Title

Essays in Psychology and Economics

Author

Wellsjo, Alexandra Steiny

Publication Date

2020

Essays in Psychology and Economics

by

Alexandra Steiny Wellsjo

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Ulrike Malmendier, Co-chair
Professor Stefano DellaVigna, Co-chair
Professor Jonathan Kolstad

Fall 2020

Essays in Psychology and Economics

Abstract

Essays in Psychology and Economics

by

Alexandra Steiny Wellsjo

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Ulrike Malmendier, Co-chair

Professor Stefano DellaVigna, Co-chair

This dissertation studies how the past influences our future choices. Leveraging insights from psychology, I study the role of past actions in forming habits and the role of past experiences in altering beliefs and preferences. I explore implications of these findings for individuals, firms, and the macroeconomy.

In the first chapter, I ask - how should organizations incentivize routine behaviors? I use a cue-based model of habit formation to describe hand hygiene behavior in hospitals and suggest implications for intervention to change behavior. Hand hygiene is the best way to prevent healthcare-associated infections which affect millions of patients and cost billions of dollars each year. However, many studies find that healthcare workers comply with hand hygiene guidelines less than half of the time. Leveraging a unique dataset that tracks compliance behavior of 13,606 workers in 123M hand hygiene opportunities, I document patterns of behavior consistent with a cue-based theory of habit. High compliers benefit from automatic hand washing and are buffered from the negative impacts of cognitive load. However, habits are broken when high compliers are taken out of their normal routines. The findings have implications for how hospitals should design, target, and implement hand hygiene interventions.

In the second chapter, in joint work with Ulrike Malmendier, we ask whether the past can explain some of the vast differences in homeownership rates across countries and cohorts. We argue that macroeconomic histories individuals have personally experienced in their home country over the past decades play a significant role in their decision to become a homeowner. We present novel survey data that reveals inflation protection as a key motivation to buy rather than rent a home. We then show that higher exposure to inflation predicts home ownership, controlling for other known determinants of homeownership as well as any differences across countries and over time. First, we find that immigrants' country-of-origin inflation ex-

periences strongly predict their homeownership rates in the US. Second, analyzing data from 22 European countries, we estimate that a one log-point increase in personally experienced inflation predicts a 18 pp increase in the average individual's likelihood of homeownership. As predicted by our model, the effects are strongest in countries with predominantly fixed-rate mortgages.

To Papa

Contents

List of Figures

List of Tables

Acknowledgments

First and foremost, I would like to thank Ulrike Malmendier and Stefano DellaVigna, whose advising began long before this dissertation. Ulrike, becoming your undergrad research assistant truly changed the trajectory of my life. Thank you for being such an inspiration and role model over the last 10 years. Stefano, thank you for being such a thoughtful and caring advisor. From the small details to the big picture, you have done so much to support me and my work at every level. I would also like to thank Jonathan Kolstad, Ned Augenblick, and Dmitry Taubinsky for the many insightful comments and thoughtful conversations. I am grateful to have been advised by a team of not only amazing scholars but incredible advisors, who have been so generous with their time and effort. Thank you all for the tremendous amount of support and guidance.

I have had many other great mentors over the years who helped to build the foundation for this dissertation. I would especially like to thank David Laibson, Dan Benjamin, Justin Sydnor, Matthew Rabin, Ricardo Perez-Truglia, David Sraer, John Beshears, James Choi, and Brigitte Madrian.

I am also incredibly thankful for the friendships and advice from my UC Berkeley peers. I am especially grateful for my first year study group, my officemates, and the psychology and economics group. In particular, I want to thank Peter Jones and Avner Strulov-Shlain for the countless hours of conversation that have turned into what I hope will be lifelong collaborations and friendships.

I would like to thank SwipeSense for graciously providing access to their electronic monitoring data, without which this dissertation would not have been possible.

And finally, I would like to thank my family, who have always been my biggest cheerleaders. To my parents, Lisa and Richard, without your love and support, I could never be where I am today. And to Sebastian, thank you for always believing in me and for pushing me to achieve my dreams. I am beyond grateful to have you in my life.

Introduction

This dissertation studies how the past influences our future choices.

In Chapter 1, “Simple Actions, Complex Habits: Lessons from Hospital Hand Hygiene,” I explore the role of habit formation in hand hygiene behavior of hospital personnel. I document evidence consistent with habit formation, in which past hand washing predicts future hand washing. I discuss the implications of habit formation for individual behavior and for organizations who aim to increase compliance with hand hygiene guidelines.

In Chapter 2, “Rent or Buy? Inflation Experiences and Homeownership within and across Countries,” joint with Ulrike Malmendier, we study the role of past macroeconomic experiences in shaping future choices. Specifically, we argue that macroeconomic histories individuals have personally experienced play a significant role in their decision to become a homeowner. We argue that differences in past experiences may contribute to the large variation we see in homeownership across countries and cohorts.

Chapter 1

Simple Actions, Complex Habits: Lessons from Hospital Hand Hygiene

1.1 Introduction

People develop habits in many situations, when they are not making thoughtful choices but instead rely on an automatic response. Studies estimate that 45% of our daily activities are routine, repeated at the same time and place each day (Neal et al. (2006)). Within organizations, these routine behaviors can be critical to performance. For example, Toyota's highly efficient production system is built on standardized tasks and processes (Suárez-Barraza et al. (2008)). In industries like manufacturing, aviation, and medicine, failure to follow safety routines can have fatal consequences.

So how should organizations incentivize these routine behaviors? By nature, monitoring of these behaviors is often difficult or too costly to enforce with standard contracts. Under a simplistic view of habits, one solution is to monitor and provide incentives for a short period of time. While costly in the short-term, if habits are formed, there will be a persistent effect even after incentives are removed. While there are some successful examples of this type of intervention (e.g., Charness and Gneezy (2009)), there are many examples in the literature of short-run interventions with minimal long-run impacts (e.g., Acland and Levy (2015), Carrera et al. (2018), Royer et al. (2015)).

In this paper, I argue that some of these results may be explained by subtleties in the habit formation process that are not captured in the typical economics framework of rational addiction (i.e., Becker and Murphy (1988)). Closer to the models of Laibson (2001), Bernheim and Rangel (2004), and Taubinsky (2014), I describe a theoretical framework that appeals directly to the cognitive science literature on habit formation as an automatic cue-response process.[1] Like other models of habit formation, experience doing an action raises the marginal utility for taking the action in the future. As in Laibson (2001), this process is cue-specific. The model differs in explicitly modeling the automaticity of habit by differenti-

[1]See, for example, Wood et al. (2014) for a review of habits in dual-process models.

ating attentive from habitual action. Rather than cues increasing cravings for consumption (as in Laibson's example of smelling warm bread), cues trigger an automatic and costless action. In order to overcome a habit, one must exert costly attention.

For a desired action, this results in a tradeoff between habits and attention. As automatic habits form, one can achieve the same level of activity while paying less attention. Because habitual actors behave automatically, they are not affected by changes in incentives or costs of attention. On the other hand, attentive actors will be responsive to these changes. This model has implications for how organizations should target, design, and implement interventions to encourage routine behaviors.

I study this question in the context of compliance with hand hygiene guidelines in hospitals, which require healthcare workers to wash or sanitize their hands each time they enter and exit a patient room.[2] This routine behavior is highly important; proper hand hygiene by healthcare providers and hospital staff is one of the most powerful tools we have to reduce healthcare-associated infections (HAI) (Jarvis (1994)). HAI are a critical public health issue; at any point in time, 1 in 25 hospital patients are affected, resulting in approximately 100k deaths and $35 billion in hospital costs each year (Klevens et al. (2007), Scott (2009), ODPHP (2019)). Since the start of the COVID-19 pandemic, hand hygiene in hospitals is arguably more important now than ever. In addition to the medical benefits and direct costs of HAI, hospitals also have large financial and reputation incentives to keep HAI low and compliance with hand hygiene guidelines high.

Despite efforts by hospitals and public health organizations to achieve high compliance, many studies find that healthcare providers wash their hands less than half the time they should (World Health Organization (2009)). Motivating hand hygiene has proven challenging; almost 200 years since the benefits of hand hygiene were first discovered, behavior change remains elusive. Many interventions are unsuccessful or have only short-lived results. Those that are successful are often multi-dimensional, making it hard to pinpoint the cause of the effects (Naikoba and Hayward (2001); Kingston et al. (2016); Gould et al. (2017)).

Growing out of this demand for higher compliance, several companies have developed technologies to electronically monitor workers' hand hygiene compliance in each patient interaction. I use data from one such company, SwipeSense, to study the hand hygiene behavior of 13,606 workers in 123M hand hygiene opportunities.

Electronic monitoring of hand hygiene represents a substantial improvement over direct observation, the "gold standard" for determining compliance rates. Built on a location-tracking system, the technology tracks each hand hygiene opportunity, or an entry to or exit from a patient room. Compliance is automatically and objectively recorded if a worker uses a soap or sanitizer dispenser within a predetermined amount of time (e.g., 30 seconds) of room entry or exit. The data generated by electronic monitoring provides a unique lens into the habit formation process for several reasons.

First, compliance is an objective, frequent, and well-measured aspect of performance. In

[2] In the remainder of the paper, I use the term washing to refer both to washing with soap and water and to the use of hand sanitizer gel.

a typical 12-hour shift, the average worker in my sample makes 34 visits to patient rooms for a total of 68 hand hygiene opportunities on room entry and exit. The nature of these interactions may vary substantially across patients, workers, and situations. For example, patient interactions may include a nurse administering medication, a technician taking vital measurements, or a surgeon providing a post-operative update. Across all of the situations and for all workers, hospital guidelines and measurement are clear and consistent: hands should be washed on entry to and exit from the patient room. This standardization allows for comparisons across people and situations in a way that is impossible in many other domains.

Second, the required timing of hand washing allows for a very specific link between the context (or cue) and whether the action was performed. For each opportunity, the data includes the specific timestamp of room entry or exit and a room-specific id. If there are physical cues that differ across rooms or departments, we can observe behavior for the same person across a variety of cues. This is much harder in other settings like gym attendance where cues, timing, and context are often less clearly defined. In one attempt to address this, Beshears et al. (2020) incentivize gym attendance within a narrow time window to facilitate formation of a time-based routine. The authors (and 77% of surveyed psychologists) anticipated that this would be more successful than a flexible incentive, but found the opposite effect. The surprising results point to the complexity of cues and the need for a better understanding of cue-based habit formation in the field.

Finally, the workers in this paper are experts in hand hygiene. The majority have likely received formal training on hand hygiene in nursing school, medical school, or another professional health education. The high number of opportunities means that even for the few workers who have not received formal training on hand hygiene before they start, they likely quickly catch up with on-the-job training. This differentiates the setting from others in which habit formation may be happening in parallel with real learning about the costs and benefits of taking a new action or consuming a new good.

Using the electronic monitoring data, I describe 5 empirical facts consistent with the assumptions and predictions of the cue theory model of habits.

Consistent with hand hygiene behavior operating as costly attention and automatic habit, I first document that hand hygiene responds to both changes in cognitive load and habit. First, I show that cognitive load, specifically fatigue and distraction, lower compliance. Similar to the findings of Dai et al. (2015), I find that as workers fatigue over a 12-hour shift, compliance falls by 6pp. As another source of fatigue, I also find that compliance is about 1.5pp lower when working night shifts. I also find that compliance is about 0.5pp lower on weekends, when there may be more distractions outside of work, higher stress from lower staffing, or higher levels of fatigue from staying out late. Second, I examine the impact of shocks to habit. Under the assumption that habit is location specific and decays over time, there are two events in the data that we might expect to disrupt habit: taking a long break from work (e.g., a 2 week vacation) or covering a shift in a different department (i.e., a float day). Indeed, I find that compliance is 2pp lower following a long break and 3pp lower on float days.

The third and fourth facts document the key predictions of the model, that attentive and habitual hand washers respond differently to shocks. Consistent with the model, I find that low compliers are responsive to cognitive load. As fatigue and distraction make it costlier to pay attention to hand washing, their compliance declines. On the other hand, the habits of high compliers allow them to wash automatically, making them unresponsive to cognitive load. I find that high compliers fatigue over the shift at about half the rate as low compliers and have no significant drop in compliance on nights or weekends. However, the habits that benefit the high types can be shocked by disruptions to their routines. When working in a different department or after taking a long vacation, high compliers do respond with a drop in compliance just as large as their low compliance peers.

In the fifth fact, I document evidence consistent with a location-specific cue-based habit formation process. I show that compliance is increasing with experience, and specifically with experience in the same physical location. Consistent with location-specific cues, compliance is higher in rooms that have been visited more often.

These findings differ from much of the existing literature on habit formation in the field by focusing on the role of cues and automaticity rather than tests of the long-run effects of short-term incentives (e.g., Charness and Gneezy (2009), Acland and Levy (2015), Royer et al. (2015), Loewenstein et al. (2016), Carrera et al. (2018), Harris and Kessler (2019)) or tests of rational addiction models (e.g., Hussam et al. (2017), Gruber and Köszegi (2001), Becker et al. (1994)).

Accounting for the automatic processes at play in routine behaviors can help us to better describe and predict individual behavior at a micro-level and can have important implications for individuals and organizations seeking behavior change.

For hospitals who aim to increase compliance, this research has several policy implications. First, the heterogeneity across individuals suggests potential gains from targeting interventions. I document large differences in both levels and elasticities of compliance across individuals. Taking these into account, hospitals can target interventions to people and situations in which we expect them to have low compliance.

Second, the habit formation model provides guidance on the types of interventions that hospitals should run. Many hand hygiene campaigns are "attention-grabbing," such as an educational session designed to remind healthcare workers about the importance of hand washing. This likely doesn't provide new information, but instead reminds workers about their organization's commitment to hand hygiene. These types of interventions may increase compliance, but will do so by encouraging workers to pay costly attention to hand hygiene. The short-run welfare costs may be negligible if these interventions help to build habits, but the lack of persistent effects in the literature suggests this is often not the case.

The link between compliance and costly attention highlights another interesting use of hand hygiene data as a measure of cognitive load. For example, consider a hospital that is changing their electronic health records system to one that they hope will make patient care more efficient. Does learning the new technology cause a short run drain on attention that, like fatigue, lowers compliance? In the long-run, does compliance improve in response to the more efficient system? Using compliance as a measure of cognitive load among workers, hospitals can quantify the distraction caused by administrative tasks, changing protocols, or learning new technologies.

Third, while focusing attention on hand hygiene is costly, alternative interventions designed to lower the costs of compliance would raise both compliance and worker welfare. Through systemic changes, these interventions could be maintained over a long time frame, allowing workers the opportunity to build habit. Standardization of room layouts and location of sanitizer dispensers, for example, could help to facilitate habit formation and the portability of habits across locations. Concentrating work in a smaller set of rooms may allow for faster habit formation. Taking advantage of heterogeneity across individuals, compliance could be increased by allocating shifts according to workers' strengths. The ability for these types of interventions to be sustained over the long term could help to maintain high compliance.

This paper is not the first to suggest solutions for organizations who attempt to improve performance in routine behaviors. A long literature on organizational routines highlight the benefits of coordination and standardization in an organization's processes and responses (Becker (2004)). This paper is more closely related to the subset of those that focus on changes to individual behavior. For example, checklists have been demonstrated to be highly effective in improving surgical outcomes and reducing airplane accidents (Haynes et al. (2009); Guwande (2011)). This paper contributes to this literature by improving our understanding of cue-based habit formation and the role it plays in designing successful routines.

The paper proceeds as follows. Section 1.2 describes the cue theory of habit and predictions of the model. Section 1.3 describes the electronic monitoring data. Section 1.4 provides empirical evidence consistent with the predictions in Section 1.2. Section 1.5 discusses the implications for organizations. Section 1.6 concludes.

1.2 Theoretical Framework

In this section, I describe a theoretical framework of habit formation in hand hygiene. Drawing from the psychological literature on the cue theory of habits, habit formation involves the creation of an automatic cue-response process. In situations when a habit has formed, hand washing is an automatic and costless response to a cue (i.e., entering or exiting a room). Because the cue-response process is non-cognitive, it is unaffected by cognitive load (e.g., fatigue and distraction). In situations when habits have not formed, remembering to wash requires costly attention and is more difficult under cognitive load. This model makes several predictions that are supported by the empirical evidence.

Cue Theory of Habits in Hand Hygiene

Consider a healthcare worker who should wash in opportunity o, characterized by the time t and the location l. Washing has a fixed benefit b, which I normalize to 1.

In each opportunity, the worker chooses how much attention, a, to pay to hand washing, conditional on her cognitive load at the time, d_t, and her level of automaticity, or habit h_{lt}, which depends on the location and time.

According to cue theory, habits are developed by repeated linking of a cue (i.e., entering or exiting a patient room) with an action (i.e., washing or sanitizing). Therefore, habits should be increasing in the number of washes in response to the cue. We also expect habits to decay if not reinforced, so that recent washes will increase habit more than washes in the distant past. If cues vary across physical locations, habits will also be location-specific. Specifically, I assume habit evolves automatically according to

$$h_{lt} = \gamma h_{lt-1} + (1-\gamma)\delta(l,l')w_{l't-1}$$

with $\gamma \in [0,1)$ and $h_0 \in [0,1]$. $w_{l't-1}$ is washing behavior in the last period, equal to 1 if the worker performed hand hygiene and 0 if the worker was not compliant or did not have an opportunity in $t-1$. Habit in location l at time t is a weighted average of the location-specific habit stock and hand hygiene behavior in the last period. $\delta(l,l') \in [0,1]$ captures the location-specificity of habit. If habits are not location-specific $\delta(l,l') = 1\ \forall l,l'$. If habits are location-specific, $\delta(l,l) = 1$ and $\delta(l,l') < 1\ \forall l \neq l'$. Based on the psychology, we should expect that locations with a higher degree of similarity in cues would have a δ closer to 1. With $\gamma < 1$, habits (and the importance of past washes) decay over time.

