 153 subjects.

Table A7. Separate Estimates of Non-stationarity and Inconsistency, Full Sample

Violation type:	Stationarity		Consistency	
	(1)	(2)	(3)	(4)
Constant	0.9614	0.9613	0.9617	0.9619
	(0.0032)	(0.0032)	(0.0031)	(0.0031)
Choice delay $((t^L - t^S) - 8)$	-0.0005***	-0.0004**	-0.0005***	-0.0005***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Front-end delay $(1 - \mathbb{1}(t^S = \tau))$	0.0007	0.0010	0.0042***	0.0044***
	(0.0012)	(0.0012)	(0.0014)	(0.0015)
$((t^L - t^S) - 8) \cdot (1 - \mathbb{1}(t^S = \tau))$		-0.0002**		-0.0001
		(0.0001)		(0.0001)
Survey-week FEs	Y	Y	N	N
Drawing-week FEs	N	N	Y	Y

Notes: Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Standard errors are clustered at the subject level. The data consist of 6,367 price lists from 153 subjects.

Table A8. Subject- and Decision-triangle-level Classification of Time Preference Properties, Full Sample

Margin of Error:	None	± 6 tickets
Properties of choices	(1)	(2)
Panel A: Subject level		
A) Invariant, consistent, stationary	6.38%	20.57%
B) Invariant only	1.42%	0%
A) + B)	7.80%	20.57%
C) Stationary only	3.55%	7.09%
D) Consistent only	2.13%	1.42%
E) None	86.52%	70.92%
C) + D) + E)	92.2%	79.43%
Panel B: Triangle level		
A) Invariant, consistent, stationary	33.88%	61.61%
B) Invariant only	6.26%	7.30%
A) + B)	40.14%	68.91%
C) Stationary only	15.20%	15.20%
D) Consistent only	8.34%	8.30%
E) None	36.31%	7.58%
C) + D) + E)	59.86%	31.09%

Notes: Data 141 subjects that completed at least two consecutive surveys. Column(1) assumes any deviation from is a time preference property violation. Column(2) assumes any deviation greater than six tickets from an analogous choice is a time preference violation, capturing a range of a standard deviation around a choice ($\sigma/2 \approx 7$, which due to interval of two tickets between choice rows corresponds to equality within six tickets).

Table A9. Aggregate Parameter Estimates for the Nested Exponential Model, Full Sample

Restrictions:	$\mu = \eta = \gamma = 1$	$\mu = \gamma = 1$	$\mu = \eta = 1$	$\eta = \gamma = 1$	None
Properties					
Invariant	✓	✓	X	X	X
Stationary	✓	X	✓	X	X
Consistent	✓	X	X	✓	X
	(1)	(2)	(3)	(4)	(5)
δ	0.9973	0.9778	0.9963	0.9756	0.9735
	(0.0002)	(0.0020)	(0.0004)	(0.0032)	(0.0027)
η	1	0.3811	1	1	0.4291
	.	(0.0257)	.	.	(0.0408)
γ	1	1	1.0460	1	0.9918
	.	.	(0.0168)	.	(0.0168)
μ	1	1	1	0.4801	0.7664
	.	.	.	(0.0342)	(0.1207)
Log-likelihood	-26,484.8	-26,134.2	-26,474.4	-26,312.6	-26,131.9
AIC	52,971.7	52,272.3	52,952.7	52,629.1	52,271.7
r_0	0.1481	0.1065	0.2113	0.1792	0.1033
	(0.0140)	(0.0102)	(0.0264)	(0.0145)	(0.0124)
r_{12}	$= r_0$	$= r_0$	0.1183	0.1058	0.1064
			(0.0174)	(0.0115)	(0.0142)
PB_S	1	0.9755	1	0.9735	0.9806
	.	(0.0023)	.	(0.0038)	(0.0109)
PB_C	1	$= PB_S$	1.0025	1	0.9737
	.		(0.0011)	.	(0.0023)
H_0: $\delta = 1$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$	$p < 0.0001$
H_0: $\eta = 1$	.	$p < 0.0001$	.	.	$p < 0.0001$
H_0: $\gamma = 1$	.	.	$p = 0.0070$	.	$p = 0.6248$
H_0: $\mu = 1$	.	.	.	$p < 0.0001$	$p = 0.0548$
H_0: $\mu = \gamma = 1$	.	.	.	.	$p = 0.0531$
H_0: $\mu = \eta = 1$	.	.	.	.	$p < 0.0001$
H_0: $\eta = \gamma = 1$	.	.	.	.	$p < 0.0001$
H_0: $\mu = \eta = \gamma = 1$	.	.	.	.	$p < 0.0001$
H_0: $r_0 = r_{12}$	.	.	$p = 0.0059$	$p < 0.0001$	$p = 0.8682$
H_0: $PB_S = 1$	.	$p < 0.0001$	.	$p < 0.0001$	$p = 0.0743$
H_0: $PB_C = 1$	.	$p < 0.0001$	$p = 0.0200$	.	$p < 0.0001$
H_0: $PB_S = PB_C$	.	.	.	.	$p = 0.5267$

Notes: standard errors are clustered at the subject level. $r_\tau = D(\tau, \tau + 52)^{-1} - 1$ is a measure of the annual discount rate. PB_S and PB_C are stationarity-based and consistency-based measures of present-bias. The data consist of 6,367 price lists from 153 subjects. All estimates are from non-linear least squares regressions.

Table A10. Subject-level Structural Classification of Time Preference Properties, Full Sample

Properties of choices	
A) Invariant, consistent, stationary	6.52%
B) Invariant only	39.13%
A) + B)	45.65%
C) Stationary only	6.52%
D) Consistent only	7.61%
E) None	40.22%
C) + D) + E)	54.35%

Notes: Data comes from 92 subjects for which we were able to obtain parameter estimates from all models.

Table A11. Drawing Week Fixed Effects Models for Nicotine Variables

	Nicotine Use	Nicotine Regret
Choice delay $((t^L - t^S) - 8)$	-0.0005***	-0.0005
	(0.0002)	(0.0004)
Front-end delay $(1 - \mathbb{1}(t^S = \tau))$	0.0051***	0.0031
	(0.0016)	(0.0033)
Variable	-0.0184	0.0044
	(0.0091)	(0.0162)
Variable * $(1 - \mathbb{1}(t^S = \tau))$	0.0015	0.0094**
	(0.0028)	(0.0041)
N	6,367	1400

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$. Coefficients are from linear models with standard errors are clustered at the subject level. Drawing-week fixed effects are present in both models. The data consist of 5,175 price lists from 97 subjects.

Table A12. Aggregate Parameter Estimates for Asymptotic Standard Exponential Model

Sample:	Full	Balanced
δ	0.9980	0.9977
	(0.0003)	(0.0003)
β	-0.0043	-0.0025
	(0.0007)	(0.0010)
Log-likelihood	-26,438.0	-18,141.3
AIC	52,881.9	36,288.5
r_0	0.3915	0.2876
	(0.0470)	(0.0567)
r_{12}	0.1298	0.1405
	(0.0145)	(0.0183)
H_0: $\delta = 1$	$p < 0.0001$	$p < 0.0001$
H_0: $\beta = 0$	$p < 0.0001$	$p = 0.0107$

Notes: standard errors are clustered at the subject level. The Full Sample data consist of 6,367 price lists from 153 subjects.The Balanced Sample data consist of 4,345 price lists from 79 subjects. $r_\tau = D(\tau, \tau + 52)^{-1} - 1$ is a measure of the annual discount rate.