Given her level of automaticity, or habit h_{lt}, and the other drains on attention at the time, d_t, the worker chooses how much costly attention a to allocate to hand washing at time t. Let a translate to a probability of paying attention, such that with full attention $a = 1$, she will always wash. With no attention $a = 0$, she will wash according to her habit level h_{lt}. At attention level a, she will wash with probability $p(a,h_{lt}) = a(1-h_{lt}) + h_{lt}$.

Attention is costly; I assume cost of attention $c(a,d_t) = \frac{d_t a}{1-a}$, where d_t is the cognitive load or distractions at time t. Thus, costs are convex in attention. Cognitive load increases the marginal cost of attention. The cost of attention approaches infinity as $a \to 1$, guaranteeing a solution $a \in [0,1)$.[3]

The healthcare worker chooses attention to maximize expected utility or

$$E[U(a,h_{lt})] = p(a,h_{lt})b - c(a,d_t).$$

Note, this solution assumes the worker is myopic when selecting her attention level and does not anticipate the impact of today's behavior on future habit levels. Without loss of generality, I normalize $b = 1$.

If interior, the optimal solution, a^*, satisfies

$$\frac{\partial c(a^*,d_t)}{\partial a} = 1 - h_{lt}$$

with probability of compliance $p(a^*,h_{lt}) = a^*(1-h_{lt}) + h_{lt}$.

[3]See A.1 for results under a more general cost function.

With cost $c(a, d_t) = \frac{d_t a}{1-a}$, the worker chooses attention

$$a^* = 1 - \sqrt{\frac{d_t}{1 - h_{lt}}} \tag{1.1}$$

resulting in a probability of compliance of

$$p(a^*, h_{lt}) = 1 - \sqrt{1 - h_{lt}} \tag{1.2}$$

when $\sqrt{\frac{d_t}{1-h_{lt}}} < 1$ and $a^* = 0$, $p(a^*, h_{lt}) = h_{lt}$ otherwise. This yields our first set of predictions.

Prediction 1. *Compliance is decreasing in the level of distraction, d_t.*

Proof. The probability of compliance is decreasing in distractions as

$$\frac{\partial p(a^*, h_{lt})}{\partial d_t} = -\frac{1}{2}\sqrt{\frac{1 - h_{lt}}{d}} \leq 0.$$

□

Prediction 2. *Compliance is increasing in habit, h_{lt}, and thus decreasing in any disruptions to habit.*

Proof. Compliance is increasing in habit as

$$\begin{aligned}\frac{\partial p(a^*, h_{lt})}{\partial h_{lt}} &= \frac{1}{2}\sqrt{\frac{d_t}{1 - h_{lt}}} \\ &\geq 0.\end{aligned}$$

□

Prediction 3. *If $\delta(l, l') < 1$ when $l \neq l'$, habit formation will be location-specific. Compliance will be higher in locations with more (or more recent) past hand washing.*

Proof. By assumption, if $\delta(l, l') < 1$, habit formation is location-specific with

$$\delta(l, l')\frac{\partial h_{lt}}{\partial w_{lt-i}} = \frac{\partial h_{lt}}{\partial w_{l't-i}} \implies \frac{\partial h_{lt}}{\partial w_{lt-i}} > \frac{\partial h_{lt}}{\partial w_{l't-i}}$$

for $i \geq 1$ and $l \neq l'$. □

While compliance may be increasing in the habit level, the attention allocated to hand washing is declining in habit. To show this graphically, I simulate the model and plot compliance levels, attention, and habit in Figure 1.1. The black line at the top of the area chart plots compliance $p(a^*, h_{lt})$ at the optimal attention level for each habit level, h_{lt} (the

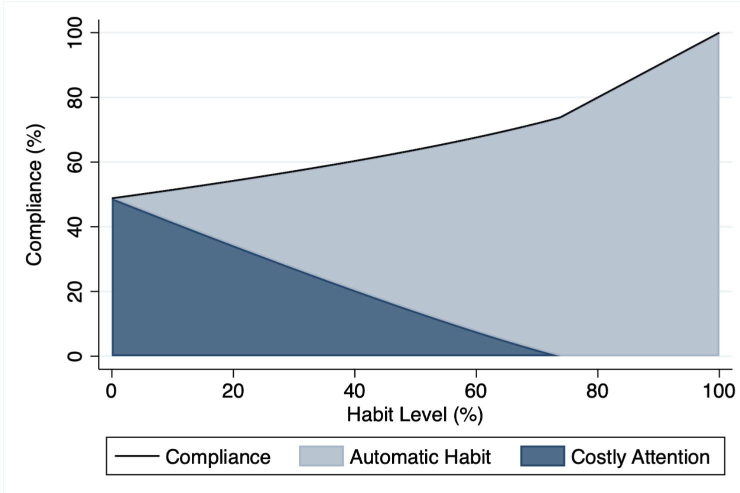

Figure 1.1: Simulated compliance and attention

Plot shows simulated compliance and attention by habit level.

x-axis). The dark region represents the share of compliance coming from effortful attention a^* while the light region is the remaining share of compliance coming from habit. This figure demonstrates the key feature of the model; habit and attention serve as substitutes. Compliance levels rise with habit, but the share of compliance driven by costly attention declines.

Distractions impact costly attention, but have no impact on habitual compliance. In the Figure, adding distractions make the dark region (attention) more costly. Because this makes up a smaller share of overall compliance at higher habit levels, distractions will be less harmful for these types. Prediction 4 formalizes this idea.

Prediction 4. *Distractions ($\uparrow d_t$) will have a larger impact on compliance when habit is low.*

Proof. Distractions have a negative effect on compliance. The prediction is true if $\frac{\partial p(a^*,h_{lt})}{\partial d_t}$ is increasing (i.e., becoming less negative) in the habit level.

$$\frac{\partial^2 p(a^*, h_{lt})}{\partial d \partial h_{lt}} = \frac{1}{4}\sqrt{\frac{1}{d(1-h)}} \geq 0.$$

□

The opposite intuition holds for disruptions to the habit level. Because high compliers rely more on habit than costly attention, their compliance will be more affected by shocks to habit than the low compliers.

Prediction 5. *Disruptions to habit ($\downarrow h_{lt}$) will have a larger impact on compliance when habit is high.*

Proof. If compliance is increasing in habit, $\frac{\partial p(a^*,h_{lt})}{\partial h_{lt}} > 0$. A positive second derivative implies that changes in habit are more impactful at high levels of habit.

$$\frac{\partial^2 p(a^*, h_{lt})}{\partial h_{lt}^2} = \frac{1}{4}\sqrt{\frac{d_t}{(1-h)^3}} \geq 0.$$

□

Figure 1.2 presents Predictions 4 and 5 graphically. From the baseline simulation in Figure 1.1, I plot the change in compliance for an increase in distraction (Figure 1.2a) or a decrease in habit level (Figure 1.2b) from different baseline compliance levels. Low compliers will have a larger response to an increase in distraction level and a smaller response to a change in habit. High compliers will have the opposite response.

(a) Distraction (b) Habit Disruption

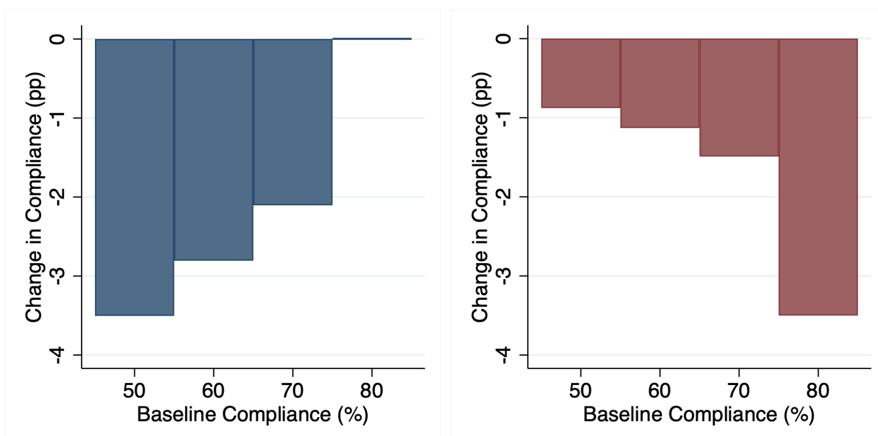

Figure 1.2: Simulated response to distraction and habit disruption

Simulated response to a distraction (a) or habit disruption (b) in the cue theory model of habit.

Alternative Logit Utility Model

An interesting benchmark is to compare these predictions to those of a logit discrete choice model. Let the utility from washing be given by

$$U_{lt} = b - c(d_t) + \sigma h_{lt} + \epsilon_t,$$

where b is the benefit from washing, $c(d_t)$ is the cost of washing (as above, a convex function of cognitive load), and h_{lt} is the location-specific habit which evolves as described above (scaled by σ). Assume the noise term, ϵ is *iid* extreme value and the utility from not washing is normalized to 0. Compliance, or the probability of washing $p(w_{lt})$ in location l at time t, is described by a logit with

$$p(w_{lt}) = \frac{e^{b-c(d_t)+\sigma h_{lt}}}{e^{b-c(d_t)+\sigma h_{lt}} + 1}.$$

Like the cue model, compliance will be declining in cognitive load, d_t, and increasing in the habit level, h_{lt}. However, Predictions 4 and 5 will not hold. Consider the derivatives of the probability of washing with respect to distractions and habit. Defining $x_{lt} = b - c(d_t) + \sigma h_{lt}$,

$$\frac{\partial p(w_{lt})}{\partial d_t} = \frac{e^{x_{lt}}}{(e^{x_{lt}}+1)^2}\left(-c'(d_t)\right) \text{ and}$$
$$\frac{\partial p(w_{lt})}{\partial h_{lt}} = \frac{e^{x_{lt}}}{(e^{x_{lt}}+1)^2}\sigma.$$

This model demonstrates the typical logit properties, with the largest elasticities for those with x_{lt} close to 0, or probability of washing closest to 50%. Above 50%, where most of the data lie, the logit implies higher compliance levels will have smaller elasticities, consistent with Prediction 4. However, the model predicts that individuals with larger response to cognitive load (i.e., those with large $\frac{e^{x_{lt}}}{(e^{x_{lt}}+1)^2}$) will also be more responsive to changes in habit. Figure 1.3 shows this result graphically with simulated responses from the logit model.

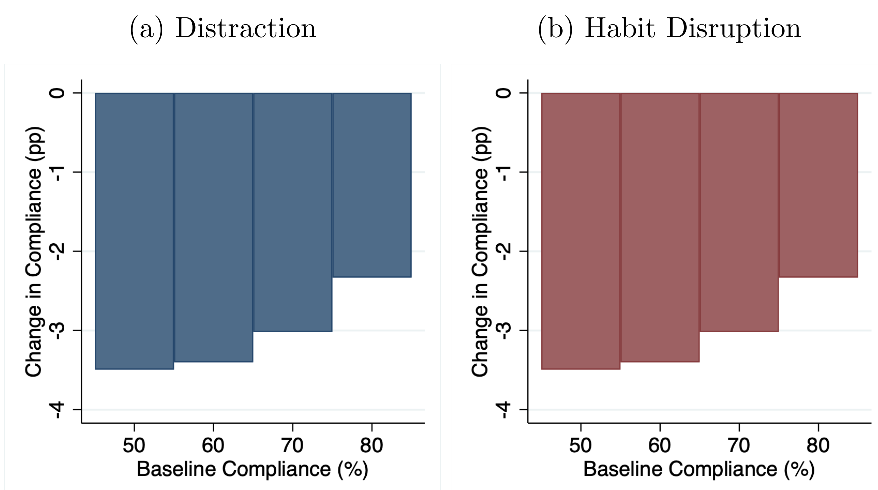

Figure 1.3: Simulated response to distraction and habit disruption

Simulated response to a distraction (a) or habit disruption (b) in the logit model.

Ties to the Empirics

In the empirics, I document patterns consistent with the assumptions and predictions of the cue theory model described above.

To test Prediction 1, I look at three proxies for cognitive load. Based on the assumption that fatigue acts as a cognitive load, I look at time at work (i.e., hours on the shift) and performance on the night shift (vs. the day shift).

Second, I look at the difference between hand hygiene behavior on weekdays and on the weekend. A long literature beginning with Macfarlane (1978) find that patients admitted on weekends have worse outcomes that those admitted on weekdays. Explanations in the literature include patient selection (e.g., Dobkin (2003)) and lower staffing ratios on weekends (e.g., Bell and Redelmeier (2001)). Staffing ratios, and weekends in particular, have been shown to correlate with higher reported stress levels, another form of cognitive load (Aiken et al. (2002), Purcell et al. (2011)).

Another difference between weekends and weekdays may be the level of distractions for workers. For example, on weekends there may be more activity on social media or texts and phone calls from friends and family. Workers on weekend shifts may also be associated with higher levels of fatigue if workers are out late on a Friday or Saturday night. This is conceptually similar to findings in finance that investors are distracted by the weekend and thus are slower to act on Friday earnings announcements (DellaVigna and Pollet (2009)).

In addition to the effects on cognitive load, shifts worked on nights and weekends likely have less managerial oversight. Indeed, those with identifiable leadership roles appear in the data less frequently on night and weekend shifts.[4] Reduced oversight may result in a lower benefit to washing, b, which demonstrates similar properties as an increase to cognitive load (see Appendix A.1 for derivations).

To test Prediction 2, I consider two shocks to habit. First, consistent with the assumption that habit decays over time, I test the impact of breaks from work on compliance. Time off work has two opposing effects. As documented by Dai et al. (2015), time off work can help lower fatigue and therefore increase compliance. However, if habits fade over time, then breaks from work can also lower compliance. I examine this tradeoff in Section 1.4.

As a second shock to habit, I consider variation in compliance across locations. Specifically, I examine changes in compliance on float days, or days when a worker covers a shift in a different department. If cues are location-specific, then habits in this new department should be lower than habits in the original department.

1.3 Data

Hospitals in the data use a "gel-in, gel-out" policy, which requires hospital workers to wash or sanitize their hands on entry to and exit from a patient room. Hand hygiene compliance data was obtained from SwipeSense, a company that developed a technology to monitor individual compliance with this policy. The technology uses three key components: RFID badges worn by workers, sensors on soap and sanitizer dispensers, and sensors in each patient room. Badges are assigned to a unique worker, which allows for tracking of individuals. To

[4] I define leadership as those with roles such that include terms such as leadership, director, manager, supervisor, chief, VP, discharge planner, and division head.

identify compliance with the gel-in, gel-out policy, room sensors record the time of worker entry to and exit from a patient room. Dispenser sensors identify when a soap or sanitizer dispenser is used and assign each use to the closest worker at the time of use. Each entry to or exit from a patient room is time stamped and recorded in the data as a hand hygiene opportunity. Compliance in each opportunity is defined as a binary outcome, equal to 1 if the worker used a soap or sanitizer dispenser within a 60 second window of entry to or exit from the patient room.

For each hand hygiene opportunity, the data includes the type of event (room entry or exit), time, location information (room number, department, unit, and facility), worker information (anonymized individual identifier, role, department), and compliance.

The dataset includes any worker at the facilities who wears an RFID badge including, for example, physicians, registered nurses, nursing assistants, technicians, administrators, and environmental service workers. The nursing staff make up the majority of the data; they work 63% of shifts and have 73% of patient interactions.

I use the pattern of patient interactions to infer shift start and end times. Following the procedure used in Dai et al. (2015), I define a new shift as occurring if 7 or more hours elapse between the end of one patient interaction and the start of the next.

The distribution of resulting shift lengths, shown in Figure 1.4, has three peaks. There is a large mass of shifts that last around 12 hours and a smaller peak close to 8 hours, which are common shift times in the industry. The third peak is close to 0, with about 10% of shifts lasting under 15 minutes. These short times are likely snapshots from a longer shift with few patient interactions (e.g., a physician who spends the majority of their day on research but has one patient visit). While a large fraction of shifts, this represents only a small share of opportunities. I limit to the set of shifts with sufficient patient interaction, namely those that last 5 to 12 hours. These shifts make up 68% of shifts, but 90% of all interactions.

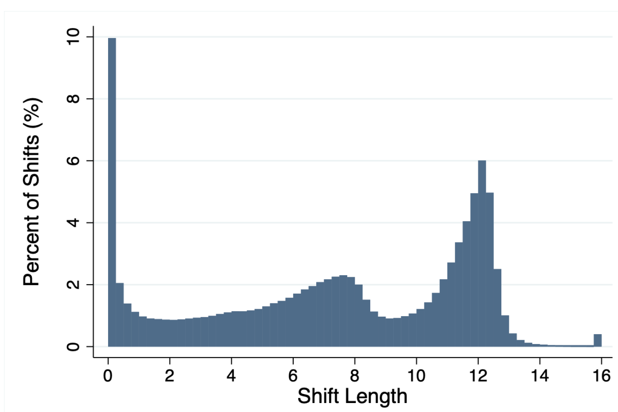

Figure 1.4: Distribution of shift length in the full sample

Plot shows a histogram of shift length in hours in the full sample. Length is capped at 16 hours for the figure.

The goal of the analysis is to study hand hygiene behavior among those for whom this is a common and frequent aspect of their job. For example, I want to include a nurse who has 60 hand hygiene opportunities per day but not an HR representative with only the occasional visit to a patient room. I limit the data to workers who average at least one opportunity per hour at work, at least one shift per week, and are in the data for at least 6 months. These make up only 40% of workers, but 80% of shifts and opportunities. Appendix A.2 summarizes the full data and the main sample.

The data covers hand hygiene behavior of 13,606 workers in 123M opportunities across 2.1M shifts. These shifts take place in 276 departments, 32 facilities, and 21 networks from August 2016 to February 2020. For the main analyses, I exclude all data from March 2020 on to exclude the COVID-19 pandemic. I discuss the interesting trends since the start of the pandemic in Appendix A.3. While intriguing, the potential changes in hospital procedures during this time make it difficult to draw conclusions without more information.

In Table 1.1, I describe the data. Panel A summarizes the shifts included in the sample. The median shift lasts 10.6 hours, during which a worker has about 58 hand hygiene opportunities (5.5 per hour).

Figure 1.5 shows the distribution of shift start times in the final sample. The vast majority of shifts start around 7am and 7pm with smaller masses at 11am, 3pm, and 11pm. I define day shifts as those that start between 2am and 2pm (64% of shifts). Consistent with typical hospital schedules, 56% of shifts appear to be on a 12-hour 7am/7pm schedule and 24% on an 8-hour 7am/3pm/11pm schedule.[5]

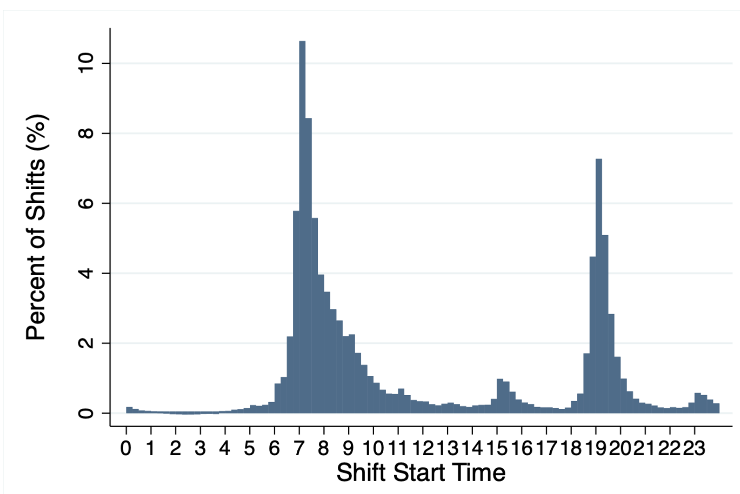

Figure 1.5: Distribution of shift start times

I define the shift date at the beginning of the shift, except for the small fraction of shifts that start between 12 and 2am. For these, I use the day before in order to assign consistent dates for night shifts that start just before and after midnight. Across days of the week,

[5]I bin start times according to the following: 7am (6-8am), 11am (9am-12pm), 3pm (2-3pm), 7pm (6-8pm), 11pm (10pm-12am).

weekdays have a slightly higher than proportional share of shifts (77% of shifts worked on weekdays vs. weekends).

Shifts are typically concentrated in one department where workers visit patients in about 8 different rooms. I associate shifts with a department if at least 75% of opportunities occur in a single department.[6] Workers are typically in the same department over time; the median person works 90% of their shifts in the same department. As is common in hospitals (and especially for nurses), 1.8% of shifts appear to be float days, when a worker temporarily works a shift in a different department. I define a float day as a 1 to 2 shift switch from one department to another.[7]

Consistent with these being active workers, the median shift is worked after only one day off. While typically active in the data, I also observe long breaks with 1.4% of shifts occurring after at least 2 weeks off.

Panel B of Table 1.1 describes the workers in the data. The average worker is compliant in 58% of their hand hygiene opportunities, with a standard deviation across workers of 22%. Figure 1.6 graphically depicts the large amount of heterogeneity in the distribution of compliance across individuals. The distribution drops steeply above 80% compliance; 17% workers have compliance rates above 80% and only 2% have compliance above 90%. This lack of mass at the very high end likely reflects an upper bound on attainable compliance due to either the circumstances of the interaction (e.g., carrying supplies into a room) or limitations of the technology (e.g., mismeasurement in the time of room entry/exit or misattribution of a dispenser use). Many workers have compliance well below this threshold, with 35% of workers at an average compliance under 50%.

Consistent with the selection on active workers, the average individual works 21.7 hours per week, for a total of 156 shifts over about a year and half in the data. The type of workers are also disproportionately weighted toward jobs with high levels of patient interaction. The majority of individuals are nursing staff; 46% are registered nurses, 13% are nursing assistants, and 4% are other nurses. The remaining 37% of workers cover a variety of roles including technicians (9%), housekeeping and food services (5%), and doctors (4%).

1.4 Empirical Evidence for Cue Theory of Habits

In this section, I provide descriptive evidence consistent with the assumptions and predictions of the cue theory model of habits.

[6] 85% of shifts are associated with a department.

[7] Specifically, I define a switch from from department 1 to 2 if a worker has been working in department 1 for at least three shifts, then switches to department 2 for either one or two shifts (the float days), followed by a shift that is not worked in department 2 (most commonly returning back to department 1).

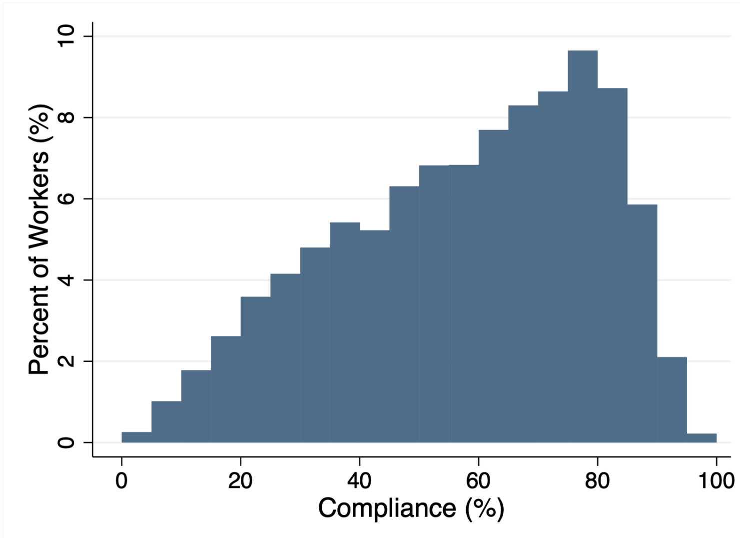

Figure 1.6: Heterogeneity in compliance

Distribution of compliance rates across individuals.

Fact 1: Compliance falls with cognitive load

According to Prediction 1, cognitive load should increase the cost of attention, lowering the amount of attention workers pay to hand hygiene and thus compliance rates. As described in Section 1.2, I test this prediction using three proxies for cognitive load: hours at work, night shifts, and weekend shifts.

In Column 1 of Table 1.2, I report the results from an OLS regression of compliance in an opportunity (a binary outcome rescaled to 100) on hours at work (divided by 12 so the coefficient reports the difference between the start and end of a 12-hour shift) and indicators for whether the shift was on a night or weekend. I control for individual fixed effects to account for variation in the perceived net benefit by different individuals. I include department fixed effects to control for variation in benefits across patient types (e.g., neonatal intensive care vs. emergency room). I include network-month-year fixed effects to account for time varying trends in actual or perceived changes in benefits of hand hygiene (e.g., flu season or the CEO sending an email reminding everyone to wash). In addition to the fixed effects, I also control for other potential proxies for benefits: the worker's total number of opportunities in the hour (normalized within person), an indicator for room entry (vs. exit), and the number of months the worker has been in the data to account for learning over time. I cluster standard errors by individual. I estimate all regressions using the Stata command reghdfe (Correia (2016)).

Consistent with Prediction 1, all estimated coefficients on the proxies for cognitive load are negative. Figure 1.7 shows the results graphically. Consistent with prior work by Dai et al. (2015), I find a large drop in compliance over the course of a shift, with compliance falling almost 6pp from the beginning to the end of a 12-hour shift. Compliance on the night shift is about 1.5pp lower than on the day shift and 0.5pp lower on weekends compared to weekdays.

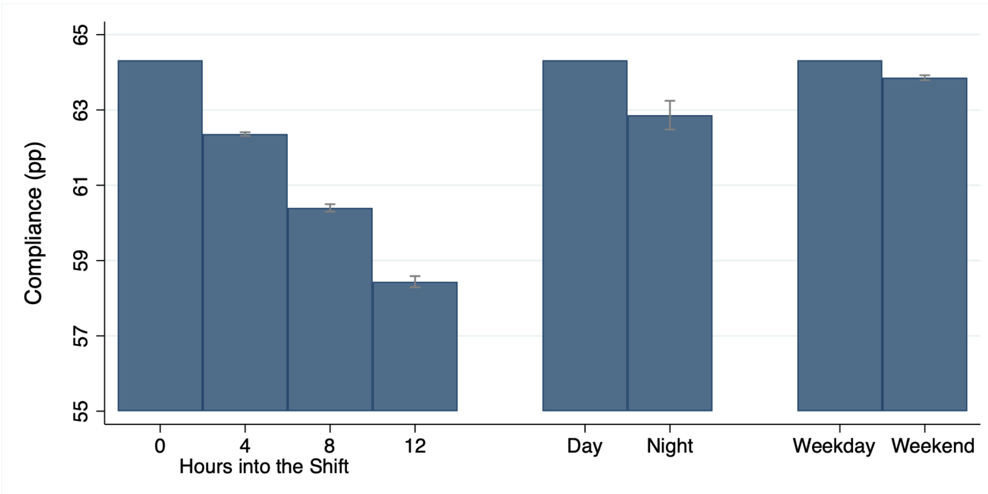

Figure 1.7: Estimated predictors of compliance

Graph plots estimated coefficients from Column 1 of Table 1.2 relative to the baseline level of compliance.

Fact 2: Compliance falls with disruptions to habit

According to Prediction 2, compliance is increasing in the level of habit. Therefore, any negative shocks to habit should lower compliance. As discussed in Section 1.2, I consider two disruptions to habit: float days and time off work.

I begin by examining disruptions with an event study design. Focusing on the variation in compliance across (rather than within) shift, I regress average compliance over the shift on event time indicators (described below) controlling for day shift (vs. night), 12-hour shifts (lasting 9-13 hours vs. 5-9 hours), number of months a worker has been in the data, as well as individual, main department, and network-day fixed effects. I weight shifts by the total number of opportunities and cluster standard errors by individual.

For float days, I estimate event time indicators for shifts relative to a switch that lasts either 1, 2, 3, and 4 or more shifts. I include indicators for the 4 shifts before the float day (with 2 shifts before as the reference category), the first float day, and the 4 shifts after. I also include indicators for all shifts more than 4 shifts prior to a float day and 4 shifts after the first float day.[8]

[8]Following Sandler and Sandler (2014), to account for the fact that there may be multiple events, rather than binary indicators, the binned categories take on a value equal to the number of events that are more than +/-4 shifts away. For example, if a shift is 10 shifts after one float day and 30 shifts after another, the event time indicator for more than 4 shifts after a float day would be equal to 2. I also include an indicator for the first 3 shifts, which by construction cannot be float days.

Figure 1.8 plots the event time coefficients for department switches that last 1 or 2 days, the most commonly observed in the data.[9] Despite the potential endogenous assignment of float days, the estimated event time coefficients reveal no major pre-trends in the shifts prior to the float day. On the float days, I observe a large drop in compliance on the order of 2.5 to 3pp when working in a different department from the previous 3 shifts. After the float day (+1 for the 1 day switch and +2 for the 2 day switch), compliance essentially returns to pre-float levels as workers go back to their main department.[10]

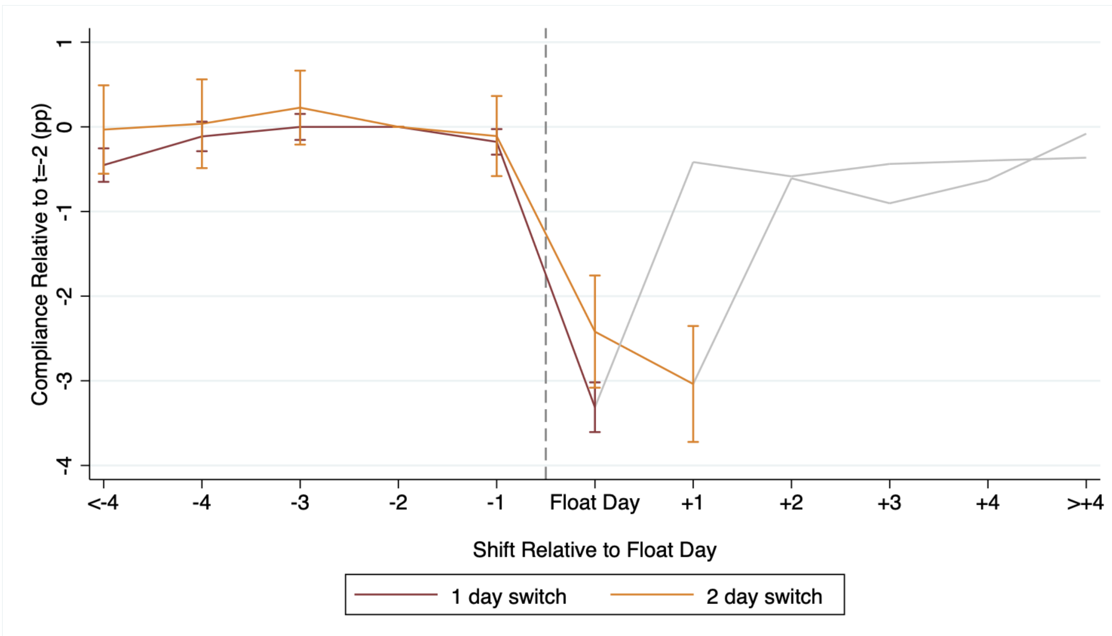

Figure 1.8: Event study graphs: float days (department switches)

See text for description of plotted coefficients.

I use the same event study methodology to look at variation in compliance around breaks from work of different lengths. Specifically, I include event time indicators for shifts around a 1 day break between shifts, a 2 day break, a 3 to 4 day break, and so on up to a 4+ week break between shifts. Figure 1.9 plots the event time coefficients. As discussed in Section 1.2, we may expect time off work to have two conflicting effects on compliance. As documented by Dai et al. (2015), time off work can reduce fatigue and increase compliance. Indeed, I confirm these effects for shorter breaks from work. Relative to a consecutive shift, compliance is 0.7pp higher after 1 or 2 days off and almost 1pp higher after 3 to 4 days off. For longer breaks the trend starts to decline, with no significant change in compliance after a break of 2 weeks or more. This pattern is consistent with habits decaying over time. Over

[9]Of the 39,151 switches that occur, 80% last only one shift, 10% last two shifts, 3% last 3 shifts, and 7% last 4 or more shifts. Appendix Figure A3 also plots the event study coefficients for switches that last more than 2 shifts.

[10]In 83% of 1 and 2 day switches, workers return to the original department.

the first few days, additional rest from more time off increases compliance. After only a few days, these restorative benefits appear to max out with additional time off resulting in lower compliance on return.

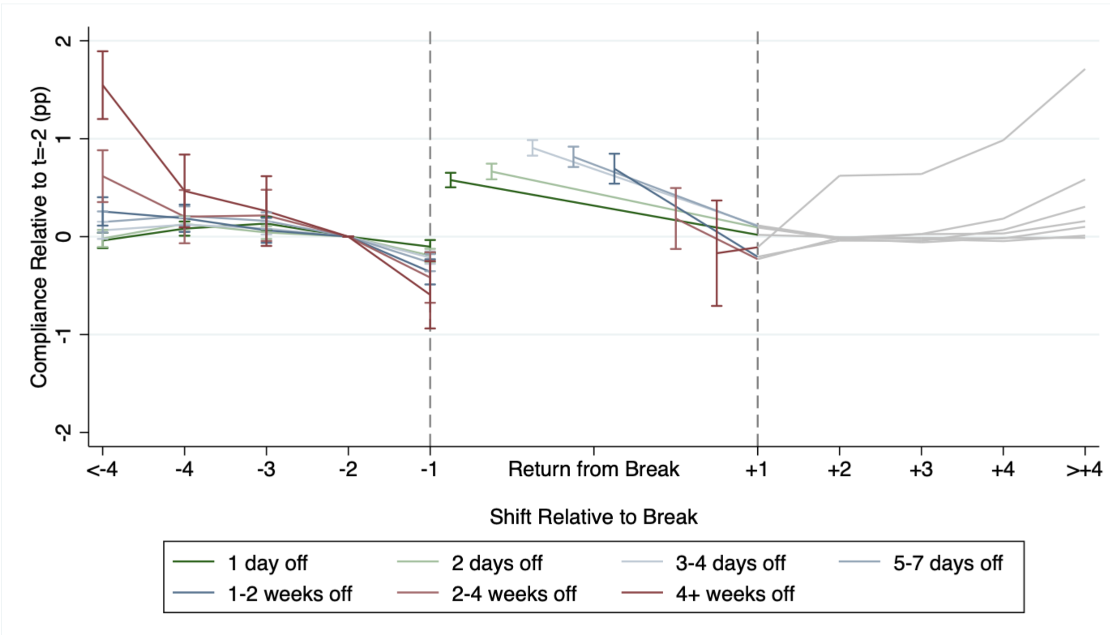

Figure 1.9: Event study graphs: days off between shifts

See text for description of plotted coefficients.