B FIGURES

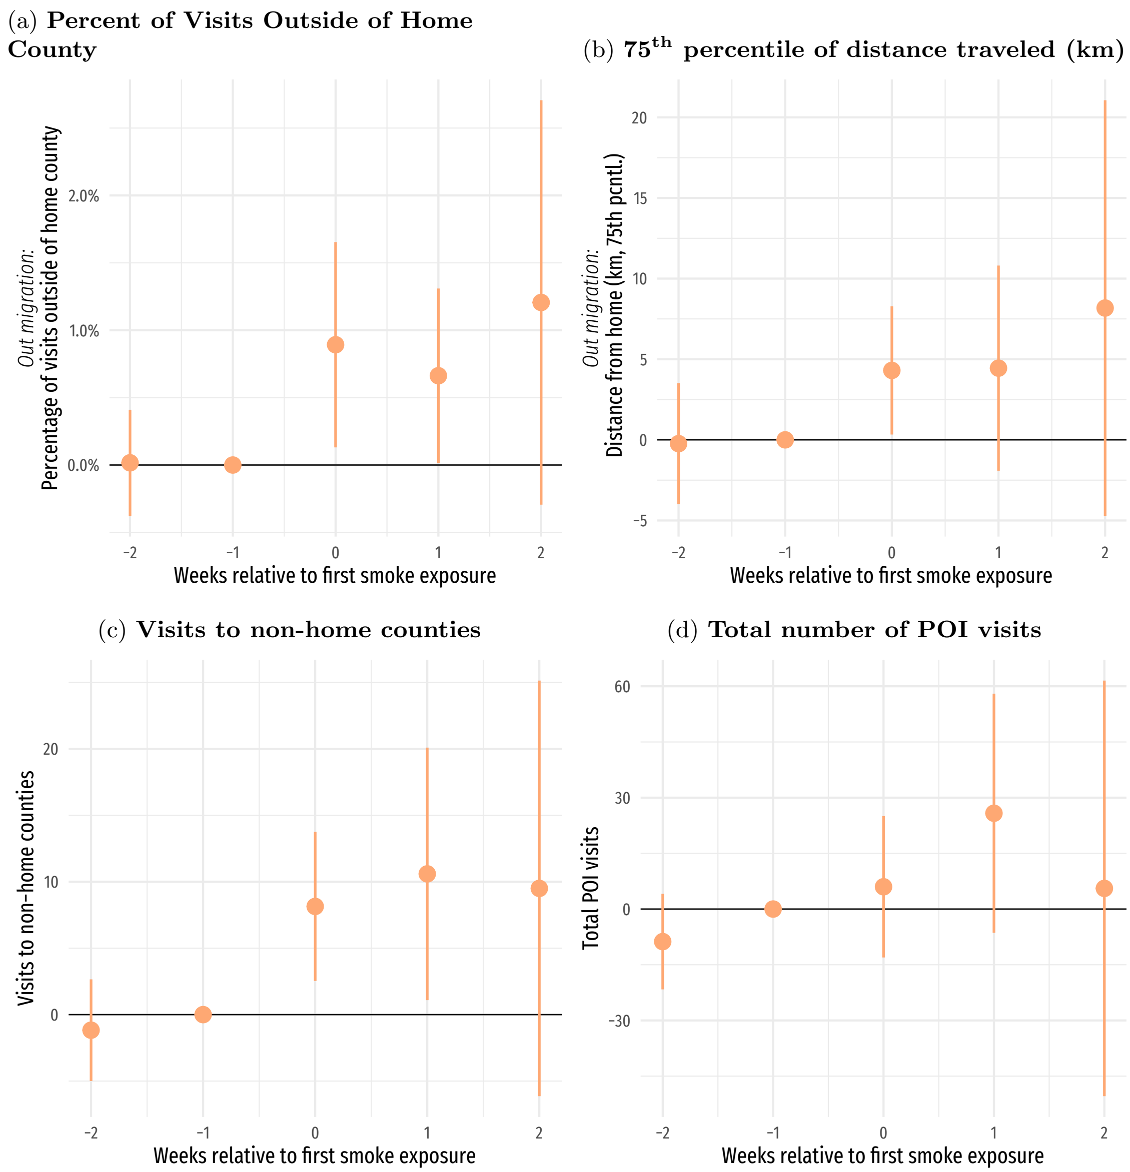

Figure B1. Event Studies of Smoke Exposure

Each event study is centered on the first week that CBGs encounter smoke for each calendar year (excluding 2021). The regression includes a fixed effect for the week the smoke started, county-month (of calendar), and county-year. We cluster the standard errors by county and month. Panels (a)–(d) only differ in their outcome variable: (a) CBG residents' shares of visits beyond their home counties, (b) CBGs' 75^{th} percentiles of distance traveled in the week, (c) the number of visits to POIs in non-home counties, and (d) the total number of POI visits. Panel (a)'s outcome is the outcome in (c) divided by the outcome in (d).