I synthesize these effects in Table 1.2 by adding indicators for simplified versions of the events to the baseline regression in Column 1. To capture the effects of time off work, I bin shifts into those that are worked consecutively after another shift (no day off), after a short break (1-13 days off), and after a long break (14+ days off). I include indicators for consecutive days and long breaks, with short breaks as the reference category. Relative to a short break, consecutive days have 0.77pp lower compliance. Consistent with long breaks acting as a disruption to habit, compliance is 2pp lower after a break of at least 2 weeks.

To capture float days, I include indicators for the 1 or 2 days that are worked in a department different from the last three. As in the event study, these float days have almost 3pp lower compliance than other shifts.

Fact 3: Cognitive load has a larger impact on low compliers

According the cue theory of habits, cognitive load like fatigue and distraction should impact attentive hand washing but not habitual hand washing. Therefore, cognitive load should have a larger effect when habit is low and hand washing is driven by effortful attention (Prediction 4).

To test this prediction empirically, I split workers into low, medium, and high compliers by compliance in their first three months of the data.[11] I then estimate the baseline regression in Column 2 of Table 1.2 separately for each of the three sets of workers. I report the results in Table 1.3 and show the results graphically in Figure 1.10. There are 4,495 workers in each bin with average compliance of 45%, 63%, and 79%.

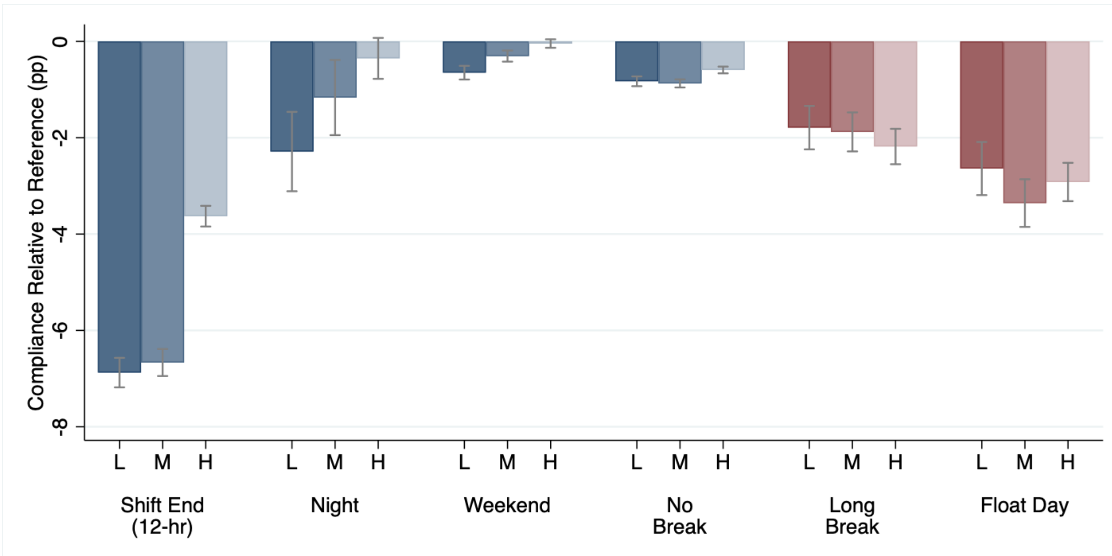

Figure 1.10: Effect of fatigue, distractions, and habit disruptions on compliance
Graph plots estimated coefficients from Table 1.3 for low, medium, and high compliers.

Consistent with Prediction 4, the negative impacts of time at work, working the night shift, and working weekends are concentrated among low compliance workers. The magnitudes of the differences across types are large. High compliers decline roughly half as much as low compliers over the course of a shift (6.88pp vs. 3.63pp). High compliers show no significant declines on nights and weekends, compared to a 2.29pp drop on nights and 0.65pp drop on weekends for low compliers. The negative effects of working a consecutive shift are also 40% larger for low compliers compared to high compliers (0.83pp vs. 0.59pp).

While these results are consistent with the predictions from the cue theory of habit, they are also qualitatively consistent with the predictions of the logit model described in Section 1.2. The results in the following section, on the impact of disruptions to habit, differentiate the models.

[11]The main results are robust to alternative ways of categorizing compliance levels including, for example, adjusting for start month, sorting within network and role-type, and sorting on endline rather than initial compliance. See Appendix A.5 for additional information on robustness and Appendix Table A1 for summary statistics for each of the subsamples.

Fact 4: Disruptions impact both low and high compliers

According to Prediction 5, disruptions to habit should affect exactly the opposite people from shocks to cognitive load. Those who rely more on their habit to wash will be more impacted by a shock to habit than those who need to pay costly attention to perform hand hygiene. The logit model, on the other hand, would predict that individuals who responded more to cognitive load will also respond more to disruptions to habit.

As described above, I estimate the impact of float days and long breaks from work separately for low, medium, and high compliers. I report the results in Table 1.3, shown graphically in Figure 1.10.

Contrary to the effects of cognitive load, shocks to habit appear to affect all workers, not just low compliers. Thus, high compliers are less responsive to cognitive load, but are just as responsive as their low compliance peers to shocks to habit.

Note that the theory predicts that shocks to habit should have a *larger* impact on high compliers. However, if shocks to habit are paired with increased cognitive load, we should expect to see responses from both low and high compliers. As these are unusual events, they are likely paired with distractions such as increased socializing with coworkers after returning from vacation or complications of trying to navigate in a new department (e.g., trying to find supplies or figuring out who to ask for a consult).

In addition to the cognitive load directly caused by these events, there may also be cognitive load caused by effortful attention that must be expended for other habitual behaviors. Both of these disruptions (switching departments and long breaks from work) should affect not only hand hygiene habits, but any habitual behavior. With drops in habit across the board, workers would have to pay more attention to any behaviors that have become habitual in their usual routine. The increased attention to these other domains could act as an additional cognitive load, making attention to hand hygiene more costly and thus eliciting a response from both low and high compliers.

Fact 5: Location-specific habit formation

In this section, I explore the formation of habits and provide evidence of location-specific habits.

In Figure 1.11, I plot estimated coefficients and the 95% confidence interval from a regression of compliance on the worker's month in the data and individual fixed effects. I estimate effects over the first year in the data and limit to the 7,904 workers with tenure of at least one year. Compliance increases substantially over the first year in the data, with compliance after 1 year about 10pp higher than starting compliance. This is consistent with Staats et al. (2017) who find that compliance increases by about 10pp over the first 20 months after introduction of a similar monitoring technology.[12]

[12]Interestingly, the impact of monitoring appears to wear off over time with compliance levels declining back to initial levels in months 20 to 40 after introduction of the technology.

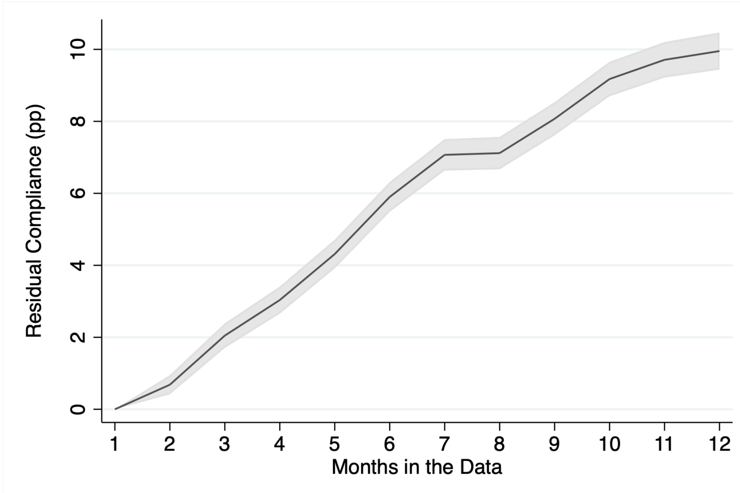

Figure 1.11: Compliance over the first year in the data

Plot shows estimated coefficients and 95% confidence interval from a regression of compliance on months in the data and individual fixed effects. Sample limited to workers in the data for at least 12 months. Standard errors clustered by individual.

This increase in compliance over a worker's tenure could be explained by many versions of learning and habit formation. Time in the data may capture learning about the benefits of hand hygiene through, for example, education on the importance of hand hygiene, campaigns to increase compliance, or feedback about individual performance.

Cue-based habit formation, on the other hand, should be specific to experience with washing in response to a particular cue. The results on switching departments suggest that some cues are department-specific. If these are visual cues tied to physical location, habit formation may be even more localized to a particular room. As a simplistic example, consider two rooms that have a sanitizer dispenser placed on either the right or left side of the door. A caregiver who frequently visits the right-side room develops a habit to wash when they see the dispenser to the right of the door. If they then visit the left-side room, the different visual set-up does not evoke habitual hand washing. In reality, there are likely a number of differences between rooms that may result in different cues to wash and therefore room-specific habits.

In Figure 1.12, I show how compliance evolves within room-specific experience. Splitting the first year data into four 3-month intervals, I plot estimated compliance levels based on the number of prior opportunities that person has had in the room. The plots show compliance relative to the first 10 opportunities in a room during the first three months of tenure, controlling for both individual and room fixed effects. The left graph includes all rooms visited while the right graph is limited to each person's most frequently visited department.

These graphs demonstrate two key points. First, the lines are shifting up with time. Thus, compliance is increasing with time in the data, for example as learning is going on (e.g., hand

(a) All Rooms (b) Main Department

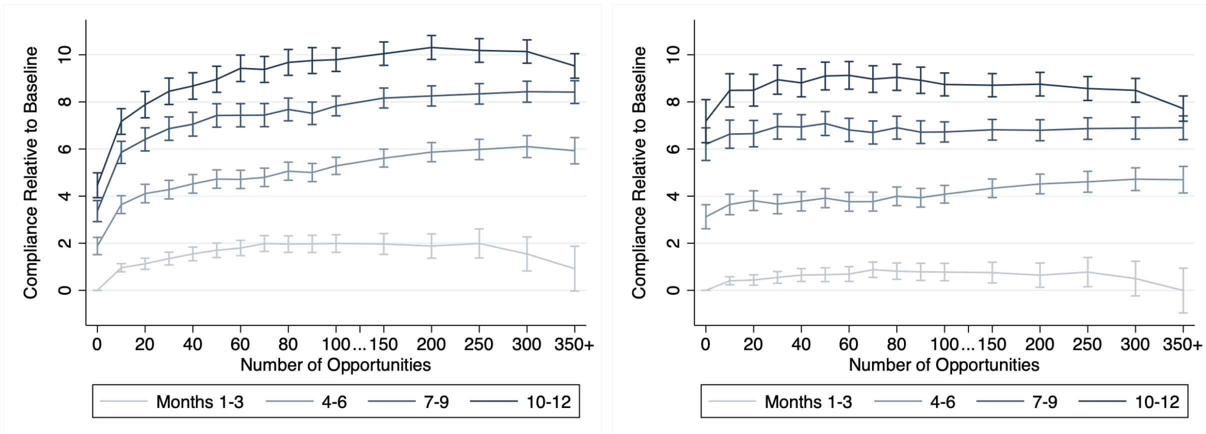

Figure 1.12: Compliance by number of prior opportunities in a room

Figure plots estimated coefficients from a regression of compliance on binned values of the number of prior opportunities in a room as well as individual and room fixed-effects. The scale of the x-axis is compressed above 100 opportunities.

hygiene training sessions) or as general habits are building with experience. However, within a time period, the upward slope indicates that compliance is higher in rooms that have been visited more frequently. Across all rooms (left graph), this increase is apparent in all time periods. In the first three months, compliance increases by almost 2pp over the first 50 opportunities in a room. In months 10-12, compliance increases by about 5pp over the first 50 opportunities.

The left graph includes variation both within and across departments. In the right graph, I limit to rooms in the worker's most frequently visited department. Even within this main department, compliance increases about 1pp after the first 50 opportunities in a room. The magnitude of this effect is similar over time, though the estimates become noisy as there are fewer rooms in the main department that have not been visited. For example, 50 prior opportunities is the 45th percentile in months 1 to 3 but only the 5th percentile in months 10 to 12.

1.5 Implications for Organizations

The facts above are consistent with two types of compliance. Automatic compliance, formed through repeated hand washing, is habitual and not impacted by drains on attention like fatigue or distraction. Attentive compliance, on the other hand, requires costly effort by the worker and is responsive to competing demands on attention. This variation has several implications for hospitals aiming to increase compliance.

Targeting Interventions

First, the large amount of heterogeneity in the data point to benefits from targeting interventions to individuals and at times when they will be most effective. For example, consider a text message sent to providers to remind them to wash their hands.

Receiving a text message reminder likely carries some annoyance cost. Receiving multiple messages may overload the recipient, reducing the impact of the message. In one of the few studies that examines repeated messaging, Allcott and Rogers (2014) find that initial social-comparison messages are more effective than subsequent ones. Thus, we may want to target even a low-cost intervention like text message reminders.

When will reminders be most effective? Reminders lower the cost of attention, increasing compliance in situations that would otherwise require costly effort. Reminders may be even more effective when costs are high, for example at the end of the shift compared to the start. According to the cue theory, reminders will not affect habitual hand washing, which is unaffected by the cost of attention. Thus, the findings in this paper suggest that reminders will be most effective for those with low habit levels (low compliers) at times when attention is most costly (late in the shift, on nights and weekends, and on consecutive shifts).

These predictions could be tested with a simple experimental design randomizing the timing of text message reminders sent to healthcare workers. Using observational data, we can predict for which workers and in which situations a reminder would be most effective and compare the effectiveness of a one-size-fits-all intervention to a targeted messaging approach.

Attention-Grabbing Interventions

The second lesson from the cue theory of habits is that the nature of the intervention matters. Many hand hygiene interventions are designed to attract attention to hand hygiene compliance (e.g., educational or social norms interventions). If sustained over the long run, these interventions may help to build habits. However, the model implies they will have short run welfare costs as low habit workers have to exert more attention to increase compliance.

Thinking beyond hand hygiene, if total attention capacity is limited, attention-grabbing hand hygiene interventions may have perverse effects in other domains. For example, Son et al. (2011) describe a highly successful 12-week intensive hand hygiene intervention that increased compliance from 60-70% to over 90%. The program involves convening multi-disciplinary hand hygiene teams in each department who identify barriers to compliance, evaluate workflows, develop and provide training to peers, develop and implement a system of peer-observation, data collection, and feedback. While successful in increasing compliance, an intervention of this scale is likely a huge drain on attention. Did this organization's emphasis on hand hygiene have any negative impacts on other behaviors?

Unfortunately, this is a difficult question to answer as many aspects of healthcare worker behavior are difficult to observe and evaluate. However, the question highlights an exciting future avenue for research using this hand hygiene compliance data. Rather than focusing on hand hygiene as the key outcome of interest, compliance can serve as a measure of the drains on attention. The objective and high-frequency nature of electronic monitoring data allows for a reliable and consistent measure of cognitive load that we typically don't have in field settings. For example, changes in compliance could provide a benchmark for testing the success of different technology and management interventions aimed to improve efficiency or to quantify the costs of stresses like patient load or noise (Aiken et al. (2002), Morrison et al. (2003), Dean (2020)).

Cost-Reducing Interventions

Rather than increasing the attention paid to hand hygiene, the model implies that reducing the costs of attention would both increase compliance and individual welfare. As with other types of interventions, short-term interventions will likely have minimal impact. However, the results provide several potential avenues for hospitals to make organizational changes for sustained impact.

The results on location-specificity of habits suggests that cues vary across departments and rooms. Standardizing cues across locations can improve the portability of habits. Given the large number of rooms that healthcare workers visit, this has the potential to have a substantial impact on compliance. The most obvious way to standardize cues is to make rooms as physically similar as possible. Even a small change such as having the sanitizer dispenser to the left or right of the door may change the nature of the cue-response. A literature in standardization and simplification in hospital design supports this idea (e.g., Price and Lu (2013), Ulrich et al. (2004)). It is also possible that new cues can be introduced to the environment. King et al. (2016) find that hand hygiene among healthcare workers and hospital visitors was higher when primed with either a clean scent or a picture of eyes above sanitizer dispensers.

Alternatively, rather than changing cues, habits could be made more effective by concentrating work in similar environments. While float days may allow for more scheduling flexibility, there is a clear cost to working in a different department. Even within a department, assigning workers to the same rooms across shifts may facilitate faster habit accumulation and limit habit decay over time.

Taking cues as fixed, another way to lower the costs of hand hygiene would be to reallocate shifts based on individual heterogeneity. Consider allocating shifts with high attention costs (e.g., long shifts, nights, weekends) to high habit types. This would increase overall compliance as the high habit types will not respond as negatively to the increased cognitive load. Because of the benefit from automaticity, the high types are able to maintain high compliance without paying high attention costs. Over the long-run, low compliers would have the chance to build better habits as they practice high compliance on low-cost shifts. This simplistic swap of course ignores other costs of working undesirable shifts, such as coordinating child care on nights and weekends. A more equitable approach may be to swap undesirable shifts, for example by assigning night shifts to high compliers and float days to low compliers. Relative to the alternative shift assignment, this would increase compliance by 2pp on these shifts. On a larger scale, using these insights to optimize shift allocation could have sustained effects on compliance at minimal cost to the organization.