(a) **Income quantiles** (b) **Percent White quantiles**

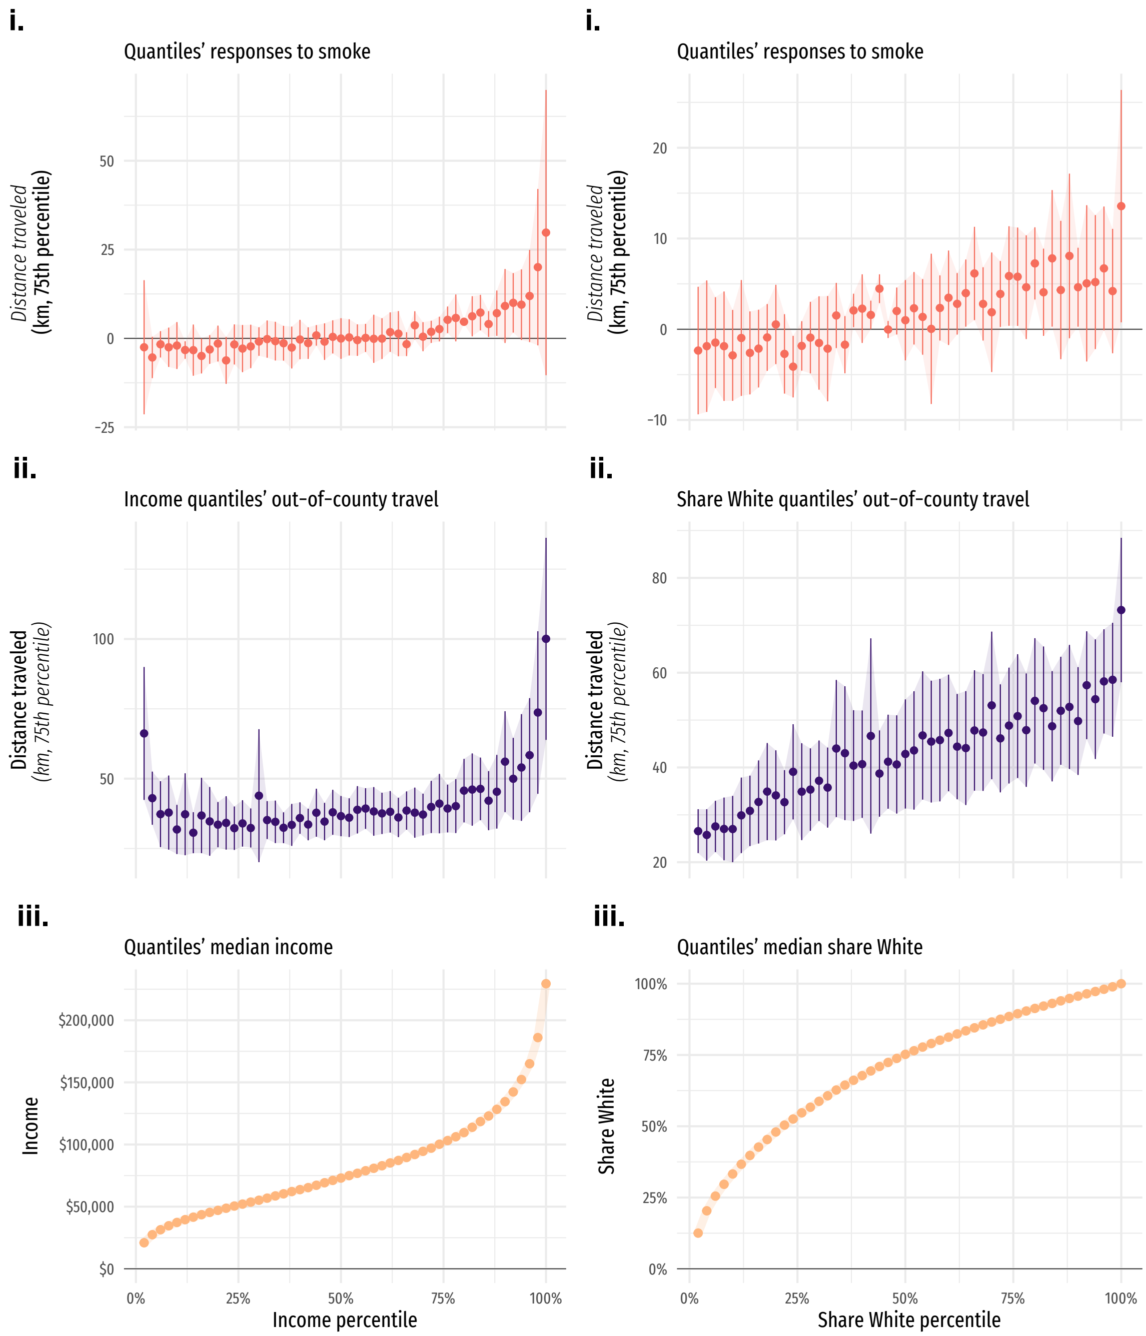

Figure B2. Unequal Smoke-induced Travel: By Income and Percent White

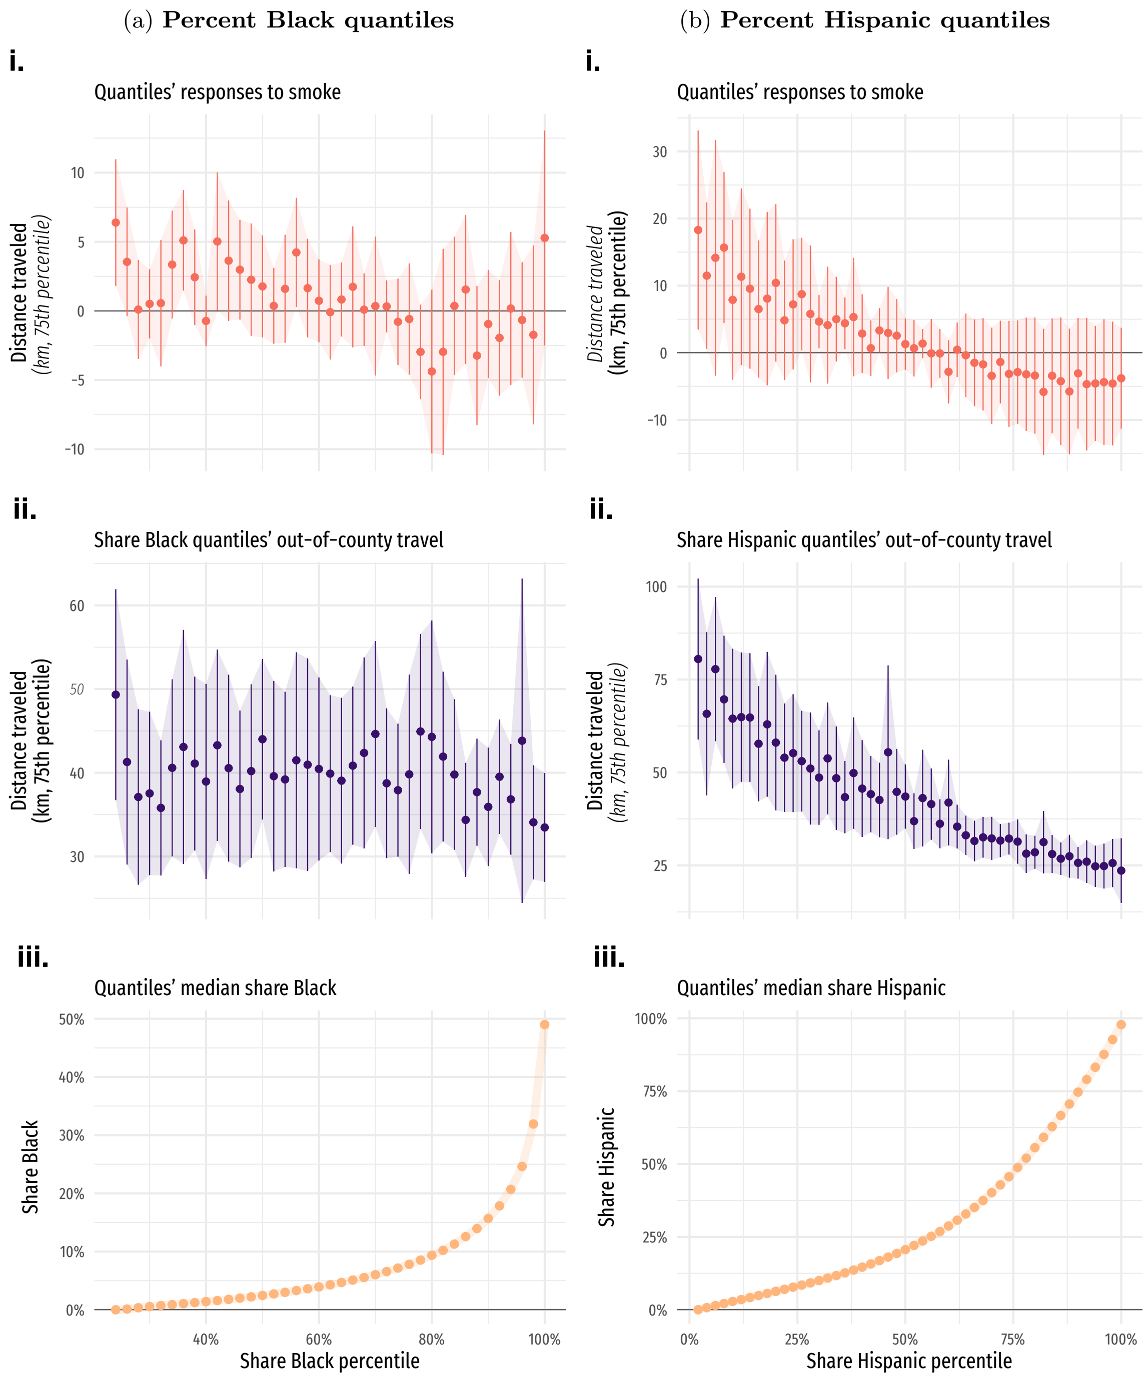

Figure B3. Unequal Smoke-induced Travel: By Percent Black and Percent Hispanic

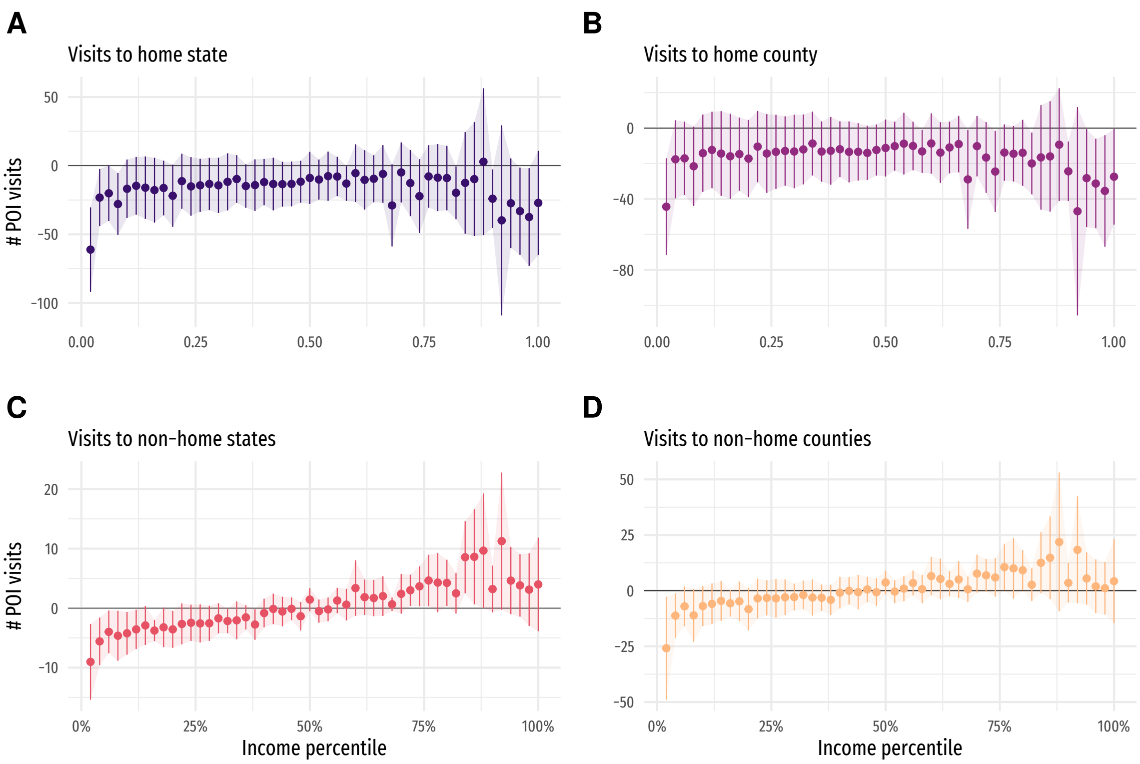

Figure B4. Smoke Exposure's Effect on Visits to Home and Non-home Counties/States

This figure estimates smoke exposure's impact on the *level* of visits (the count)—rather than the percentage—across income percentiles (grouped into two-percentile bins). Panel **A** shows the effect of smoke exposure on the number of visits to residents' home **states**; Panel **B** shows the effect of smoke exposure on the number of visits to residents' home **counties**. Panels **C** and **D** show the effect of smoke exposure on visits to other states (**C**) and counties (**D**).

	Prize Draw today	or	Prize Draw in 2 weeks
Choice 1:	202		242
Choice 2:	204		242
Choice 3:	206		242
Choice 4:	208		242
Choice 5:	210		242
Choice 6:	212		242
Choice 7:	214		242
Choice 8:	216		242
Choice 9:	218		242
Choice 10:	220		242
	I prefer 222 tickets today to 242 tickets in 2 weeks		
Choice 11:	222		242
Choice 12:	224		242
Choice 13:	226		242
Choice 14:	228		242
Choice 15:	230		242
Choice 16:	232		242
Choice 17:	234		242
Choice 18:	236		242
Choice 19:	238		242
Choice 20:	240		242
Choice 21:	242		242
Choice 22:	244		242
	Prize Draw today	or	Prize Draw in 2 weeks

Next

Figure B5. Example Price List

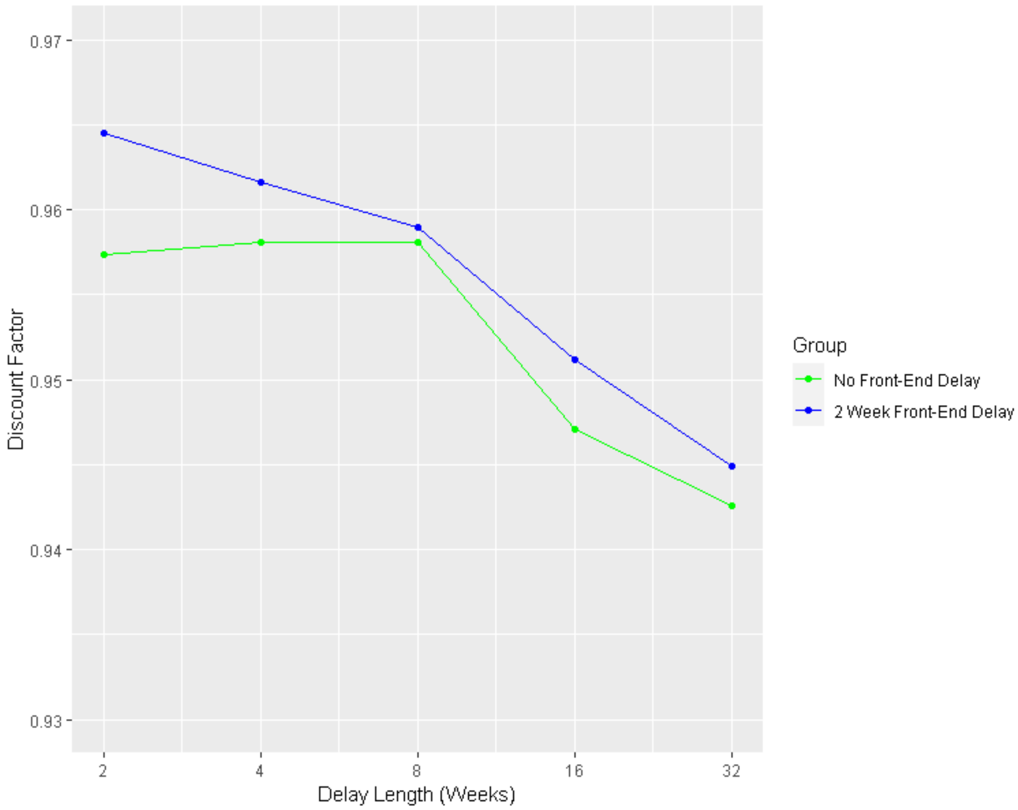

Figure B6. Average Discount Factor by Choice Delay, Full Sample

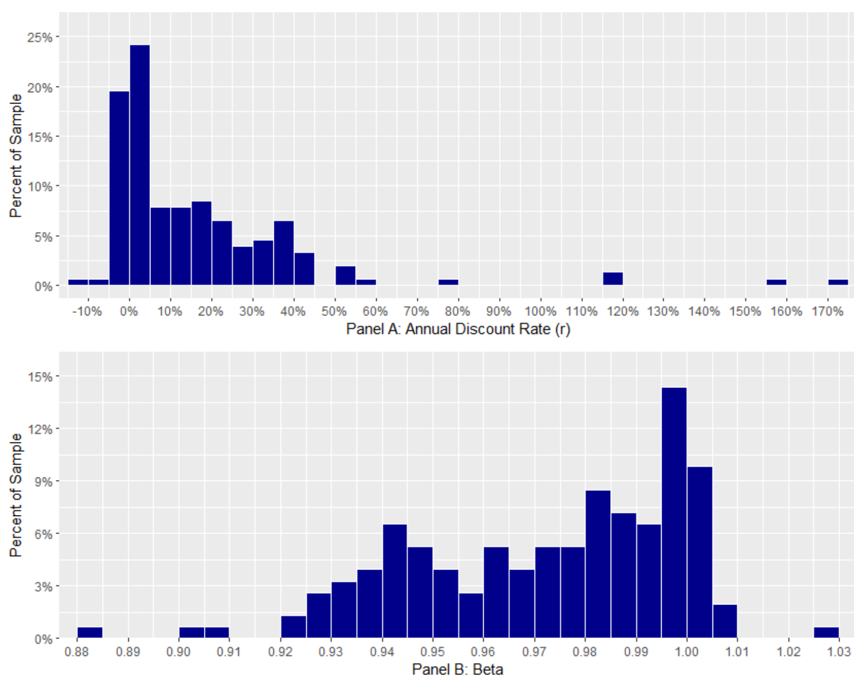

Figure B7. Distributions of Individual Estimates From the β-δ Model, Full Sample

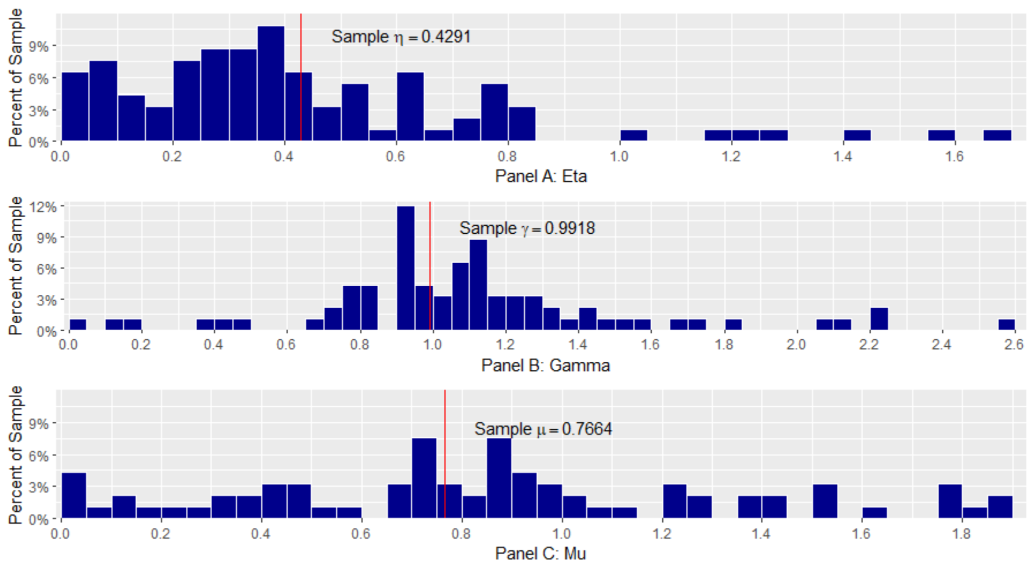

Figure B8. Individual Parameter Estimates for the Unrestricted Nested Exponential Model

Notes: 18 estimates are suppressed from the right tails of the distributions in Panel (B) and Panel (C) for visual clarity.

Imagine you have won a prize, but the longer you wait to redeem that prize, the bigger it gets. Which of the following options would you most prefer?

- $50 in 8 weeks
- $53 in 9 weeks
- $54 in 10 weeks
- $55 in 12 weeks
- $56 in 14 weeks
- $57 in 16 weeks
- $58 in 18 weeks
- $59 in 20 weeks
- $60 in 23 weeks
- $61 in 26 weeks
- $62 in 30 weeks
- $63 in 34 weeks

→ Next

Figure B9. Rotating Question Week 4

Imagine you have won a prize, but the longer you wait to redeem that prize, the bigger it gets. Which of the following options would you most prefer?

- $50 today
- $53 in 1 week
- $54 in 2 weeks
- $55 in 4 weeks
- $56 in 6 weeks
- $57 in 8 weeks
- $58 in 10 weeks
- $59 in 12 weeks
- $60 in 15 weeks
- $61 in 18 weeks
- $62 in 22 weeks
- $63 in 26 weeks

→

Figure B10. Rotating Question Week 12 (Week 4 Remeasure)

For each of the time horizons below, do you consider events in your life that you expect at that time to be in the "present", the "near future" or the "distant future"?

	Present	Near future	Distant future
The end of this survey	○	○	○
In 2 weeks	○	○	○
In 4 weeks	○	○	○
In 8 weeks	○	○	○
In 16 weeks	○	○	○
In 32 weeks	○	○	○

Figure B11. Exit Survey Question - Patience: Time Horizons

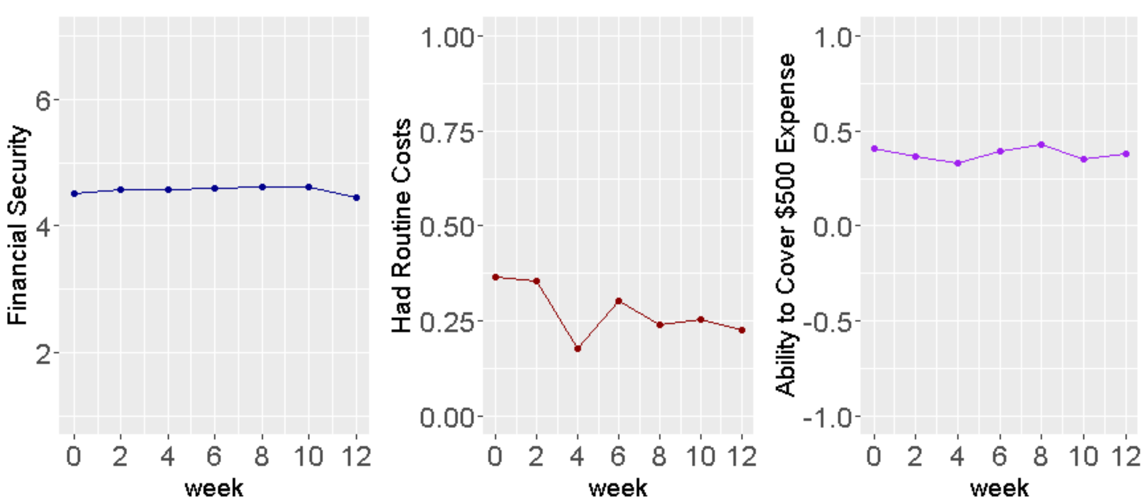

Figure B12. Time-Trends of Dynamic Variables: Balanced Sample

C DATA

C.1 Privacy, Noise, and Censoring in SafeGraph Weekly Patterns Data.

The following quote from SafeGraph's *Patterns* documentation SafeGraph (2022b) describes the company's approach to manipulating their aggregated data products so as to better protect individual privacy SafeGraph (2022a).

> To preserve privacy, we apply differential privacy techniques to the following columns: `visitor_home_cbgs`, `visitor_home_aggregation`, `visitor_daytime_cbgs`, `visitor_country_of_origin`, `device_type`, `carrier_name`. We have added Laplacian noise to the values in these columns. After adding noise, only attributes (e.g., a census block group) with at least two devices are included in the data. If there are between 2 and 4 visitors this is reported as 4.

As described in Movement Data, our outcome variables use the count of visitors decomposed by the visitors' home CBGs, i.e.the variable `visitor_home_cbgs`. SafeGraph's differential-privacy approach likely has little effect on our estimates of the level and equity of smoke-induced migration. First, because we aggregate to CBG (across many POIs within each CBG) and SafeGraph's manipulation mainly affects low-count observations at the POI level, we still accurately account for the vast majority of visits. Second, our distance-based measurement of out-migration uses the 75^{th} percentile—i.e.a measure that is relatively robust to small changes in the tails of a distribution. Finally, the differential-privacy approach affects our outcome variable (rather than an explanatory variable), so any measurement error merely ends up in the error term (rather than biasing our point estimates).

C.2 Five-Year American Community Survey. The 2019 five-year ACS estimates aggregate the prior five years of survey data collected by the US Census Bureau in the ACS. The five-year estimates offer the advantage of supplying CBG-level data spanning the entire US—the shorter time span 1-year estimates are restricted to higher population areas. Five-year estimates likely better match the 'real-time' demographics of the sample period than the 2010 decennial census. The relevant 2020 decennial data were not available at the time of analysis.

Finally, the Census censors both ends of the ACS data on CBG-level median household income (below \$2,500 and above \$250,000)—as is evident in the summary-statistic tables (Tables A1 and A2).

D THEORY

D.1 Proof of Proposition II.

Proof. ($\Leftarrow$) Suppose a discount function $D(\tau, t)$ is only a function of $t - \tau$. Given bundles $(\varepsilon,\, t + \Delta_1) \sim_t (\psi,\, t + \Delta_2)$ which the agent is indifferent between at evaluation time t, we have

$$\begin{aligned} D(t,\, t + \Delta_1)\, \varepsilon &= D(t, t + \Delta_2)\, \psi \\ \implies D(t + \Delta_1 - t)\, \varepsilon &= D(t + \Delta_2 - t)\, \psi \\ \implies D(\Delta_1)\, \varepsilon &= D(\Delta_2)\, \psi \end{aligned}$$

Note that for any t', we can certainly write $\Delta_i = t' + \Delta_i - t'$ for $i = 1,\, 2$. Thus, the above equation implies

$$D(t' + \Delta_1 - t')\, \varepsilon = D(t' + \Delta_2 - t')\, \psi$$

By assumption, this is equivalent to

$$D(t',\, t' + \Delta_1)\, \varepsilon = D(t',\, t' + \Delta_2)\, \psi$$

which establishes Time Invariance.

$(\impliedby)$

($\Rightarrow$) Let $D(\tau, t)$ be a time-invariant function. Suppose there exists bundles $(\varepsilon,\, t + \Delta_1) \sim_t (\psi,\, t + \Delta_2)$ which the agent is indifferent between at evaluation time t. According to the definition of Time Invariance, the agent's discounted utility must satisfy

$$D(t,\, t + \Delta_1)\, \varepsilon = D(t,\, t + \Delta_2)\, \psi$$

for <u>all</u> $t \geq 0$. Namely, let us choose $t = 0$. Since

$$\begin{aligned} D(0,\, \Delta_1)\, \varepsilon &= D(0,\, \Delta_2)\, \psi \\ \implies D(0,\, t - \tau)\, \varepsilon &= D(0,\, t' - \tau)\, \psi \end{aligned}$$

$(\implies)$

■

D.2 Proof of Proposition IV.

Proof. ($\Leftarrow$) It is straightforward to verify that $D(\tau,\, t) \;=\; \alpha^{-(t-\tau)f(\tau)}$ is a stationary function. Consider bundles $(\varepsilon,\, t)$ and $(\psi,\, t+\Delta)$ which the discounter is indifferent between at time τ. Let $t' \geq \tau$ with $t' - t = k$. Exponent rules tell us

$$
\begin{aligned}
D(\tau,\, t')\, \varepsilon &= \alpha^{-(t'-\tau)f(\tau)}\, \varepsilon \\
&= \alpha^{-kf(\tau)}\, D(\tau,\, t)\, \varepsilon
\end{aligned}
$$

The indifference between $(\varepsilon,\, t)$ and $(\psi,\, t+\Delta)$ further tells us that this equals $\alpha^{-kf(\tau)}\, D(\tau,\, t + \Delta)\, \psi$. Expanding upon the particular form of the discount, it is straightforward to refactor this expression into $D(\tau,\, t')\, \psi$. Thus, stationarity has been established.