1.6 Conclusion

In this paper, I document evidence consistent with a cue theory of habit formation in hand hygiene. I show that compliance grows with location-specific hand washing experience. Consistent with the theoretical predictions, I find high compliers are buffered from negative impacts of cognitive load but are impacted by disruptions to habits.

These findings have direct implications for hospitals, for whom even small changes in hand hygiene can have large effects. Estimates on the benefits of compliance suggest that even a 1pp increase in compliance would result in 23 fewer infections, 1.3 fewer deaths, and $0.5M in savings for the average hospital each year. Across the country, these estimates imply a reduction of 140k infections, 8k fewer deaths, and $2.9B in savings.[13]

The findings highlight the benefits of targeted interventions, the differing welfare implications across types of interventions, the value of sustained intervention, and suggest a focus on interventions that can help to build habits. Testing these predictions would be a valuable area for future research.

[13]Following calculations by Dai et al. (2015), I assume a 1pp increase in compliance results in 3.9 fewer infections per 1000 patients (Pittet et al. (2000)), that 5.82% of infections are fatal (Klevens et al. (2007)), and that the average infection costs $20,549 Scott (2009). According the American Health Association, in 2018 there were 36,353,946 patients admitted across 6,146 hospitals.

Table 1.1: Electronic Monitoring Data

	Mean	Median	Standard Deviation
Panel A: Across shifts (N=2,123,320)			
Opportunities per hour	5.7	5.5	2.9
Shift length (hours)	9.7	10.6	2.4
Day shift (vs. night)	64%		
Weekday (vs. weekend)	77%		
Number of departments per shift	1.7	1	1.2
Number of rooms per shift	8.2	7	4.6
Float days (switch department)	1.8%		
Days off between shifts	2.3	1	7.9
Consecutive shift (no days off)	46%		
Long break (shift following 14+ days off)	1.4%		
Panel B: Across individuals (N=13,606)			
Compliance	58%	61%	22%
Hours worked per week	21.7	21.5	8.0
Number of shifts	156.1	129	98.5
Tenure (months)	17.5	14	9.9
Roles:			
Registered Nurse	46%		
Nursing Assistant	13%		
Other Nurse	4%		
Technician	8%		
Housekeeping + Food Services	5%		
Doctors	4%		
Other Jobs	20%		

Notes: Table displays summary statistics of the electronic monitoring data. Panel A summarizes the data across shifts and Panel B across individuals.

Table 1.2: Impact of Fatigue, Distractions, and Disruptions on Compliance

Dependent Variable: Compliance	(1)	(2)
Shift end (12 hr)	-5.88***	-5.89***
	(0.08)	(0.08)
Night shift	-1.46***	-1.38***
	(0.19)	(0.20)
Weekend	-0.46***	-0.35***
	(0.03)	(0.03)
Consecutive shift (no day off)		-0.77***
		(0.02)
Long break (14+ days off)		-2.00***
		(0.12)
Float day (department switch)		-2.95***
		(0.13)
Individual FE	X	X
Department FE	X	X
Network-Month-Year FE	X	X
Benefit Proxies	X	X
Opportunities	122,804,319	122,083,817
Workers	13,606	13,606
Mean Compliance	60.5	60.6
R^2	0.195	0.195

Notes: Table shows coefficients from OLS regressions of compliance (rescaled to 100) in each hand hygiene opportunity on attributes of the interaction. Shift end is a linear trend of time on the shift (scaled to show the difference between the beginning and end of a 12-hour shift). Night shift indicates that the shift started between 2pm and 2am. Weekend indicates the shift started on a weekend. Consecutive shifts are worked with no days off in between. Long break indicates shift that follows a 14+ day break. Reference group is 1 to 14 days off. A float day is a 1 to 2 day change from one department to another (see text for detailed definition). Benefit proxies include opportunities per hour (normalized within individual), room entry (vs. exit), and a linear trend over months in the data. Standard errors are clustered by individual.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.3: Heterogeneity in Response to Fatigue, Distractions, and Habit Disruptions

Dependent Variable: Compliance	(1) Low Average Compliance	(2) Medium Average Compliance	(3) High Average Compliance
Shift end (12 hr)	-6.88***	-6.67***	-3.63***
	(0.16)	(0.14)	(0.11)
Night shift	-2.29***	-1.17***	-0.35
	(0.42)	(0.40)	(0.22)
Weekend	-0.65***	-0.31***	-0.04
	(0.07)	(0.06)	(0.05)
Consecutive shift (no day off)	-0.83***	-0.87***	-0.59***
	(0.05)	(0.04)	(0.04)
Long break (14+ days off)	-1.79***	-1.88***	-2.18***
	(0.23)	(0.21)	(0.19)
Float day (department switch)	-2.64***	-3.36***	-2.92***
	(0.28)	(0.25)	(0.20)
Individual FE	X	X	X
Department FE	X	X	X
Network-Month-Year FE	X	X	X
Benefit Proxies	X	X	X
Opportunities	32,426,886	32,559,200	32,117,233
Workers	4,495	4,495	4,495
Mean Compliance	44.5	63.1	79.3
R^2	0.171	0.108	0.071

Notes: Table shows coefficients from OLS regressions of compliance (rescaled to 100) in each hand hygiene opportunity on attributes of the interaction. Data split by average compliance in the first three months in the data. Shift end is a linear trend of time on the shift (scaled to show the difference between the beginning and end of a 12-hour shift). Night shift indicates that the shift started between 2pm and 2am. Weekend indicates the shift started on a weekend. Consecutive shifts are worked with no days off in between. Long break indicates shift that follows a 14+ day break. Reference group is 1 to 14 days off. A float day is a 1 to 2 day change from one department to another (see text for detailed definition). Benefit proxies include opportunities per hour (normalized within individual), room entry (vs. exit), and a linear trend over months in the data. Standard errors are clustered by individual.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Transitory Remarks

In Chapter 1, I explore the role of habit formation in hand hygiene behavior of hospital workers. Habit formation leads to inertia in actions, with past hand hygiene behavior predictive of future hand washing.

In Chapter 2, I explore a distinct mechanism through which the past influences future actions. Joint with Ulrike Malmendier, Chapter 2 explores the role of past macroeconomic experiences on homeownership decisions. Specifically, we show that people who live through periods of high inflation are more likely to own their home, in an effort to hedge against high inflation. In this Chapter, we demonstrate how past experiences (rather than actions) can have lasting effects on individual choice. These individual differences can aggregate up to meaningful cross-country differences.

Chapter 2

Rent or Buy? Inflation Experiences and Homeownership within and across Countries[1]

2.1 Introduction

Participation in the housing market varies widely, both across and within countries. As in the case of other asset markets, notably the stock market, researchers have struggled to fully explain participation decisions. Within Europe, for example, less than half of all households own their home in Germany and Austria, compared to 85% or more in Lithuania, Slovakia, and Croatia. Only 57% of households own their home in France, but 82% do in neighboring Spain. There are also sizable cross-sectional differences within countries. In Italy, for example, 48% of 30-year-olds own their home, but 79% of 60-year-olds. In the Netherlands, instead, the homeownership rate is nearly identical for 30- and 60-year-olds.

What explains this "housing-market participation puzzle"? Why do households with similar demographics and in similar financial situations make systematically different tenure decisions? Clearly, institutional differences play an important role, at least for the cross-country differences, as do variations in house prices, housing supply, and population demographics.[2] In this paper, we identify a novel and economically meaningful determinant of homeownership decisions both within and across countries: We show that *past* macroeco-

[1]This paper is joint with Ulrike Malmendier. We thank Victor Couture, Anthony De Fusco, Amir Kermani, Herve Le Bihan, Charles Nathanson, Nick Sander, Nancy Wallace and workshop participants at Berkeley, Chicago, Duke, Stanford as well as the AEA, Finish Economic Association, ESA conferences, the Household and Behavioral Finance Symposium at Cornell conference, the CEPR Network Event on Household Finance, the BdF/ECB Conference on Household Finance and Consumption, and the UCSD Spring School in Behavioral Economics for helpful comments. We thank Luisa Cefala, Pablo Munoz Henriquez, and Murilo Ramos for help with translations. This paper uses data from the Eurosystem Household Finance and Consumption Survey, the American Community Survey (from IPUMS), and the Survey of Health, Ageing and Retirement in Europe. We also thank John Landon-Lane for providing historical house-price data.

[2]Prior evidence includes Andersen (2011), Andrews and Caldera Sánchez (2011), Andrews et al. (2011),

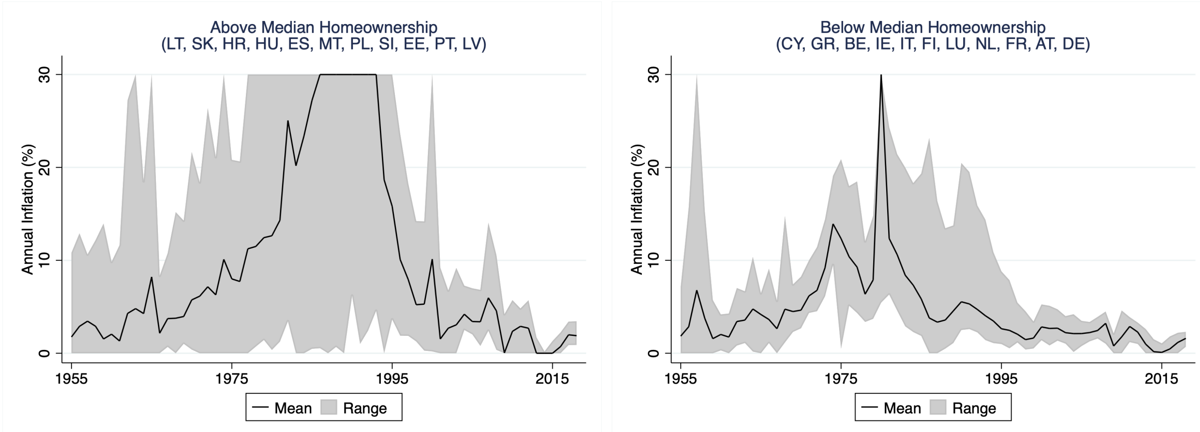

Figure 2.1: Inflation history above and below median homeownership rate

Note: Inflation data sources described in the text. Above and below median homeownership based on average country homeownership rate across all three HFCS waves. Figure plots the mean and range of inflation across listed countries. Inflation for chart capped above at 30% and below at 0%.

nomic and institutional conditions experienced by the population of potential homeowners strongly predict their investment in the housing market, even decades later, and above and beyond the influence of contemporaneous policies and institutions. Within the emerging literature on experience-based learning, we are the first to show that the effect of past exposures is strong enough to influence individuals even as they move to different countries with different market conditions and macroeconomic histories.

Our argument builds on the notion that past experiences of political, institutional, and economic conditions exert a longlasting influence on attitudes and beliefs (Alesina and Fuchs-Schündeln (2007) ,Luttmer and Singhal (2011),Giuliano and Spilimbergo (2009)). In our context, the conjecture is that exposure to high inflation triggers the desire to protect financial wealth from devaluation and encourages home purchases.

We present two types of motivating evidence for this conjecture. On the macro level, consider the relationship between homeownership and historical inflation in Figure 2.1. The left graph plots annual inflation for European countries with above-median homeownership rates (averaging at 81%); the right graph shows inflation in those countries with below-median homeownership (averaging at 56%).[3] The graphs illustrate that high-homeownership countries have witnessed significantly higher historical inflation over the past 60 years, which homeowners in the data have personally lived through.

On the micro level, a second piece of motivating evidence comes from a homeownership survey that we fielded in several European countries. We asked 700 homeowners in Austria, Germany, Ireland, Italy, Portugal, and Spain what they believe to be good reasons for buying

Chiuri and Jappelli (2003), Clark and Dieleman (1996), Doling (1973), Fisher and Jaffe (2003), Follain and Ling (1988), Gwin and Ong (2008), Haurin et al. (1996), Henderson and Ioannides (1987), Hilber (2007), Earley (2004), Ioannides (1987), Painter et al. (2001), and Sinai and Souleles (2005).

[3]Data from the 2008-2018 Household Finance and Consumption Surveys, discussed in detail below.

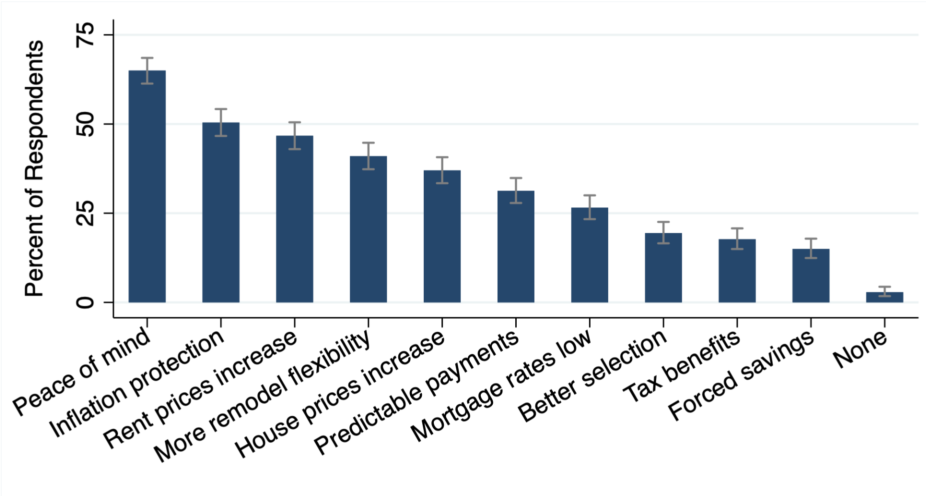

Figure 2.2: What do you think are good reasons for buying a home?

Note: Respondents were asked to select all options that apply. Order of options was randomized. Figure shows percent of respondents selecting each option and 95% confidence intervals. Survey responses from 700 homeowners in Germany, Ireland, Italy, Portugal, and Spain. See Appendix B.2 for more details.

a home. Out of 10 options, shown in Figure 2.2, inflation protection was selected by 50% of respondents, ranking second among all options (after "peace of mind"). When asked about their personal homeownership decision, a third of all respondents say that concerns about inflation impacted their own decision to buy, and the latter result is concentrated among homeowners who personally experienced high inflation (see Appendix-Figure B6).[4] The inflation-hedge motive turns out to be more important to homeowners than reasons such as tax benefits to owners, better selection of homes to buy versus rent, low mortgage rates, and even increasing house prices. In fact, if we take the motivation to "protect against inflation" and to "protect against rent price increases" together, concerns about price increases dominate all other categories (72% of respondents selected at least one of these options). In other words, regardless of whether or not real estate is indeed a suitable inflation hedge, as proposed in the classic Gordon (1962) growth model and as some authors argue empirically,[5] households appear to *believe* it to be true and important.

Our key research question in this paper is whether we can detect a long-lasting influence of past exposure to inflationary periods on the composition of housing markets. We test whether differences in inflation experiences over the past decades help predict household tenure choice, both across and within countries, and beyond the influence of other known determinants of individual tenure decisions.

We start from a simple theoretical framework that motivates the link between histories

[4] 283 respondents report experiencing high inflation, and 391 not. See Appendix B.2 for more details.

[5] Empirical tests of whether real estate and real estate investment trusts (REITs) act as inflation hedges have mixed results; cf. Anari and Kolari (2002), Brounen et al. (2014), Case and Wachter (2011), Fama and Schwert (1977), and Liu et al. (1997).

of past inflation, beliefs about future inflation, and homeownership. Building on recent formalizations of experience-based learning (cf. Collin-Dufresne et al. (2017), Malmendier et al. (2016b), and Schraeder (2015)), we assume that the histories of inflation an individual has personally experienced, both in general and in house prices, shape her beliefs about future realizations of the same variables.

The model illustrates two channels through which high inflation expectations induce a higher likelihood of homeownership: (1) the desire to protect oneself from high inflation and (2) the perceived attractiveness of a fixed-rate mortgage. As for the first channel, the underlying mechanism mirrors our survey results that households perceive real estate as an inflation hedge: When they experience an inflationary period, they anticipate higher future inflation, and therefore higher real rates of return on real estate, compared to non-inflation hedging assets. Similarly, if past experiences in the housing market induce households to anticipate higher future house prices, they are more likely to purchase their home today. As for the second channel, the perceived attractiveness of fixed-rate borrowing reflects that individuals who overestimate future inflation also expect higher nominal interest rates in the future, and therefore perceive fixed mortgage rates to be too low in real terms. As a result, they are more likely to purchase a home if they can finance it with a fixed-rate mortgage.

The second channel also generates a second model prediction: In countries with predominantly variable-rate mortgages, the link between prior exposure to inflation and the decision to purchase a (mortgage-financed) home should be weaker. We note that the latter prediction helps tease out a belief- versus a preference-based model of how macro-economic histories affect consumer decisions in the long-run: While prior experiences may also affect *preferences* for owning a home, those should be orthogonal to the prevalence of variable-rate mortgages.[6]

We test the model predictions using household microdata from 22 countries participating in the European Central Bank's Household Finance and Consumption Survey (HFCS). We collect historical inflation and house-price data for these 22 countries from Reinhart and Rogoff (2009), the International Monetary Fund, Global Financial Data, Apostolides (2011), Michal (1960), the Federal Reserve Bank of Dallas, Knoll et al. (2017), and Bordo and Landon-Lane (2014). For each country, we also build a dataset of housing-market conditions that are relevant to homeownership decisions, using data from Andrews et al. (2011) and the OECD. To measure the macroeconomic histories individuals in these countries have experienced over their lifetimes so far, we calculate each individual's weighted lifetime average experience at each point in time, using the learning-from-experience parameter estimates of Malmendier and Nagel (2015).