$(\Longleftarrow)$

($\Rightarrow$) Suppose $D(\tau,\, t)$ admits stationary preferences. The definition of stationarity implies that if an agent is indifferent between $(\varepsilon,\, t)$ and $(\psi,\, t+\Delta)$ when evaluated at time τ, then they are also indifferent between bundles $(\varepsilon,\, t')$ and $(\psi,\, t'+\Delta)$ when evaluated at time τ, for any $t' \geq \tau$. In terms of discounted utility, we can summarize this as

$$
D(\tau,\, t)\, \varepsilon = D(\tau,\, t+\Delta)\, \psi
$$

for all $t \geq \tau$. Taking the log of both sides, we obtain

$$
\log\left[D(\tau,\, t)\right] + \log\, \varepsilon = \log\left[D(\tau,\, t+\Delta)\right] + \log\, \psi
$$

Since this holds for arbitrary (sensible) t, we can take the derivative on either side with respect to t. Letting $\dot{D}$ denote the derivative of D with respect to payoff time, t, we obtain

$$
\frac{\dot{D}(\tau,\, t)}{D(\tau,\, t)} = \frac{\dot{D}(\tau,\, t+\Delta)}{D(\tau,\, t+\Delta)}
$$

Therefore, the rate of change of the discount function is independent of payoff time. In other words, given an arbitrary function $g(\cdot)$, stationarity implies that the rate of change of the discount function is solely a function of τ:

$$
\frac{\dot{D}(\tau,\, t)}{D(\tau,\, t)} = g(\tau)
$$

Review of an entry-level differential equations textbook reveals that one particular function which satisfies $h'(x) \;=\; k \cdot h(x)$ is the exponential function $h(x) \;=\; a \cdot e^{kx+b}$. The extension to our multivariate world is straightforward: one solution to the differential equation $\dot{D}(\tau,\, t) = g(\tau)D(\tau,\, t)$ is given by

$$D(\tau, t) = a \cdot e^{\,f_1(\tau)\,\cdot\, t \;+\; f_2(\tau)}$$

where $f_1' = g$ and a is a constant. Since we require $D(0,\, 0) = 1$, we can immediately say that $a = 1$. Furthermore, we require that $D(\tau,\, t) = 0$ whenever $\tau = t$, restricting our exponent to satisfy

$$\begin{aligned} & f_1(\tau) \cdot \tau + f_2(\tau) = 0 \\ \Longrightarrow & \boxed{f_2(\tau) = -\tau \cdot f_1(\tau)} \end{aligned}$$

Consequently, we can drop the subscripts on f and our exponent can be reduced to $tf(\tau) - \tau f(\tau) = (t - \tau)f(\tau)$. That is, our stationary discount function is given by

$$D(\tau,\, t) = e^{(t-\tau)f(\tau)}$$

To allow for full generality, we replace e with some $\alpha > 0$. We would like D to be clearly decreasing in t, but this depends on the sign of $1 - \alpha$ (as well as the sign of f). Since f is arbitrary, we can factor -1 out of it when $\alpha < 1$. As this is often the case, our function now takes the form

$$D(\tau,\, t) = \alpha^{-(t-\tau)f(\tau)}$$

Hence, we have shown that stationarity implies a discount function of this form.

$(\Longrightarrow)$

■

Aizer, A., & Currie, J. (2014, May). The intergenerational transmission of inequality: Maternal disadvantage and health at birth. *Science*, *344*(6186), 856–861. Retrieved from https://doi.org/10.1126/science.1251872 doi: 10.1126/science.1251872

Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., & Yang, D. (2020, November). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics*, *191*, 104254. Retrieved from https://doi.org/10.1016/j.jpubeco.2020.104254 doi: 10.1016/j.jpubeco.2020.104254

Almond, D., & Chay, K. Y. (2006). The long-run and intergenerational impact of poor infant health: Evidence from cohorts born during the civil rights era. *Working Paper*. Retrieved from https://users.nber.org/~almond/chay_npc_paper.pdf

Alsan, M., Chandra, A., & Simon, K. (2021, August). The great unequalizer: Initial health effects of COVID-19 in the united states. *Journal of Economic Perspectives*, *35*(3), 25–46. Retrieved from https://doi.org/10.1257/jep.35.3.25 doi: 10.1257/jep.35.3.25

Andersen, S., Harrison, G. W., Lau, M. I., & Rutström, E. E. (2014). Discounting behavior: A reconsideration. *European Economic Review*, *71*, 15–33.

Andersen, S., Harrison, G. W., Lau, M. I., & Rutström, E. E. (2008). Eliciting Risk and Time Preferences. *Econometrica*, *76*(3), 583–618.

Anderson, S., Plantinga, A. J., & Wibbenmeyer, M. (2022, August). Inequality in agency response: Evidence from salient wildfire events. *The Journal of Politics*. Retrieved from https://doi.org/10.1086/722044 doi: 10.1086/722044

Anderton, D. L., Anderson, A. B., Oakes, J. M., & Fraser, M. R. (1994, May). Environmental equity: The demographics of dumping. *Demography*, *31*(2), 229–248. Retrieved from https://doi.org/10.2307/2061884 doi: 10.2307/2061884

Andreoni, J., Kuhn, M. A., & Sprenger, C. (2017). Measuring time preferences: A comparison of experimental methods. *Journal of Economic Behavior and Organization*, *116*, 451–464.

Andreoni, J., & Sprenger, C. (2012). Estimating Time Preferences from Convex Budgets. *American Economic Review*, 50.

Aneja, A., & Avenancio-Leon, C. F. (2022). The effect of political power on labor market inequality: Evidence from the 1965 voting rights act. *American Economic Review*, *Forthcoming*.

Aneja, A., & Xu, G. (2021, October). The costs of employment segregation: Evidence from the federal government under woodrow wilson. *The Quarterly Journal of Economics*, *137*(2), 911–958. Retrieved from https://doi.org/10.1093/qje/qjab040 doi: 10.1093/qje/qjab040

Ashton, L. (2015). *Hunger games: Does hunger affect time preferences?* (SSRN Working Paper)

Attema, A. E., Bleichrodt, H., Rohde, K. I. M., & Wakker, P. P. (2010). Time-tradeoff sequences for analyzing discounting and time inconsistency. *Management Science*, *56*(11), 2015-2030.

Augenblick, N., Niederle, M., & Sprenger, C. (2015). Working over time: Dynamic inconsistency in real effort tasks. *Quarterly Journal of Economics*, *130*(3), 1067-1115.

Autor, D., Figlio, D., Karbownik, K., Roth, J., & Wasserman, M. (2019, July). Family disadvantage and the gender gap in behavioral and educational outcomes. *American Economic Journal: Applied Economics*, *11*(3), 338–381. Retrieved from `https://doi.org/10.1257/app.20170571` doi: 10.1257/app.20170571

Aycinena, D., Blazsek, S., Rentschler, L., & Sprenger, C. (2022). Intertemporal choice experiments and large-stakes behavior. *Journal of Economic Behavior and Organization*, *196*, 484-500.

Banzhaf, S., Ma, L., & Timmins, C. (2019, February). Environmental justice: The economics of race, place, and pollution. *Journal of Economic Perspectives*, *33*(1), 185–208. Retrieved from `https://doi.org/10.1257/jep.33.1.185` doi: 10.1257/jep.33.1.185

Barreca, A., Clay, K., Deschênes, O., Greenstone, M., & Shapiro, J. S. (2015, May). Convergence in adaptation to climate change: Evidence from high temperatures and mortality, 1900–2004. *American Economic Review*, *105*(5), 247–251. Retrieved from `https://doi.org/10.1257/aer.p20151028` doi: 10.1257/aer.p20151028

Barreca, A., Clay, K., Deschenes, O., Greenstone, M., & Shapiro, J. S. (2016, February). Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century. *Journal of Political Economy*, *124*(1), 105–159. Retrieved from `https://doi.org/10.1086/684582` doi: 10.1086/684582

Bayer, P., Keohane, N., & Timmins, C. (2009, July). Migration and hedonic valuation: The case of air quality. *Journal of Environmental Economics and Management*, *58*(1), 1–14. Retrieved from `https://doi.org/10.1016/j.jeem.2008.08.004` doi: 10.1016/j.jeem.2008.08.004

Baylis, P., & Boomhower, J. (2021, December). Mandated vs. voluntary adaptation to natural disasters: The case of u.s. wildfires. *NBER Working Paper 29621*. Retrieved from `https://doi.org/10.3386/w29621` doi: 10.3386/w29621

Baylis, P., & Boomhower, J. (2022). The economic incidence of wildfire suppression in the united states. *American Economic Journal: Applied Economics*, *Forthcoming*.

Bchir, M., & Willinger, M. (2013). *Does the exposure to natural hazards affect risk and time preferences? some insights from a field experiment in perú.* (RePEc Working Paper #13-04)

Beine, M., Charness, G., Dupuy, A., & Joxhe, M. (2020). *Shaking things up: on the stability of risk and time preferences.* (CESifo. Working Paper No. 8187)

Belot, M., Kircher, P., & Muller, P. (2021). *Eliciting time preferences when income and consumption vary: Theory, validation and application to job search.* (IZA Discussion Paper No. 14091)

Bernheim, B. D., Braghieri, L., Martínez-Marquina, A., & Zuckerman, D. (2021). A theory of chosen preferences. *American Economic Review*, *111*(2), 720–754.

Bernheim, B. D., & Taubinsky, D. (2018). Behavioral public economics. *Handbook of behavioral economics: Applications and Foundations 1*, *1*, 381–516.

Bertocchi, G., & Dimico, A. (2012, November). The racial gap in education and the legacy of slavery. *Journal of Comparative Economics*, *40*(4), 581–595. Retrieved from `https://doi.org/10.1016/j.jce.2012.04.001` doi: 10.1016/j.jce.2012.04.001

Blanden, J., Doepke, M., & Stuhler, J. (2022, April). Educational inequality. *NBER Working Paper 29979*. Retrieved from `https://doi.org/10.3386/w29979` doi: 10.3386/w29979

Borgschulte, M., Molitor, D., & Zou, E. Y. (2022, September). Air pollution and the labor market: Evidence from wildfire smoke. *The Review of Economics and Statistics*, 1–46. Retrieved from `https://doi.org/10.1162/rest_a_01243` doi: 10.1162/rest_a_01243

Bourdeau-Brien, M., & Kryzanowski, L. (2020). Natural disasters and risk aversion. *Journal of Economic Behavior and Organization*, *177*, 818–835.

Boustan, L. P., Kahn, M. E., & Rhode, P. W. (2012, May). Moving to higher ground: Migration response to natural disasters in the early twentieth century. *American Economic Review*, *102*(3), 238–244. Retrieved from `https://doi.org/10.1257/aer.102.3.238` doi: 10.1257/aer.102.3.238

Brownback, A., Imas, A., & Kuhn, M. A. (2023). Behavioral food subsidies. *The Review of Economic and Statistics*. (Forthcoming)

Bullard, R. D., Mohai, P., Saha, R., & Wright, B. (2008). Toxic wastes and race at twenty: Why race still matters after all of these years. *Envtl. L.*, *38*, 371.

Burke, M., Driscoll, A., Heft-Neal, S., Xue, J., Burney, J., & Wara, M. (2021, January). The changing risk and burden of wildfire in the united states. *Proceedings of the National Academy of Sciences*, *118*(2). Retrieved from `https://doi.org/10.1073/pnas.2011048118` doi: 10.1073/pnas.2011048118

Burke, M., & Emerick, K. (2016, August). Adaptation to climate change: Evidence from US agriculture. *American Economic Journal: Economic Policy*, *8*(3), 106–140. Retrieved from `https://doi.org/10.1257/pol.20130025` doi: 10.1257/pol.20130025

Burke, M., Heft-Neal, S., Li, J., Driscoll, A., Baylis, P., Stigler, M., ... Gould, C. F. (2022, 7). Exposures and behavioural responses to wildfire smoke. *Nature Human Behaviour*. Retrieved from `https://doi.org/10.1038/s41562-022-01396-6` doi: 10.1038/s41562-022-01396-6

Burness, H. S. (1976). A note on consistent naive intertemporal decision making and an application to the case of uncertain lifetime. *The Review of Economic Studies*, *43*(3), 547–549.

Callen, M., Isaqzadeh, M., Long, J., & Sprenger, C. (2014). Violence and risk preference: Experimental evidence from afghanistan. *American Economic Review*, *104*(1), 123-148.

Cameron, L., & Shah, M. (2013). *Risk-taking behavior in the wake of natural disasters.* (NBER Working Paper #19534)

Carleton, T., Jina, A., Delgado, M., Greenstone, M., Houser, T., Hsiang, S., ... Zhang, A. T. (2022, April). Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. *The Quarterly Journal of Economics*, *137*(4), 2037–2105. Retrieved from `https://doi.org/10.1093/qje/qjac020` doi: 10.1093/qje/qjac020

Carrera, M., Royer, H., Stehr, M., Sydnor, J., & Taubinsky, D. (2022). Who chooses commitment? evidence and welfare implications. *Review of Economic Studies*, *89*, 1205–1244.

Cascio, W. E. (2018, May). Wildland fire smoke and human health. *Science of The Total Environment*, *624*, 586–595. Retrieved from `https://doi.org/10.1016/j.scitotenv.2017.12.086` doi: 10.1016/j.scitotenv.2017.12.086

Casey, J. A., Morello-Frosch, R., Mennitt, D. J., Fristrup, K., Ogburn, E. L., & James, P. (2017, July). Race/ethnicity, socioeconomic status, residential segregation, and spatial variation in noise exposure in the contiguous united states. *Environmental Health Perspectives*, *125*(7), 077017. Retrieved from `https://doi.org/10.1289/ehp898` doi: 10.1289/ehp898

Cassar, A., Healy, A., & Kessler, C. V. (2017). Trust, risk, and time preferences after a natural disaster: Experimental evidence from thailand. *World Development*, *94*, 90–05.

Castillo, M., Ferraro, P. J., Jordan, J. L., & Petrie, R. (2011). The today and tomorrow of kids: Time preferences and educational outcomes of children. *Journal of Public Economics*, *95*(11), 1377-1385. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0047272711001381` (Special Issue: International Seminar for Public Economics on Normative Tax Theory) doi: https://doi.org/10.1016/j.jpubeco.2011.07.009

Chandra, A., Kakani, P., & Sacarny, A. (2020, October). Hospital allocation and racial disparities in health care. *NBER Working Paper 28018*. Retrieved from `https://doi.org/10.3386/w28018` doi: 10.3386/w28018

Charness, G., Chemaya, N., & Trujano-Ochoa, D. (2023). Learning your own risk preferences. *Journal of Risk and Uncertainty*, *67*, 1–19.

Chen, M. K., Haggag, K., Pope, D. G., & Rohla, R. (2022, September). Racial disparities in voting wait times: Evidence from smartphone data. *The Review of Economics and Statistics*, 1–10. Retrieved from `https://doi.org/10.1162/rest_a_01012` doi: 10.1162/rest_a_01012

Chen, M. K., & Rohla, R. (2018, June). The effect of partisanship and political advertising on close family ties. *Science*, *360*(6392), 1020–1024. Retrieved from `https://doi.org/10.1126/science.aaq1433` doi: 10.1126/science.aaq1433

Chen, S., Chen, Y., Lei, Z., & Tan-Soo, J.-S. (2021, January). Chasing clean air: Pollution-induced travels in china. *Journal of the Association of Environmental and Resource Economists*, *8*(1), 59–89. Retrieved from `https://doi.org/10.1086/711476` doi: 10.1086/711476

Chen, S., Oliva, P., & Zhang, P. (2022, May). The effect of air pollution on migration: Evidence from china. *Journal of Development Economics*, *156*, 102833. Retrieved from `https://doi.org/10.1016/j.jdeveco.2022.102833` doi: 10.1016/j.jdeveco.2022.102833

Chetty, R., Hendren, N., Kline, P., & Saez, E. (2014, September). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics*, *129*(4), 1553–1623. Retrieved from `https://doi.org/10.1093/qje/qju022` doi: 10.1093/qje/qju022

Chetty, R., Hendren, N., Lin, F., Majerovitz, J., & Scuderi, B. (2016, May). Childhood environment and gender gaps in adulthood. *American Economic Review*, *106*(5), 282–288. Retrieved from `https://doi.org/10.1257/aer.p20161073` doi: 10.1257/aer.p20161073

Chetty, R., Stepner, M., Abraham, S., Lin, S., Scuderi, B., Turner, N., ... Cutler, D. (2016, April). The association between income and life expectancy in the united states, 2001-2014. *JAMA*, *315*(16), 1750. Retrieved from `https://doi.org/10.1001/jama.2016.4226` doi: 10.1001/jama.2016.4226

Chung, C., Myers, S. L., & Saunders, L. (2001, May). Racial differences in transportation access to employment in chicago and los angeles, 1980 and 1990. *American Economic Review*, *91*(2), 174–177. Retrieved from `https://doi.org/10.1257/aer.91.2.174` doi: 10.1257/aer.91.2.174

Clark, L. P., Harris, M. H., Apte, J. S., & Marshall, J. D. (2022, August). National and intraurban air pollution exposure disparity estimates in the united states: Impact of data-aggregation spatial scale. *Environmental Science & Technology Letters*, *9*(9), 786–791. Retrieved from `https://doi.org/10.1021/acs.estlett.2c00403` doi: 10.1021/acs.estlett.2c00403

Colmer, J., Hardman, I., Shimshack, J., & Voorheis, J. (2020, July). Disparities in $PM_{2.5}$ air pollution in the united states. *Science*, *369*(6503), 575–578. Retrieved from `https://doi.org/10.1126/science.aaz9353` doi: 10.1126/science.aaz9353

Currie, J. (2009, March). Healthy, wealthy, and wise: Socioeconomic status, poor health in childhood, and human capital development. *Journal of Economic Literature*, *47*(1), 87–122. Retrieved from `https://doi.org/10.1257/jel.47.1.87` doi: 10.1257/jel.47.1.87

Currie, J., & Schwandt, H. (2016a, May). Inequality in mortality decreased among the young while increasing for older adults, 1990–2010. *Science*, *352*(6286), 708–712. Retrieved from `https://doi.org/10.1126/science.aaf1437` doi: 10.1126/science.aaf1437

Currie, J., & Schwandt, H. (2016b, May). Mortality inequality: The good news from a county-level approach. *Journal of Economic Perspectives*, *30*(2), 29–52. Retrieved from `https://doi.org/10.1257/jep.30.2.29` doi: 10.1257/jep.30.2.29

Dave, D., Friedson, A. I., Matsuzawa, K., & Sabia, J. J. (2020, September). When do shelter-in-place orders fight covid-19 best? policy heterogeneity across states and adoption time. *Economic Inquiry*, *59*(1), 29–52. Retrieved from `https://doi.org/10.1111/ecin.12944` doi: 10.1111/ecin.12944

Davis, J. M. V., & Mazumder, B. (2022, January). The decline in intergenerational mobility after 1980. *Federal Reserve Bank of Chicago Working Paper 2014-28*. Retrieved from https://doi.org/10.2139/ssrn.3138979 doi: 10.2139/ssrn.3138979

Dedoussi, I. C., Eastham, S. D., Monier, E., & Barrett, S. R. H. (2020, February). Premature mortality related to united states cross-state air pollution. *Nature*, *578*(7794), 261–265. Retrieved from https://doi.org/10.1038/s41586-020-1983-8 doi: 10.1038/s41586-020-1983-8

DeJarnette, P. (2020). Temptation over time: Delays help. *Journal of Economic Behavior and Organization*, *177*, 752–761.

Deschênes, O. (2014, November). Temperature, human health, and adaptation: A review of the empirical literature. *Energy Economics*, *46*, 606–619. Retrieved from https://doi.org/10.1016/j.eneco.2013.10.013 doi: 10.1016/j.eneco.2013.10.013

Deschênes, O., & Greenstone, M. (2011, October). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics*, *3*(4), 152–185. Retrieved from https://doi.org/10.1257/app.3.4.152 doi: 10.1257/app.3.4.152

Deschênes, O., Greenstone, M., & Shapiro, J. S. (2017, October). Defensive investments and the demand for air quality: Evidence from the NOx budget program. *American Economic Review*, *107*(10), 2958–2989. Retrieved from https://doi.org/10.1257/aer.20131002 doi: 10.1257/aer.20131002

Dieleman, J. L., Chen, C., Crosby, S. W., Liu, A., McCracken, D., Pollock, I. A., ... Murray, C. J. L. (2021, August). US health care spending by race and ethnicity, 2002-2016. *JAMA*, *326*(7), 649. Retrieved from https://doi.org/10.1001/jama.2021.9937 doi: 10.1001/jama.2021.9937

Drouhin, N. (2020, January). Non-stationary additive utility and time consistency. *Journal of Mathematical Economics*, *86*, 1–14.

Eckel, C., El-Gamal, M., & Wilson, R. (2009). Risk loving after the storm: A bayesian-network study of hurricane katrina evacuees. *Journal of Economic Behavior and Organization*, *69*(2), 110–124.

Elder, T. E., Figlio, D. N., Imberman, S. A., & Persico, C. L. (2021, January). School segregation and racial gaps in special education identification. *Journal of Labor Economics*, *39*(S1), S151–S197. Retrieved from https://doi.org/10.1086/711421 doi: 10.1086/711421

Fajgelbaum, P. D., Khandelwal, A., Kim, W., Mantovani, C., & Schaal, E. (2021, December). Optimal lockdown in a commuting network. *American Economic Review: Insights*, *3*(4), 503–522. Retrieved from https://doi.org/10.1257/aeri.20200401 doi: 10.1257/aeri.20200401

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives*, *19*(4), 25–42.