Our identification exploits variation in individual exposure to inflationary periods across three dimensions: age, country, and survey year. We first conduct an analysis of households with different histories of exposure that holds the housing market constant. The sample

[6]There is direct evidence on the belief-based mechanism in the empirical literature on experience effects, including Malmendier and Nagel (2015), Malmendier et al. (2016a), Malmendier and Shen (2017), and Botsch and Malmendier (2019). It is similar in spirit to other papers that explore the implications of potential homeowners not being fully rational, e. g., Glaeser and Nathanson (2015).

consists of 165,387 households in the American Community Survey (ACS) who immigrated to the U.S. from one of our sample countries (the same countries that we will analyze in the HFCS) and who make tenure decisions in the U.S. housing market. We find that inflation histories in their countries of origin significantly predict their homeownership decisions, even after they moved to a new and common housing market.

Turning to the HFCS data, we estimate the relationship between individual exposure to inflation and homeownership both across and within countries. We find that a 1 log-point increase in experienced inflation corresponds to an 18 pp increase in the likelihood of ownership for the average household. Moreover, as predicted by the model, experienced inflation is a stronger predictor of ownership in countries with access to fixed-rate financing. We also show that the effect is stronger among singles than couples, who may differ in their experiences.

Experienced inflation remains a significant predictor of homeownership after controlling for country fixed effects and even country-time fixed effects. The fixed effects specification rules out housing-market or other contemporaneous macroeconomic conditions as the sole driver of the link between macroeconomic histories and ownership. Relying on only the smaller within-country (or within-country-year) variation, we still find that experienced inflation significantly predicts homeownership. Our estimates imply that a one log-point increase in experienced inflation increases the predicted probability of homeownership for the average household by 3 pp. That is, even when we focus on individuals who have been exposed to the same country's macro-history and inflation, we find that personal exposure predicts significant, albeit smaller, differences in tenure decisions depending on which parts of a country's time-series occurred during an individual's lifetime.

The magnitude of the estimated effect is significant, even in comparison to the effect sizes of other factors discussed in the literature cited above. In the estimation controlling for housing-market measures, one log-point change in experienced inflation is associated with a change in homeownership roughly half as large as a one standard deviation change in measures of rent control, tax benefits to homeowners, and price-to-rent ratio and 1.2 times as large as a one standard deviation increase in buyer transaction costs.[7]

How do we interpret the estimated relationship between past exposure to inflation and homeownership decisions? On the one hand, the results support the hypothesis that macroeconomic shocks people experience in their home countries have a lasting impact on their decision making. The robustness of the results to the inclusion of country-time fixed effects, as well as their robustness across data sets, address alternative explanations based on housing-market features. The personal experience of seeing prices increase and the purchasing value of money fall appears to induce the willingness to invest in housing. On the other hand, other lifetime experiences may play a role. For example, countries with higher inflation over the last 20 years have had lower GDP per capita and household expenditure per capita.[8]

[7]For the effect sizes of these and other factors in the literature, see Andrews and Caldera Sánchez (2011), Andrews et al. (2011), Earley (2004), Fisher and Jaffe (2003), Gwin and Ong (2008), Hilber (2007).

[8]2000-2018 GDP and household expenditure data obtained from the World Bank.

While we cannot easily disentangle inflation experiences from all other macroeconomic experience effects, our focus on inflation reflects the prior evidence on inflation affecting inflation beliefs, interest rate beliefs, home purchases and mortgage decisions, including the new survey evidence discussed earlier. We do, however, consider prior experiences in the housing market as an additional predictor of homeownership decisions. Using a similarly constructed measure of lifetime experiences of *house* prices, we estimate a positive, albeit weaker and less robust, relationship compared to the effect of experienced inflation. One reason is the more limited availability of historical house-price data across countries, compared to inflation data. Even within the more limited set of countries and time spans, though, inflation experiences dominate house-price experiences. This may reflect the direct impact of house-price growth on affordability. In fact, we find that in countries where average house-price growth has been high, the price-to-rent ratio is higher than in countries with relatively low house-price growth.

Overall, our analyses reveal a novel, strong, and very robust determinant of the composition of housing market participants—the long-lasting effects of past inflation conditions experienced by the cohorts of potential homeowners. These influences appear to be a significant factor in explaining the large differences in housing markets across countries as well as the within-country changes over time.

Relation to previous literature. Our paper relates to the literature on the determinants of tenure choice, broadly classified as market factors and household characteristics. Among the market factors, homeownership has been linked to tax benefits, rent prices, transaction costs, housing supply, and other government policies.[9] Most closely related to our work, several of these papers argue that historical influences have a long-lasting impact on housing markets. Earley (2004) links the cultural tradition of passing property through family in Southern Europe to the high homeownership rates today and argues that the dowry laws in Greece continue to contribute to a culture of high homeownership, despite their repeal in 1983. Andrews et al. (2011) argue that differences in the timing and extent of historical mortgage market reforms help explain persistent cross-country differences in the availability of mortgage financing today. Our approach differs from these prior studies of historical influences in that we focus on a person's lifetime experiences. We show that personal experiences are predictive of homeownership even when controlling for current macroeconomic conditions, institutions, regulations, and country-specific cultural influences.

Household-level characteristics that have been linked to homeownership include household demographics (such as age, marital status, presence of children, and employment status), measures of household financial status (such as income, wealth, and access to mortgage debt), and preferences over types of home (e.g., apartment vs. single-detached unit).[10] We show that our analyses are robust to controlling for a vast array of both household and

[9] As discussed, for example, in Andrews et al. (2011), Andrews and Caldera Sánchez (2011), and Earley (2004).

[10] See for example Andersen (2011), Andrews and Caldera Sánchez (2011), Bracha and Jamison (2012), Collins and Choi (2010), and Drew and Herbert (2013).

market determinants.

Our paper also builds on the growing literature on experience effects, which shows that life experiences of macroeconomic events such as inflation and stock returns have significant impacts on expectations and financial decisions. Most closely related is the paper by Ampudia and Ehrmann (2014), who also use household data from the HFCS. They exploit its cross-sectional variation to demonstrate that experiencing higher stock-market returns increases households' self-reported tolerance of financial risk and their stock-market participation. Relative to the results estimated on U.S. data by Malmendier and Nagel (2011), European households tend to weight recent experiences more highly, i. e., exhibit stronger recency bias. They also find that "extreme" experiences have lasting effects on behavior. Relatedly, Ehrmann and Tzamourani (2012) find that the experience of hyperinflation has a lasting effect on beliefs about the importance of price stability. We note that in our data, the estimated effect is not driven by countries who experienced hyperinflation. Past work has also shown that inflation experiences predict future interest-rate expectations and the choice of fixed- vs. adjustable-rate mortgage financing (Botsch and Malmendier (2019)).

The rest of the paper proceeds as follows. Section 2.2 presents a simple theoretical framework demonstrating how inflation and house-price expectations can influence tenure choice. Section 2.3 describes the key data sets we employ. In Section 2.4, we present the analyses of the relationship between individual experiences of past (aggregate and house price) inflation and homeownership. Section 2.5 concludes.

2.2 Theoretical Framework

We present a stylized model of household tenure choice to demonstrate how macroeconomic experiences can influence the decision to rent or buy a home.

Real estate has classically been viewed as an inflation hedge, for example, in the seminal Gordon growth model (Gordon (1962)). Our model embeds the possibility of experience-based belief formation into Gordon's theoretical setting to analyze the perceived attractiveness of real estate as a real asset, as well as the perceived attractiveness of fixed-rate mortgages.

Consider an agent born at time t who lives for one period. The agent is endowed with wealth w_t and consumes all of her wealth at $t+1$. We distinguish between nominal and real values, and denote inflation in the price of consumption from t to $t+1$ as π_{t+1}. Agents have log utility over consumption in $t+1$, i. e., over real terminal wealth.

The decision of interest is the agent's choice between buying and renting a home to live in. Households maximize expected real terminal wealth subject to the constraint that they must either rent or own a home from t to $t+1$. Any wealth not spent on housing is invested in an alternative asset, which pays a nominal interest rate n_t. This assumption implies that housing is the only inflation-protected investment opportunity. We discuss below how our results differ in the presence of alternative inflation hedges. (See Appendix B.3 for details on these and other results.)

Rent. If the agent decides to rent her home, her expected utility is

$$\begin{aligned} E_t\left[U_{t+1}(R)\right] &= E_t\left[u\left(\frac{w_{t+1}(R)}{1+\pi_{t+1}}\right)\right] \\ &= \log\left((w_t - h_t)(1+n_t)\right) - E_t[\log(1+\pi_{t+1})], \end{aligned} \tag{2.1}$$

where h_t is the rental price at t and $w_{t+1}(R)$ is the expected nominal wealth in $t+1$ conditional on renting.

Buy. If the agent decides to buy a home, she pays the current house price M_t at time t, and sells the house at price M_{t+1} at time $t+1$. The change in house prices in each period can be decomposed into inflation in the price of consumption π and an exogenous housing-specific component g, $M_{t+1} = M_t(1+\pi_{t+1})(1+g_{t+1})$, where g_{t+1} is the real house-price growth between t and $t+1$.[11]

We assume the agent can finance a home purchase by borrowing amount $m_t \leq M_t$. Under a fixed-rate mortgage, she borrows at a nominal rate n_t^f, and then has to repay $(1+n_t^f)m_t$ in $t+1$. Under a variable-rate mortgage, she borrows at a real rate r_t^v (and thus nominal rate $(1+r_t^v)(1+\pi_{t+1})$), and has to repay $(1+r_t^v)(1+\pi_{t+1})m_t$.[12] We analyze each scenario separately.

Under fixed-rate financing, the expected utility of ownership is given by

$$\begin{aligned} E_t\left[U_{t+1}(FR)\right] &= E_t\left[u\left(\frac{w_{t+1}(FR)}{1+\pi_{t+1}}\right)\right] \\ &= E_t\left[\log\left(\frac{M_t(1+\pi_{t+1})(1+g_{t+1}) - m_t(1+n_t^f) + (w_t-(M_t-m_t))(1+n_t)}{1+\pi_{t+1}}\right)\right] \end{aligned} \tag{2.2}$$

where $w_{t+1}(FR)$ is expected nominal wealth in $t+1$ conditional on buying and financing with a fixed-rate mortgage at the prevailing prices.

The equation highlights the two channels through which expected inflation affects the value of homeownership. The first is the classic real-asset motivation: If house prices move with inflation, investment in real estate protects households from high inflation in the future. As inflation rises, the real value of the alternative investment is reduced while the real value of the home stays constant. The second motivation comes from a desire to borrow at a fixed-rate when inflation is high. With a nominal fixed rate, the *real* mortgage rate, $(1+n_t^f)/(1+\pi_{t+1})$,

[11]The exogenous process for home prices abstracts from house prices reacting to demand and supply. It allows us to illustrate the main effects of experience-based learning without complicating the model.

[12]In practice, variable-rate mortgages take many forms; here we assume the nominal rate adjusts one-for-one with inflation.

is decreasing in inflation. Therefore, homeownership is attractive as it allows households to borrow "cheaply."

Under variable-rate financing, the corresponding expression for the expected utility of ownership (relegated to Appendix-Section B.3) reveals that the real-asset motivation remains the same, but that the variable rate removes the latter channel: households no longer benefit from fixed-rate borrowing at what they expect to be a low real rate.

In the Appendix, we extend the set-up to include a cost of ownership which captures, for example, maintenance costs or property taxes (see Appendix-Section B.3). We also explore a larger parameter space, including "housing crisis" scenarios (see Appendix B.3).[13] In the main text, we restrict the analysis to the simpler setting that assumes that the value of the house tomorrow is greater than the outstanding loan for all possible realizations of inflation and of house-price growth.

Household Beliefs. At time t, the agent observes rental price h_t, house price M_t, mortgage rate n_t^f or r_t^v, current inflation π_t, and current (real) house-price growth g_t.

We allow beliefs to be influenced by agents' personal experiences: They believe that future inflation and (real) house-price growth will be more similar to what they have experienced in the past than rational learning implies. For our purposes, it suffices to assume that the inflation beliefs of an agent who has experienced high inflation at time t first-order stochastically dominate beliefs of an agent who has experienced lower inflation at time t. Similarly, an agent who has experienced higher real house-price growth at t has beliefs about g_{t+1} that first-order stochastically dominate the beliefs of an agent who has experienced lower g_t. (Beliefs about π and g are uncorrelated.) Agents take rent and mortgage rates as given and do not use them to draw inferences about future inflation or house-price growth.

Note that, in the empirical analysis, we will account for experiences over agents' actual lifetimes and for previously documented features of experience-based learning, including the weighting function used in experience-based learning models such Malmendier et al. (2016b). We will also allow all other historical data to matter; the key feature is that lifetime experiences receive some extra weight.

Predictions. Agents choose to buy or rent in order to maximize their expected utility, given their experience-based beliefs. We show how this decision is affected by past experiences.

Prediction 6. *Homeownership is increasing in experienced inflation.*

Proof. An increase in experienced inflation at time t shifts beliefs about $(t + 1)$ inflation to a first-order stochastically dominant distribution. Hence, homeownership is increasing

[13] In Appendix-Section B.3, we simulate the model to demonstrate the robustness of our predictions. We identify the bounds of the parameter space where each prediction holds, and then demonstrate the robustness of the predictions under a variety of alternative assumptions about the distribution of inflation beliefs and alternative levels of risk-aversion. We also use the simulations to demonstrate that these patterns do not hold if past inflation affects only the variance of inflation beliefs.

in experienced inflation if the expected utility difference between owning and renting is increasing in expected inflation. We check whether this difference is positive for any given realization of future inflation and future house-price growth, $\frac{\partial U(buy)-U(rent)}{\partial \pi_{t+1}} \geq 0 \ \forall \pi_{t+1}, g_{t+1}$, separately for each of the two mortgage types:

$$\begin{aligned} \frac{\partial}{\partial \pi_{t+1}} \left[U_{t+1}(FR) - U_{t+1}(R)\right]\bigg|_{\pi,g} &= \frac{M_t(1+g)}{w_{t+1}(FR|\pi,g)} \qquad (2.3) \\ &> 0 \quad \forall \pi, g. \\ \frac{\partial}{\partial \pi_{t+1}} \left[U_{t+1}(VR) - U_{t+1}(R)\right]\bigg|_{\pi,g} &= \frac{M_t(1+g) - m_t(1+r_t^v)}{w_{t+1}(VR|\pi,g)} \qquad (2.4) \\ &> 0 \quad \forall \pi, g \ \ s.t. \ \ M_t(1+g) > m_t(1+r_t^v). \end{aligned}$$

where $w_{t+1}(FR|\pi, g)$ and $w_{t+1}(VR|\pi, g)$ is the wealth in $t+1$ under fixed- and variable-rate financing, respectively.

Under fixed-rate financing, the derivative is positive for all possible realizations of future inflation and future house-price growth. Under variable-rate financing, the derivative is positive under our assumption that $M_t(1+g) > m_t(1+r_t^v) \ \ \forall g$. Thus, the expected utility difference is also increasing in experienced inflation. We simulate the model in Appendix-Section B.3 to confirm that this prediction is robust to a broader parameter space. □

Our second prediction hones in on the difference between variable- and fixed-rate financing, namely, that households no longer benefit from a perceive-to-be-low real rate under variable-rate mortgages:

Prediction 7. *Among households with comparable wealth, the effect of experienced inflation is weaker for households who only have access to variable-rate mortgages.*

Proof. We compare the magnitudes of the point-wise derivatives in equations (2.3) and (2.4). Assuming $(t+1)$-wealth is similar when financing with either mortgage ($w_{t+1}(FR) \approx w_{t+1}(VR)$ for any π and g), homeownership will react more to experienced inflation under fixed- than variable-rate financing as (2.3)-(2.4) $\approx m_t(1+r_t^v) > 0$. We show that this prediction also holds without the similar-wealth assumption using simulations under a broad range of conditions in Appendix-Section B.3. □

Thus far, we have focused on the effect of past periods of inflation on housing markets. Our model also makes a clear prediction about the effect of past house-price growth.

Prediction 8. *Homeownership is increasing in experienced real house-price growth.*

Proof. The utility of ownership is strictly increasing in g, while the utility of renting is independent of g (see Appendix B.3). Therefore, a first-order stochastic dominating shift in beliefs about g unambiguously increases homeownership. □

In Section 2.4, we test these three predictions, relaxing some of the theoretical simplifications of our model. For example, we control for household characteristics that may shift the relative utility of ownership (e. g., family structure) or ability to buy (e. g., income and wealth). We also control for factors that may shift the relative cost of ownership, c, including tax benefits and tenant protections. Controlling for variation in homeownership rates due to these factors, we test whether prior macroeconomic realizations have a long-lasting effect on homeownership by exploiting variation in the exposure to past macroeconomic realizations and in access to different types of mortgages across cohorts and countries.

To capture households' access to variable- versus fixed-rate mortgages in Prediction 2, we would ideally measure the supply of different types of mortgages. Our empirical proxy will rely on the prevalence of variable-rate mortgages in a country.

We also note that, while variable-rate financing shuts down the cheap-borrowing motivation for ownership, we can further shut down the real-asset channel by allowing for inflation-protected investment other than housing. In Appendix-Section B.3, we show that, with an alternative inflation hedge, Prediction 1 continues to hold for fixed-rate financing, through the perceived cheap borrowing motivation, but there is no predicted relationship under variable-rate financing. Therefore, in the presence of alternative inflation hedges, Prediction 1 is weakened (homeownership is only weakly increasing in experienced inflation) while Prediction 2 remains robust.

This discussion implies yet another prediction of our model: the effect of experienced inflation should be weaker in markets with access to alternative inflation hedges. To empirically test this prediction, we would need a convincing measure of households' access to alternative inflation hedges and its variation over time and across countries. Lacking such a measure we leave further exploration of this prediction to future research.

2.3 Data and Empirical Measures

Data Sets

American Community Survey (ACS) Data. We obtain data on households in the American Community Survey (ACS) from IPUMS (2020). The analysis of this sample will allow us to hold the housing market, in which individuals make their tenure decisions, constant while varying macroeconomic histories due to different countries of origin. We include all households headed by individuals aged 20-80 who immigrated to the U.S. from one of the 22 European countries we study in the HFCS.[14] This main sample consists of 165,387 immigrant household heads surveyed from 2008 to 2018. For robustness, we replicate the results also including all American-born household heads in the ACS data. Homeownership (defined

[14]This includes 15,663 household heads who immigrated in the implied birth year (calculated as the survey year minus age) or in the year prior to implied birth year (which we assume is due to the timing of the survey relative to birth date, and we treat as immigrating in the birth year). We exclude 40 household heads whose year of immigration is more than 1 year earlier than the implied birth year.

as owning the property where surveyed) among this immigrant population is approximately 69%, which is slightly higher than the U.S. native average of 66%.

To construct household experience measures, we use information on the head of household's birth country, age, and year immigrated to the U.S.. As demographic controls, we also use data on household head age, gender, marital status, children in the home, educational attainment, employment status, and total household income. Unlike the HFCS data, the ACS does not survey households on their total wealth.

The data are summarized in Panel A of Table 2.1. The average household head is 53 years old. 54% of household heads are male and 37% have children living in the household. The average household income, expressed in 2010 USD, is about 85,000. 56% of household heads are married (29% to a native spouse), 15% are single, and the remaining heads are widowed or divorced. We group educational attainment into three buckets: 39% of household heads have obtained a Bachelor's degree or higher, 52% completed high school (includes household heads who completed grade 12, have a high school diploma or GED, completed some college, or have an Associate's degree), and the remaining 9% have not completed high school. 65% of household heads were employed at the time of the survey, 3% unemployed, with the remainder retired or out of the workforce.

All analyses using the ACS data are weighted using household weights designed to be representative of the U.S. population.

Household Finance and Consumption Survey (HFCS) Data. We obtain three waves of household-level microdata from the Eurosystem Household Finance and Consumption Network's Survey (HFCS). Conducted by the European Central Bank (ECB), the goal of the HFCS is to collect harmonized data on finances and consumption across the euro area, that is representative at both the euro area aggregate and the individual country level. The target population is all private households and their current members residing in the national territory. The first wave was conducted in 2008-2011 (primarily in 2010) and includes 15 countries: Austria, Belgium, Cyprus, Finland, France, Germany, Greece, Italy, Luxembourg, Malta, the Netherlands, Portugal, Slovakia, Slovenia, and Spain. The second wave of the survey was conducted in 2011-2015 (primarily in 2014) and, in addition to the 15 countries from the first wave, also includes Estonia, Ireland, Latvia, Hungary, and Poland. The third survey was conducted in 2016-2018, covers all earlier countries with the exception of Spain, and adds Croatia and Lithuania.

Table 2.1, Panel B, shows the summary statistics of household characteristics in the merged HFCS data, based on the microdata from each country and weighted to be representative both within and across countries. Our sample includes 220,697 households across 22 countries. The average household head is 51 years old. 55% of household heads are male, and 45% have children. 54% of household heads are married (or in a consensual union on a legal basis), 24% are single, and the remaining heads are widowed or divorced. The highest level of education attained by the household head is measured in the HFCS using the International Standard Classification of Education (ISCED-97), with 27% of household heads being educated at the tertiary ISCED-97 level (college in the U.S.), 44% at the upper secondary level (high school in the U.S.), and the remaining 29% at the lower secondary level or below. 56% of household heads were employed at the time of the survey, 6% unemployed, with the remainder retired or out of the workforce.

We measure net wealth and total gross income at the household level, converting all monetary values to 2010 euros using country-specific inflation to the reference year.[15] The average net wealth, in 2010 euros, is about 200,000, and the average household income is about 37,000. In our analyses, we will use deciles of wealth and income, calculated across survey respondents. (We test the robustness of our analyses to several alternative specifications of wealth and income, described in Section 2.4.) We focus on HFCS household heads aged 20-80 at the time the surveys were conducted.[16] In our main analyses, we include all households surveyed, regardless of where the household head was born. The identification of natives is not available in all country-waves. When available, the data identifies almost 90% of are natives. Our main results are robust to limiting analyses to natives.[17]

Note that despite vastly different sample selection of the HFCS than the ACS sample (namely, natives in their home countries versus immigrants to the US), the similarities in household characteristics across the two samples are striking. The close match in terms of demographics such as age, gender, or marriage status might reflect country-of-origin effects as individuals in both samples come from the same countries.

For the HFCS data, we also show homeownership data by country. The leftmost column of Table 2.2 reveals a wide variation in homeownership rates across the 22 HFCS countries, as also illustrated graphically in Appendix-Figure B1. For example, less than half of households own their main residence in Austria and Germany, while homeownership rates are above 80% in countries like Croatia, Slovakia, Hungary, and Spain. The middle columns of Table 2.2 indicate the information on housing market characteristics for each country, which comes from other data sources that we will discuss below.

The rightmost columns of Table 2.2 summarizes our HFCS measure of the prevalence of variable-rate (relative to fixed-rate) mortgages. It is defined, separately for each country-wave, as the sum of all adjustable-rate mortgages on households' main residences (in euros outstanding at the time of the survey) divided by the sum of outstanding mortgages carrying either an adjustable- or fixed-rate.[18] In the second to last column, we display the percentage of mortgage-euros in adjustable-rate mortgages for each country, averaged across survey waves using the sum of household weights. We normalize this measure to have a mean of 0 and variance of 1 across the sample and summarize in the last column of Table 2.2.

[15] For households missing a survey date, we assume the survey was conducted at the start of the fielding period.

[16] For about 6% of households, we have only 5-year age buckets, in which case we use the midpoint of the age bucket.

[17] The ECB does not provide the country of origin for non-natives. This data is available in the ACS and the focus of our empirical tests there.

[18] This measure is not defined for Finland, where all mortgage rates are reported as unknown. In all other countries, the percentage of mortgage dollars with unknown rates is minimal.

We note that, to test Prediction 2 of our theoretical model, we would ideally measure the "availability" of fixed- and variable-rate mortgages, i. e., mortgage supply. Practically, our proxy is an equilibrium measure of the prevalence of variable-rate mortgages in each country. We also note the implicit assumption in the empirical analysis that variable-rate mortgages are linked to inflation in some way (e.g., targeting a real interest rate), and we recognize that there are other forms of variable-rates.

Inflation Data. Our primary source of historical inflation data is Reinhart and Rogoff (2009), who provide time series of consumer price indices (CPI) for a large number of countries until 2010. We extend the data to 2019 using inflation data from the International Monetary Fund (IMF). We note that the calculation of the CPI by the Bureau of Labor Statistics is meant to capture housing costs (Greenlees and McClelland (2008)) and has historically included house prices, while its more recent design targets housing consumption rather than investment.

For several countries not included in the Reinhart and Rogoff data, we use alternative historical inflation data. For Cyprus and Malta, we use data from Apostolides (2011) for inflation from 1922-1938 and Global Financial Data (GFD) from 1943 on for Cyprus and from 1947 on for Malta. For Luxembourg, we obtain GFD inflation data extending to 1922. For countries formerly part of the Soviet Union (Estonia, Latvia, and Lithuania), we obtain GFD for the individual countries when possible, and from the Soviet Union when the individual series are not available (1945-early 1990s). Data for Estonia and Latvia begin in 1922. Data for Lithuania begins in 1925. Similarly, for Croatia and Slovenia, we use GFD data for Yugoslavia from 1929-1943 and from the individual countries from 1952 on. For Slovakia, we use GFD data from Czechoslovakia from 1922-1948, cost-of-living index data from Michal (1960) from 1949-1959, and GFD data for Slovakia from 1964 on. For countries with gaps in the inflation series (ranging from 1 to 8 years), we linearly interpolate missing values over inflation rates.

House Price Data. We obtain our house-price indices from several sources.

Real house-price indices from 1975 onward are available from the Federal Reserve Bank of Dallas for 10 of the 22 ECB countries: Belgium, Croatia, Finland, France, Germany, Ireland, Italy, Luxembourg, the Netherlands, and Spain.[19] The house-price index for each country is chosen (by the Dallas Fed) to be most consistent with the quarterly U.S. house-price index for existing single-family houses produced by the Federal Housing Finance Agency, and is seasonally adjusted. With this data, we cannot compare relative house prices across countries so, instead, we compare house-price growth. Using the fourth quarter index values, we calculate annual house-price growth in each country and construct a (partial) measure of experienced real house-price growth.

Using historical house-price index time series from Knoll et al. (2017) and Bordo and Landon-Lane (2014), we are able to complete the time series to full experience measures

[19]The authors acknowledge use of the dataset described in Mack et al. (2011).

for 6 of the countries in our HFCS sample. From Knoll et al. (2017), we obtain nominal house-price indices for Belgium, Finland, France, Germany, and the Netherlands which we convert to real house-price growth using the inflation data described above. For Spain, we use the real house-price index data of Bordo and Landon-Lane (2014).

Housing Market. We obtain country-level measures of housing-market characteristics as of approximately the time of the ECB surveys, summarized in the middle columns of Table 2.2. We normalize all continuous comparative housing market measures to have a mean of 0 and variance of 1 in our sample.

From Andrews et al. (2011), we obtain comparative measures of tenant protection, rent control, and tax benefits to homeowners, as well as transaction costs calculated as the average cost associated with purchasing a home. We obtain annual price-to-rent ratios from the OECD for 11 of the sample countries. The baseline for each country is set equal to the country's long-run average price-to-rent ratio, where the long-run is defined as starting in 1980 or averages all available data if the data begins after 1980.

Measures of Exposure to Past Inflation

As measures of individual exposure to past price increases, we calculate weighted averages of annual inflation over each individual's past life. We follow prior literature in incorporating recency bias by assigning price changes in the most recent past the highest weights. In our main specification, we let weights decrease linearly down from the year before the survey to zero at birth. In a robustness check, we implement a modified version of Malmendier and Nagel (2015), where individuals use their past inflation experiences to recursively estimate an AR(1) model of inflation to generate one-year inflation forecasts, and extend it to longer forecasting horizons appropriate for homeownership decisions.[20] That is, our main measure of the experienced inflation history of household i as of year t is calculated as

$$\pi_{i,t} = \frac{\sum_{k=1}^{age_{i,t}-1} w_{i,t}(k)\pi_{t-k}}{\sum_{k=1}^{age_{i,t}-1} w_{i,t}(k)} \quad \text{with weights} \quad w_{i,t}(k) = age_{i,t} - k.$$

For ACS households, we calculate this measure using the past inflation rates the household heads have experienced in their birth country from the year of birth to the year of immigration to the U.S., and U.S. inflation thereafter. For HFCS households, we use inflation in their countries of residence.

Table 2.3 shows the summary statistics of actual historical inflation from 1930-2018 as well as experienced inflation of all households in each country (of origin) in the ACS and

[20] The AR(1) model is not immediately applicable here as homeownership decisions are likely based on beliefs about inflation over longer periods, and extending the AR(1) model to a long-term inflation forecast requires taking a stance on the relevant forecast horizon (e.g., inflation over 5, 10, 20 years) as well as how individuals forecast forward, e.g., whether they iterate the AR(1) forward, apply the 1-year forecast to all future periods, anticipate learning in the future etc. For these reasons, we leave this approach as a robustness exercise, see Section 2.4.

HFCS samples. Both the historical rates and individuals' exposure to past inflation (over their respective past lifetimes) average in the single digits for the majority of countries. While some countries (33-36% in terms of experienced inflation) feature much higher rates, the respective populations account for a small fraction of the sample, as illustrated in Figure 2.3: The left panel shows the distribution of lifetime average inflation for HFCS households. While the vast majority of households (78% unweighted and 91% of the weighted sample) has experienced inflation under 10%, the distribution has a long right tail. (Appendix-Figures B2 and B3 show the country-by-country histograms.) To reduce the influence of these outliers in our main analyses, we apply a log transformation to the household measure of experienced inflation, $\ln(\pi)$. The resulting distribution is shown in the right panel of Figure 2.3 for the HFCS sample.

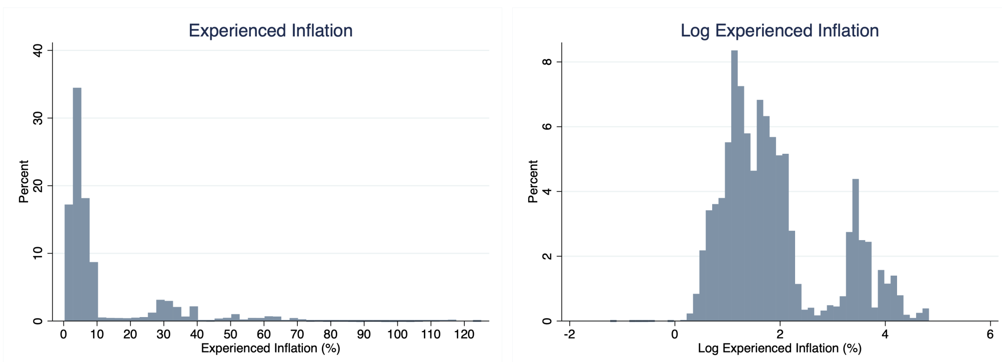

Figure 2.3: Distribution of experienced inflation

Histograms plot the distribution of experienced inflation (left) and log experienced inflation (right) in the HFCS sample.

Table 2.3 also reveals that, in some countries, inflation histories experienced by the current population differ substantially from the long-term historical averages, depending on the age structure. For example, while inflation in Italy averages at 11.3% over the last 90 years, the weighted average of inflation the Italian population has been exposed to over their respective lifetimes is only 4.9% in the HFCS data. Vice versa, the lifetime experiences of people in Poland or Estonia have been much higher (27.2% and 47.9% in the HFCS) than averaged historical inflation (19.8% and 31.5%). The same pattern is visible in the ACS data, albeit with some common exposure to U.S. inflation folded into the average. Across individuals in all countries, experienced inflation averages at 7.5% in the ACS sample and 6.8% in the HFCS sample.

We construct parallel measures for house prices. The data here is more limited. First, we obtain a complete house-price history for households in six countries from Knoll et al. (2017) and Bordo and Landon-Lane (2014). Second, we use real house-price data from the Federal Reserve, which covers nine sample countries, but is limited to house prices from 1975 onward. We construct a partial measure by re-scaling weights for the available data years. Here, the identifying variation does not utilize differences in experiences prior to

1975, but comes from the feature that recent years will matter more for younger than older households. Appendix-Figure B4 shows the distribution of both measures of experienced house-price growth across countries.

We note two assumptions implicit in the construction of our measures. First, we allow households to continuously update based on their experiences and to re-evaluate their tenure status. An alternative assumption is that, once a household has purchased a home, they will not re-evaluate their tenure status based on their most recent inflation exposure, i. e., that homeownership is sticky. In Section 2.4, we will use SHARE data, which identifies when individuals first become homeowners, to test whether prior experiences predict when individuals first purchase their home.

Second, we consider the past experiences of the household head as the relevant determinant. This is likely to be a good proxy for the relevant macroeconomic experiences in single households, in married households with spouses from the same country and of similar age, and if the head primarily makes the financial decisions. We expect the measure to be a noisier proxy among married households. Indeed, we will show below a stronger response to experienced inflation among singles (see Table 2.6).

2.4 Empirical Analysis

In this section, we test the predictions derived in Section 2.2. First, we demonstrate that households who have lived through periods of higher inflation are more likely to be homeowners (Prediction 1), both in the U.S. housing market using the ACS data, and within and across European countries using the HFCS data. We then show that the response to experienced inflation is weaker in countries with predominantly variable-rate financing (Prediction 2) and in married household mixed experiences. Finally, we test Prediction 3 by examining the relationship between experienced house-price growth and ownership.

Within the U.S. Housing Market (ACS)

We start by testing Prediction 1 in the American Community Survey (ACS): Is there a positive relationship between past inflationary periods that immigrants to the U.S. have experienced in their different countries of origin and their likelihood of becoming a homeowner, even after they have left their home country and face an *identical* housing market in the U.S.?

We first estimate logit regressions of homeownership on past exposure to inflation. We control for age, gender, educational attainment (below high school, high school, and four or more years of college), employment status (employed, unemployed, and not in the labor force), marital status (single, married, widowed, or divorced), an indicator equal to 1 if the household head is married to a U.S. native, children living in the home, and decile of household income, where income is adjusted for inflation over the survey years and deciles are calculated out of our entire ACS sample (including U.S. natives). In addition to the demographic controls, we control for years lived in the U.S. as a percent of the household head's life. In all regressions, we also control for survey year fixed effects. We weight the data using ACS-provided representative weights and report heteroskedastic-robust standard errors.