Freeman, R., Liang, W., Song, R., & Timmins, C. (2019, March). Willingness to pay for clean air in china. *Journal of Environmental Economics and Management*, *94*, 188–216. Retrieved from https://doi.org/10.1016/j.jeem.2019.01.005 doi: 10.1016/j.jeem.2019.01.005

Glaeser, E. L., Gorback, C., & Redding, S. J. (2022, January). JUE insight: How much does COVID-19 increase with mobility? evidence from new york and four other u.s. cities. *Journal of Urban Economics*, *127*, 103292. Retrieved from `https://doi.org/10.1016/j.jue.2020.103292` doi: 10.1016/j.jue.2020.103292

Gneezy, A., & Imas, A. (2014). *Poverty traps and the effect of helplessness on impatience and risk seeking: Evidence from the lab and field.* (Working Paper)

Gneezy, U., & Potters, J. (1997). An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics*, *112*(2), 631–645. Retrieved 2024-05-01, from `http://www.jstor.org/stable/2951248`

Goolsbee, A., & Syverson, C. (2021, January). Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. *Journal of Public Economics*, *193*, 104311. Retrieved from `https://doi.org/10.1016/j.jpubeco.2020.104311` doi: 10.1016/j.jpubeco.2020.104311

Goss, M., Swain, D. L., Abatzoglou, J. T., Sarhadi, A., Kolden, C. A., Williams, A. P., & Diffenbaugh, N. S. (2020, August). Climate change is increasing the likelihood of extreme autumn wildfire conditions across california. *Environmental Research Letters*, *15*(9), 094016. Retrieved from `https://doi.org/10.1088/1748-9326/ab83a7` doi: 10.1088/1748-9326/ab83a7

Graff Zivin, J., & Neidell, M. (2009, September). Days of haze: Environmental information disclosure and intertemporal avoidance behavior. *Journal of Environmental Economics and Management*, *58*(2), 119–128. Retrieved from `https://doi.org/10.1016/j.jeem.2009.03.001` doi: 10.1016/j.jeem.2009.03.001

Graff Zivin, J., Neidell, M., & Schlenker, W. (2011, May). Water quality violations and avoidance behavior: Evidence from bottled water consumption. *American Economic Review*, *101*(3), 448–453. Retrieved from `https://doi.org/10.1257/aer.101.3.448` doi: 10.1257/aer.101.3.448

Halevy, Y. (2015). Time Consistency: Stationarity and Time Invariance. *Econometrica*, *83*(1), 335–352.

Hanaoka, C., Shigeoka, H., & Watanabe, Y. (2018). Do risk preferences change? evidence from the great east japan earthquake. *American Economic Journal: Applied Economics*, *10*(2), 298-330.

Harper, S., Lynch, J., Burris, S., & Smith, G. D. (2007, March). Trends in the black-white life expectancy gap in the united states, 1983-2003. *JAMA*, *297*(11), 1224. Retrieved from `https://doi.org/10.1001/jama.297.11.1224` doi: 10.1001/jama.297.11.1224

Harrison, G. W., Lau, M. I., & Williams, M. B. (2002). Estimating individual discount rates in denmark: A field experiment. *American economic review*, *92*(5), 1606–1617.

Harrison, G. W., Lau, M. I., & Yoo, H. I. (2023). *Constant discounting, temporal instability and dynamic inconsistency in denmark: A longitudinal field experiment.* (Working Paper)

Heft-Neal, S., Driscoll, A., Yang, W., Shaw, G., & Burke, M. (2022, January). Associations between wildfire smoke exposure during pregnancy and risk of preterm birth in california. *Environmental Research*, *203*, 111872. Retrieved from `https://doi.org/10.1016/j.envres.2021.111872` doi: 10.1016/j.envres.2021.111872

Holloway, M. S., Wiegand, C., & Kuhn, M. (2024). *Time-varying time preference.* (Working Paper)

Hornbeck, R., & Naidu, S. (2014, March). When the levee breaks: Black migration and economic development in the american south. *American Economic Review*, *104*(3), 963–990. Retrieved from `https://doi.org/10.1257/aer.104.3.963` doi: 10.1257/aer.104.3.963

Hsiang, S., Oliva, P., & Walker, R. (2019, January). The distribution of environmental damages. *Review of Environmental Economics and Policy*, *13*(1), 83–103. Retrieved from `https://doi.org/10.1093/reep/rey024` doi: 10.1093/reep/rey024

(2021, June). *American Journal of Health Economics*. Retrieved from `https://doi.org/10.1086/713787` doi: 10.1086/713787

Ifcher, J., & Zarghamee, H. (2011). Happiness and time preference: The effect of positive affect in a random-assignment experiment. *American Economic Review*, *101*(7), 3109–3129.

Ilin, C., Annan-Phan, S., Tai, X. H., Mehra, S., Hsiang, S., & Blumenstock, J. E. (2021, June). Public mobility data enables COVID-19 forecasting and management at local and global scales. *Scientific Reports*, *11*(1). Retrieved from `https://doi.org/10.1038/s41598-021-92892-8` doi: 10.1038/s41598-021-92892-8

Imas, A., Kuhn, M. A., & Mironova, V. (2022). Waiting to choose: The role of deliberation in intertemporal choice. *American Economic Journal: Microeconomics*, *14*(3), 414–40.

IPCC. (2022). *Global warming of 1.5°c.* Cambridge University Press. Retrieved from `https://doi.org/10.1017/9781009157940` doi: 10.1017/9781009157940

Ito, K., & Zhang, S. (2020, May). Willingness to pay for clean air: Evidence from air purifier markets in china. *Journal of Political Economy*, *128*(5), 1627–1672. Retrieved from `https://doi.org/10.1086/705554` doi: 10.1086/705554

Jamison, D. T., & Jamison, J. (2011, April). Characterizing the Amount and Speed of Discounting Procedures. *Journal of Benefit-Cost Analysis*, *2*(2), 1–56.

Janssens, W., Kramer, B., & Swart, L. (2017). Be patient when measuring hyperbolic discounting: Stationarity, time consistency and time invariance in a field experiment. *Journal of Development Economics*, *126*, 77–90.

Jaramillo, P., LaFave, D., & Novak, L. (2023). *Extreme weather events impact risk tolerance and time preferences.* (Working Paper)

Jay, J., Bor, J., Nsoesie, E. O., Lipson, S. K., Jones, D. K., Galea, S., & Raifman, J. (2020, November). Neighbourhood income and physical distancing during the COVID-19 pandemic in the united states. *Nature Human Behaviour*, *4*(12), 1294–1302. Retrieved from `https://doi.org/10.1038/s41562-020-00998-2` doi: 10.1038/s41562-020-00998-2

Kettlewell, N., Rijsdijk, F., Siribaddana, S., Sumathipala, A., Tymula, A., Zavos, H., & Glozier, N. (2018). *Civil war, natural disaster and risk preferences: Evidence from sri lankan twins.* (IZA Discussion Paper No. 11901)

Khanna, G., Liang, W., Mobarak, A. M., & Song, R. (2021, January). The productivity consequences of pollution-induced migration in china. *NBER Working Paper 28401*. Retrieved from `https://doi.org/10.3386/w28401` doi: 10.3386/w28401

Kibris, A., & Uler, N. (2021). *The impact of exposure to armed conflict on risk and ambiguity attitudes.* (Working Paper)

Kochi, I., Champ, P. A., Loomis, J. B., & Donovan, G. H. (2012, January). Valuing mortality impacts of smoke exposure from major southern california wildfires. *Journal of Forest Economics*, *18*(1), 61–75. Retrieved from `https://doi.org/10.1016/j.jfe.2011.10.002` doi: 10.1016/j.jfe.2011.10.002

Kondo, M., Roos, A. D., White, L., Heilman, W., Mockrin, M., Gross-Davis, C., & Burstyn, I. (2019, March). Meta-analysis of heterogeneity in the effects of wildfire smoke exposure on respiratory health in north america. *International Journal of Environmental Research and Public Health*, *16*(6), 960. Retrieved from `https://doi.org/10.3390/ijerph16060960` doi: 10.3390/ijerph16060960

Kuhn, M. A., Kuhn, P., & Villeval, M. C. (2017). Decision-environment effects on intertemporal financial choices: How relevant are resource-depletion models? *Journal of Economic Behavior and Organization*, *137*, 72–89.

Kuroishi, Y., & Sawada, Y. (2019). *On the stability of preferences: Experimental evidence from two disasters.* (CREPE DISCUSSION PAPER NO. 61)

Lai, W., Li, S., Liu, Y., & Barwick, P. J. (2022, March). Adaptation mitigates the negative effect of temperature shocks on household consumption. *Nature Human Behaviour*, *6*(6), 837–846. Retrieved from `https://doi.org/10.1038/s41562-022-01315-9` doi: 10.1038/s41562-022-01315-9

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, *112*(2), 443-477.

Lennon, C. (2022). Property and suppression costs. *Working Paper*.

Liu, J. C., Wilson, A., Mickley, L. J., Dominici, F., Ebisu, K., Wang, Y., … Bell, M. L. (2017, January). Wildfire-specific fine particulate matter and risk of hospital admissions in urban and rural counties. *Epidemiology*, *28*(1), 77–85. Retrieved from `https://doi.org/10.1097/ede.0000000000000556` doi: 10.1097/ede.0000000000000556

Long, E. F., Chen, M. K., & Rohla, R. (2020, 9). Political storms: Emergent partisan skepticism of hurricane risks. *Science Advances*, *6*(37). Retrieved from `https://doi.org/10.1126/sciadv.abb7906` doi: 10.1126/sciadv.abb7906

Malmendier, U., & Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk-taking? *Quarterly Journal of Economics*, *126*(1), 373-416.

Manson, S., Schroeder, J., Van Riper, D., Kugler, T., & Ruggles, S. (2022). *Ipums national historical geographic information system: Version 17.0 [dataset].* Minneapolis, MN. doi: http://doi.org/10.18128/D050.V17.0

Massetti, E., & Mendelsohn, R. (2018, July). Measuring climate adaptation: Methods and evidence. *Review of Environmental Economics and Policy*, *12*(2), 324–341. Retrieved from `https://doi.org/10.1093/reep/rey007` doi: 10.1093/reep/rey007

Matousek, J., Havranek, T., & Irsova, Z. (2022). Individual discount rates: a meta-analysis of experimental evidence. *Experimental Economics*, *25*(1), 318–358.

Miller, N., Molitor, D., & Zou, E. (2017). Blowing smoke: Health impacts of wildfire plume dynamics. *Working Paper*.

Mohai, P., & Saha, R. (2007). Racial inequality in the distribution of hazardous waste: A national-level reassessment. *Social Problems*, *54*(3), 343–370. Retrieved 2022-10-05, from `http://www.jstor.org/stable/10.1525/sp.2007.54.3.343`

Moretti, E., & Neidell, M. (2011). Pollution, health, and avoidance behavior evidence from the ports of los angeles. *Journal of Human Resources*, *46*(1), 154–175.

Neidell, M. (2009). Information, avoidance behavior, and health the effect of ozone on asthma hospitalizations. *Journal of Human Resources*, *44*(2), 450–478.

NOAA Office of Satellite and Product Operations. (2022). *Hazard mapping system.* Retrieved 2022-09-12, from `https://www.ospo.noaa.gov/Products/land/hms.html#about`

O'Dell, K., Ford, B., Fischer, E. V., & Pierce, J. R. (2019, January). Contribution of wildland-fire smoke to US $PM_{2.5}$ and its influence on recent trends. *Environmental Science & Technology*, *53*(4), 1797–1804. Retrieved from `https://doi.org/10.1021/acs.est.8b05430` doi: 10.1021/acs.est.8b05430

O'Donoghue, T., & Rabin, M. (1999). Doing it now or later. *American Economic Review*, *89*(1), 103–124.

Page, L., Savage, D. A., & Torgler, B. (2014). ariation in risk seeking behaviour following large losses: A natural experiment. *European Economic Review*, *71*, 121–131.

Parolin, Z., & Lee, E. K. (2021, March). Large socio-economic, geographic and demographic disparities exist in exposure to school closures. *Nature Human Behaviour*, *5*(4), 522–528. Retrieved from `https://doi.org/10.1038/s41562-021-01087-8` doi: 10.1038/s41562-021-01087-8

Plantinga, A. J., Walsh, R., & Wibbenmeyer, M. (2022, July). Priorities and effectiveness in wildfire management: Evidence from fire spread in the western united states. *Journal of the Association of Environmental and Resource Economists*, *9*(4), 603–639. Retrieved from `https://doi.org/10.1086/719426` doi: 10.1086/719426

Prelec, D. (2004). Decreasing impatience: A criterion for non-stationary time preference and "hyperbolic" discounting. *The Scandinavian Journal of Economics*, *106*(3), 511-532.

Rabin, M. (2013). An approach to incorporating psychology into economics. *American Economic Review*, *103*(3), 617–622.

Radeloff, V. C., Hammer, R. B., Stewart, S. I., Fried, J. S., Holcomb, S. S., & McKeefry, J. F. (2005, June). The wildland-uban interface in the united states. *Ecological Applications*, *15*(3), 799–805. Retrieved from https://doi.org/10.1890/04-1413 doi: 10.1890/04-1413

Radeloff, V. C., Helmers, D. P., Kramer, H. A., Mockrin, M. H., Alexandre, P. M., Bar-Massada, A., ... Stewart, S. I. (2018, March). Rapid growth of the US wildland-urban interface raises wildfire risk. *Proceedings of the National Academy of Sciences*, *115*(13), 3314–3319. Retrieved from https://doi.org/10.1073/pnas.1718850115 doi: 10.1073/pnas.1718850115

Read, D., & van Leeuwen, B. (1998). Predicting hunger: The effects of appetite and delay on choice. *Organizational Behavior and Human Decision Processes*, *76*(2), 189-205.

Reynaud, A., & Aubert, C. (2019). Does flood experience modify risk preferences? evidence from an artefactual field experiment in vietnam. *The Geneva Risk and Insurance Review*, *45*, 36–74.

Richardson, L. A., Champ, P. A., & Loomis, J. B. (2012, January). The hidden cost of wildfires: Economic valuation of health effects of wildfire smoke exposure in southern california. *Journal of Forest Economics*, *18*(1), 14–35. Retrieved from https://doi.org/10.1016/j.jfe.2011.05.002 doi: 10.1016/j.jfe.2011.05.002

SafeGraph. (2022a). *Patterns: Privacy.* Retrieved 2022-09-14, from https://docs.safegraph.com/docs/monthly-patterns#section-privacy

SafeGraph. (2022b). *Weekly patterns.* Retrieved 2022-09-14, from https://docs.safegraph.com/docs/weekly-patterns

Schwandt, H., Currie, J., Bär, M., Banks, J., Bertoli, P., Bütikofer, A., ... Wuppermann, A. (2021, September). Inequality in mortality between black and white americans by age, place, and cause and in comparison to europe, 1990 to 2018. *Proceedings of the National Academy of Sciences*, *118*(40). Retrieved from https://doi.org/10.1073/pnas.2104684118 doi: 10.1073/pnas.2104684118

Sepahvand, M. H., Shahbazian, R., & Swain, R. B. (2019). *Does revolution change risk attitudes? evidence from burkina faso.* (Working Paper)

Singh, M., & Venkataramani, A. (2022, August). Capacity strain and racial disparities in hospital mortality. *NBER Working Paper 30380*. Retrieved from https://doi.org/10.3386/w30380 doi: 10.3386/w30380

Smith, A. B. (2022). *U.s. billion-dollar weather and climate disasters.* NOAA National Centers for Environmental Information. Retrieved from https://www.ncei.noaa.gov/archive/accession/0209268 doi: 10.25921/STKW-7W73

Squire, R. F. (2019a). *Measuring and correcting sampling bias in safegraph patterns for more accurate demographic analysis.* Retrieved 2022-09-14, from https://www.safegraph.com/blog/measuring-and-correcting-sampling-bias-for-accurate-demographic-analysis

Squire, R. F. (2019b). *What about bias in the safegraph dataset?* Retrieved 2022-09-14, from https://colab.research.google.com/drive/1u15afRytJMsizySFqA2EPlXSh3KTmNTQ

Storey, J. D. (2011). False discovery rate. *International encyclopedia of statistical science*, *1*, 504–508.

Strotz, R. H. (1956). (1955-1956). myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies.-1955-1956*, *23*(3), 165.

Strulik, H. (2021). Hyperbolic discounting and the time-consistent solution of three canonical environmental problems. *Journal of Public Economic Theory*, *23*(3), 462–486.

Sun, C., Kahn, M. E., & Zheng, S. (2017, January). Self-protection investment exacerbates air pollution exposure inequality in urban china. *Ecological Economics*, *131*, 468–474. Retrieved from https://doi.org/10.1016/j.ecolecon.2016.06.030 doi: 10.1016/j.ecolecon.2016.06.030

Tai, X. H., Mehra, S., & Blumenstock, J. E. (2022, May). Mobile phone data reveal the effects of violence on internal displacement in afghanistan. *Nature Human Behaviour*, *6*(5), 624–634. Retrieved from https://doi.org/10.1038/s41562-022-01336-4 doi: 10.1038/s41562-022-01336-4

Takeuchi, K. (2011). Non-parametric test of time consistency: Present bias and future bias. *Games and Economic Behavior*, *71*(2), 456-478.

Tan Soo, J.-S. (2018). Valuing air quality in indonesia using households' locational choices. *Environmental and Resource Economics*, *71*(3), 755–776.

Thompson, S. A., & Warzel, C. (n.d.). One nation, tracked. *The New York Times*. Retrieved from https://www.nytimes.com/interactive/2019/12/19/opinion/location-tracking-cell-phone.html

Thompson, Y., Stuart A. Serkez, & Kelley, L. (n.d.). How has your state reacted to social distancing? *The New York Times*. Retrieved from https://www.nytimes.com/interactive/2020/03/23/opinion/coronavirus-economy-recession.html

USDA, E. R. S. (2012). U.s. adult food security survey module: Three-stage design, with screeners.

Valentino-Devries, J., Singer, N., Keller, M. H., & Krolik, A. (n.d.). Your apps know where you were last night, and they're not keeping it secret. *The New York Times*. Retrieved from https://www.nytimes.com/interactive/2018/12/10/business/location-data-privacy-apps.html

Voors, M., Nielsen, E., Verwimp, P., Bulte, E., Lensink, R., & Soest, D. V. (2012). Violent conflict and behavior: A field experiment in burundi. *American Economic Review*, *102*(2), 941-964.

Wasserman, L. (2000). Bayesian model selection and model averaging. *Journal of mathematical psychology*, *44*(1), 92–107.

Weill, J. A., Stigler, M., Deschenes, O., & Springborn, M. R. (2020, July). Social distancing responses to COVID-19 emergency declarations strongly differentiated by income. *Proceedings of the National Academy of Sciences*, *117*(33), 19658–19660. Retrieved from `https://doi.org/10.1073/pnas.2009412117` doi: 10.1073/pnas.2009412117

Wen, J., & Burke, M. (2021). *Wildfire smoke plume segmentation using geostationary satellite imagery.* arXiv. Retrieved from `https://arxiv.org/abs/2109.01637` doi: 10.48550/ARXIV.2109.01637

Wen, J., & Burke, M. (2022, September). Lower test scores from wildfire smoke exposure. *Nature Sustainability*. Retrieved from `https://doi.org/10.1038/s41893-022-00956-y` doi: 10.1038/s41893-022-00956-y

Westerling, A. L., Hidalgo, H. G., Cayan, D. R., & Swetnam, T. W. (2006, August). Warming and earlier spring increase western u.s. forest wildfire activity. *Science*, *313*(5789), 940–943. Retrieved from `https://doi.org/10.1126/science.1128834` doi: 10.1126/science.1128834

Wibbenmeyer, M., & Robertson, M. (2022, May). The distributional incidence of wildfire hazard in the western united states. *Environmental Research Letters*, *17*(6), 064031. Retrieved from `https://doi.org/10.1088/1748-9326/ac60d7` doi: 10.1088/1748-9326/ac60d7

Wildland Fire Interagency Geospatial Service. (2022). *Wildland fire perimeters full history.* Retrieved 2022-09-12, from `https://data-nifc.opendata.arcgis.com/datasets/wfigs-wildland-fire-perimeters-full-history/`

Zhang, J., & Mu, Q. (2018, November). Air pollution and defensive expenditures: Evidence from particulate-filtering facemasks. *Journal of Environmental Economics and Management*, *92*, 517–536. Retrieved from `https://doi.org/10.1016/j.jeem.2017.07.006` doi: 10.1016/j.jeem.2017.07.006

Zhao, M., Holtz, D., & Aral, S. (2021, July). Interdependent program evaluation: Geographic and social spillovers in COVID-19 closures and reopenings in the united states. *Science Advances*, *7*(31). Retrieved from `https://doi.org/10.1126/sciadv.abe7733` doi: 10.1126/sciadv.abe7733

Zheng, S., & Kahn, M. E. (2008, March). Land and residential property markets in a booming economy: New evidence from beijing. *Journal of Urban Economics*, *63*(2), 743–757. Retrieved from `https://doi.org/10.1016/j.jue.2007.04.010` doi: 10.1016/j.jue.2007.04.010