Table 2.4 reports the results of this analysis. Coefficients are shown as odds ratios (exponentiated coefficients), so that an estimate above 1 indicates a positive relationship and an estimate below 1 a negative relationship. In column (1), we find that a 1 log-point increase in experienced inflation predicts a 22 pp increase in the odds of ownership among immigrants to the U.S. That is, for a household head with a baseline probability of ownership of 65% (roughly the average in our sample) our estimation predicts that increasing one log-point in experienced inflation (say from 2% to 5.4%), would increase the likelihood of ownership to 69%.

In column (2), we include as an additional control the homeownership rate among other (non-immigrant) households in the same state and year. This estimation addresses the concern that the positive estimate of past exposure to inflation might be (partly) explained by immigrants with higher inflation experiences moving to higher homeownership states. The coefficient estimate of the additional control confirms that immigrants are indeed more likely to own a home if they move to a state with a higher homeownership rate. However, conditional on the homeownership rate in their state and year, immigrants with higher experienced inflation are still more likely to own a home. In fact, the magnitude and significance of the coefficient remain very similar.

The estimates are also robust, though attenuated, when including country-of-birth fixed effects, as shown in column (3). In this specification, we are identifying the effect of experienced inflation solely off differences in personal exposure differences across household heads born in the same country. While country-of-birth fixed effects help rule out confounds with cultural differences and other country-specific attributes, they also eliminate the average differences in inflation experiences across countries, which is a valid source of identification.

Finally, we re-estimate the relation between experienced inflation and homeownership on a larger sample, which includes households headed by a U.S. native. Here, we control for state and state-year fixed effects as well as an indicator for being born in the U.S. (column 4) or even the full set of birth-country fixed effects (column 5). We continue to find a positive significant effect of past exposure to inflation. The magnitude of the coefficient estimate increases quite dramatically when we include U.S. natives in the analysis, most of whom have experienced relatively low inflation compared to the immigrant population (The average inflation experience of U.S. natives is 3.3% compared to 7.5% among our European immigrant population.). A one log-point in experienced inflation (e.g., 2% to 5.4%), corresponds to a predicted increase in the likelihood of ownership from 65% to 80-82%. While a one log-point change is well-within the distribution of experienced inflation among immigrants, the standard deviation of experiences among U.S. natives is only 0.5%.

In summary, from the ACS data we find that variation in experienced inflation predicts variation in homeownership both among immigrants and U.S. natives. Building on this analysis, we expand to a cross-country analysis of homeownership in Europe using household survey data from 22 European countries.

Within and Across European Markets (HFCS)

Before leveraging the wealth of information in the household-level HFCS data, we examine whether Prediction 1 holds in the aggregate: Does the population average of lifetime exposure to inflation predict aggregate homeownership across countries? We collapse the HFCS data into country averages, using the survey weights representative of the population. We then weight countries by average population across survey years (from the World Bank).

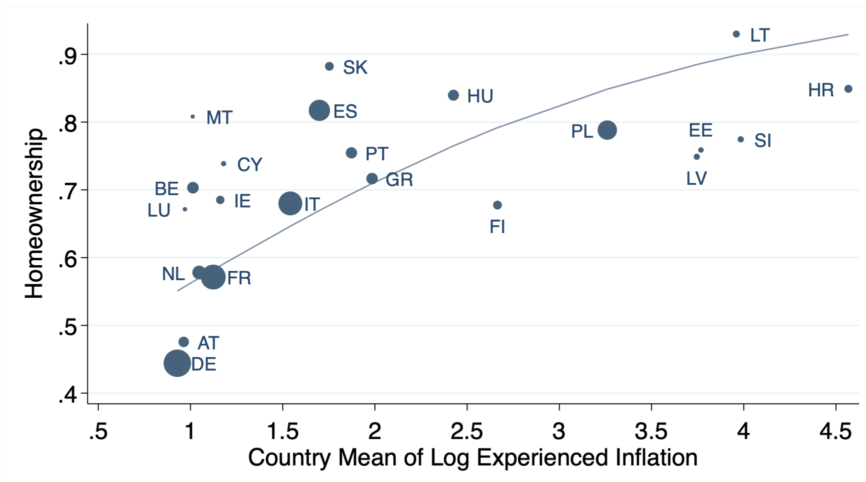

Figure 2.4: Aggregate homeownership rates by experienced inflation

Scatter plot of country average of log experienced inflation (x-axis) and homeownership rate (y-axis). Country abbreviations are in Table 2.2. Size indicates relative population. Line shows the population-weighted logit fit of a regression of homeownership on country average of log experienced inflation.

Figure 2.4 shows the resulting relationship graphically. The scatter diagram plots the country averages of log household experienced inflation on the x-axis and the homeownership rates (percent of households living in owner-occupied housing) on the y-axis for each country in our analysis. The points are coded by the size of the population. The plot reveals a positive relationship between experienced inflation and homeownership, which we confirm by fitting a logit regression to the country-level averages (estimated odds ratio of 1.92 (s.e. 0.36)). The magnitude of the predicted relationship is striking. For example, a roughly one log-point increase in experienced inflation from the Netherlands to Greece is associated with a 14 pp higher homeownership rate.

Turning to the detailed household-level data, we related individual differences in past exposure to inflation to the likelihood of homeownership, controlling for household characteristics, housing market features, and other time- and country-specific effects. In these analyses, we have variation in experiences both across individuals in different countries and also across individuals within a country (by age and survey year).

Our key dependent variable is a binary indicator of whether the household owns their primary residence.[21] The key independent variable is household experienced inflation, calculated using the household head's age, country, and survey year as described above. We include controls for household demographics that are likely to be related to homeownership: age, gender, having children, marital status, educational attainment, employment status, decile of net wealth, and decile of household gross income.[22] We also control for survey wave. We use the HFCS multiple-imputation data and the corresponding estimation techniques from Rubin (2004) to include the full imputed sample in our analyses, despite some households having missing data. In all analyses, we use the HFCS household weights that are representative of each country and the EU population (inverse probability of being sampled and non-response).

In Table 2.5, we report the odds ratios and standard errors from logit regressions. (The results are robust to probit and OLS specifications.) All of our analyses control for the full array of household demographic variables, reported in more detail in Appendix-Table B2. The estimated odds ratio of 2.66 in column (1) implies that, controlling for household demographics, a one log-point increase in experienced inflation predicts a 166% increase in the odds of being a homeowner. This relationship is quite large. A household with a likelihood of ownership at 65% (roughly the average homeownership rate in our sample) has an odds ratio of 1.86. A 1 log-point increase in experienced inflation (e.g., from 2% to 5.4%) predicts an increase in the odds ratio to 4.70, which corresponds to an 82% probability of ownership. Among the controls, being married or widowed, having children, and being employed or retired are robust predictors of higher homeownership. We also find that homeownership is increasing in wealth, but conditional on wealth, is decreasing in income.

One way to illustrate the magnitude of the estimated effect is to calculate the implied counterfactual homeownership rate if a country had a different inflation history. The hypothetical counterfactual abstracts, of course, from general-equilibrium considerations and serves merely as a back-of-the envelope calculation, though we do account for differences in population age structures. Figure 2.5 provides three examples. The left panel shows how homeownership in Austria and Greece would change if we switched their inflation histories. These countries are at opposite ends of the spectrum, with Austria having a homeownership rate under 50% (first bar) and a low inflation history, while Greece has a homeownership rate of about 70% (second bar) and a high inflation history. Had Austrians experienced Greece's inflation history, our estimates would predict a substantial 14 pp increase in the homeownership rate for Austria (third bar). Likewise, if Greece had Austria's inflation history, we would predict a 15 pp drop in homeownership (fourth bar). As a second example, consider Italy and Portugal, shown in the middle. Their homeownership rates differ less, by 7 pp. Our model predicts that the gap would almost completely close if these countries had the same inflation history. Finally, we pick an example where the hypothetical change in inflation histories would neither switch nor even out the cross-country differences in homeownership.

[21]Our main results also hold if we define the dependent variable as owning any property.

[22]Wealth and income are converted to 2010 euros, and deciles are calculated across the entire sample.

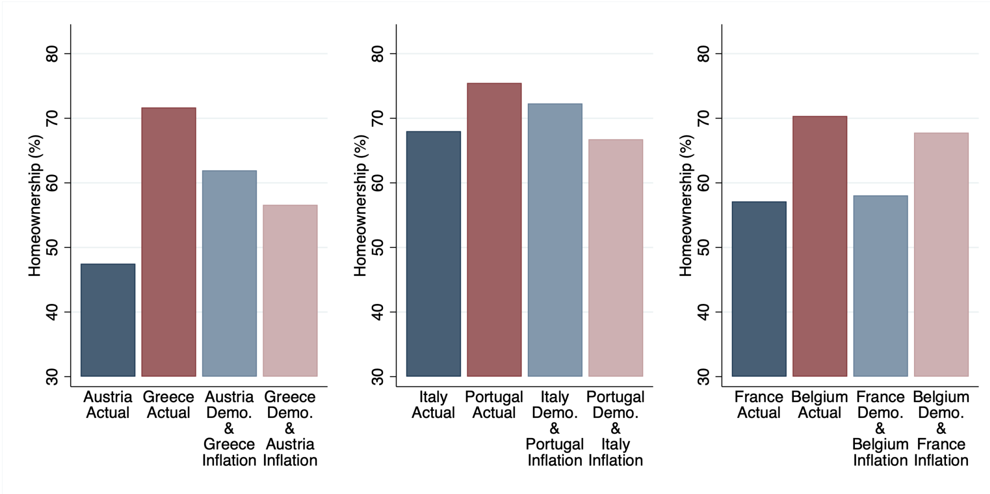

Figure 2.5: Hypothetical homeownership rates with alternate inflation histories

Actual homeownership from the HFCS data. Hypothetical homeownership rates calculated using the model estimated in Table 2.5, column (1), assuming another country's inflation history.

The last panel compares France and Belgium, where our model predicts that the large gap in homeownership would persist if inflation histories were switched. Even in that case, though, we predict a substantial role of experienced inflation, with a reduction in the homeownership gap of almost 4pp.

In the baseline estimation, we have abstracted from country differences in housing regulations, the housing stock available for purchase vs. renting, transaction costs, or cultural features that make ownership more or less appealing. In column (2) of Table 2.5, we test the robustness of Prediction 1 to the inclusion of country-level measures of tenant protection, rent control, tax benefits to homeowners, buyer transaction costs, and price-to-rent ratio. The estimated effect remains large: a one log-point increase in experienced inflation change predicts an increase in ownership from 65% to 72%. To benchmark this magnitude, note that, from a baseline of 65%, 1 SD higher rent control reduces predicted homeownership to 54%; 1 SD higher tax benefits increase predicted ownership to 76%; 1 SD higher buyer transaction costs reduce it to 61%; and a 1 SD higher price-to-rent ratio reduces it to 57%. Interestingly, controlling for the other factors, a 1 SD increase in tenant protection is associated with *higher* homeownership (increase to 67%).

In column (3), we go one step further and eliminate any country differences as a source of identification by including country fixed effects in the empirical model. This specification addresses concerns about unobserved differences such as, for example, historical difference arising from former Soviet countries having different ownership structures.[23] However, the specification also eliminates the average difference in experienced inflation across countries as

[23] In our data, average homeownership in countries that were formerly in the Soviet sphere of influence is significantly higher if we exclude Germany (Croatia, Estonia, Finland, Hungary, Latvia, Lithuania, Poland, Slovenia, Slovakia) and significantly lower if we include Germany.

a source of identification and only tests whether inflation exposure predicts homeownership *within* country. The resulting estimates reveal that a one log-point increase in experienced inflation predicts a 12% increase in the odds of homeownership (or for example, a change in probability of ownership from 65% to 68%). That is, even within country, the magnitude of the relationship between homeownership and experienced inflation is large.

Finally, in column (4), we also control for potential time-varying differences across countries and housing markets by including country-wave fixed effects. While the sample period over the three HFCS waves is relatively short, we do observe some changes in homeownership and inflation exposure across survey waves, as illustrated graphically in Appendix-Figure B5. The effect size remains roughly the same: A one log-point increase in experienced inflation is associated with a change in likelihood of ownership from 65% to 68%.

Heterogeneity in Within-Household Experiences. The hypothesis that past experiences shape beliefs and household financial decisions has several additional implications for homeownership decisions. One is that the predictive power of past exposure to inflation should be weaker when decisions are made jointly by two people, i. e., couples (with potentially different experiences). This prediction further distinguishes the experience-effect hypothesis from alternative explanations. If the correlation between experiences and homeownership is driven by unobserved correlations with market factors or financing opportunities, it is not clear why we should expect the relationship to vary across single and married household heads.

In columns (1) and (2) of Table 2.6, we test this hypothesis by interacting log experienced inflation with an indicator for being married. We limit the sample to households heads who are either married or single, excluding those who are widowed or divorced. Consistent with experience effects as a mechanism, in column (1) of Table 2.6 we find that indeed the estimated odds ratio on the interaction between experienced inflation and married is less than 1, implying that the effect of experienced inflation is weaker among married household heads. In column (2), we include country-wave fixed-effects and confirm this result using only the within-country-wave variation in experiences.[24]

Testing Prediction 2: Access to Fixed-Rate Financing. Another prediction of the model is that the relation between prior inflation histories and homeownership is not as strong in countries with variable-rate mortgage financing. Homeownership still provides an inflation hedge in that case, but mortgage financing is not perceived as more advantageous by those who have lived through periods of high inflation. Instead, all potential home buyers agree on the (real) cost of mortgage financing.

[24]We have also calculated past inflation exposure separately for the household head and the spouse. When we include both in the estimation, both coefficients positively predict ownership (with odds ratios of 1.81 for household heads and 1.26 for spouses). Controlling for country or country-wave fixed-effects, we lose statistical significance though the magnitudes are similar to what we find in Table 2.5 (odds ratios of 1.17-1.22 for household heads and 1.12-1.14 for spouses).

To proxy for the availability of fixed- versus variable-rate mortgages, we use a normalized measure of the prevalence of variable-rate mortgages (PVR) derived from the amount of euros outstanding in variable- vs. fixed-rate mortgages reported in the HFCS. To the extent that this measure reflects relative supply, rather than demand, of variable-rate mortgages, we predict that in countries with higher PVR prior inflation histories should have less predictive power (Prediction 2).

We test this prediction by estimating model specifications that mirror those of columns (1) and (2) of Table 2.6: In column (3), we include the PVR measure as well as its interaction with households prior exposure to inflation in the estimating equation. We find that, while past inflation exposure continues to predict an increased likelihood of homeownership, but the effect is attenuated in countries with higher prevalence of variable-rate mortgages, as indicated by the estimated odds ratio below 1 for the interaction with PVR. In column (4), we include country-wave fixed-effects, which subsume the level effect of PVR. The estimated interaction effect remains negative and statistically significant.

Our result confirming Prediction 2 is robust to many alternative specifications described in Section 2.4 (e.g., controlling for survey year, alternative age or cohort controls, and estimation using OLS or probit).

We re-iterate that the distinction between countries with predominantly variable-rate versus predominantly fixed-rate mortgages is an imperfect measure of supply.

For example, the composition of mortgages might affect *access* to financing and thus the homeownership rate. Or, vice versa, homeownership rates might influence the composition of mortgages in a country, e. g., since marginal homeowners are more likely to have variable-rate mortgages. Both channels would explain the positive coefficient estimate on the level effect of PVR (in column 3). It is less clear why these channels would generate an interaction effect with inflation exposure. That is, while the available data does not allow us to distinguish between the above-mentioned (and further) channels, the estimation results on the effect of past exposure to inflation and its interaction with PVR are consistent with and corroborate the hypothesis that past inflation histories have significant influence on the composition of real-estate markets within and across countries.

Testing Prediction 3: House-Price Histories and Homeownership. In addition to experienced inflation, prior experiences to rising or falling real-estate prices may also predict homeownership (Prediction 3). As illustrated in the model in Section 2.2, households that have lived through periods of high real house-price growth might believe house prices will continue to grow in the future and therefore value ownership. Other channels could generate a similar relationship. For example, experienced house-price growth may also influence preferences for homeownership or risk-taking in addition to beliefs.

Historical data on house prices is scarcer than historical data on inflation. The available data allows us to construct the two measures of experienced real house-price growth, outlined in Section 2.3: one uses a full history of lifetime experiences for a smaller sample, and one captures only a partial history for a larger sample. For both measures, we do not apply the log transformation since, unlike experienced inflation, a non-negligible fraction of households have experienced *negative* real house-price growth (see Appendix-Figure B4). We also standardize all experience measures within the sample so that we can compare magnitudes.

To ensure that the analysis of house-price experiences and its comparison to the estimated effect of general inflation experiences is not affected by the more restricted data, we first replicate our main results on the on the subsample of countries with available house-price data. As shown in columns (1) and (4) of Table 2.7, our estimation results are robust. In the samples for a measure of house price experience is available, a 1 standard-deviation increase in log experienced inflation predicts a 93% or 76% increase in the odds of ownership, or an increase in predicted ownership from 65% to 78% or 77%.

Turning to the explanatory factor of interest, in columns (2) and (5), we add the measures of past real house-price growth. We find that homeownership is significantly predicted by past house-price growth experiences, consistent with Prediction 2. The estimated magnitudes are similar for both measures of experienced house price growth, albeit smaller than that of experienced inflation. We find that a one standard deviation increase in experienced house-price growth predicts an increase in the probability of homeownership from 65% to 70% in the full-history sample, and to 68% in the partial-history sample. In both of these regressions, the effect of experienced inflation remains relatively stable and statistically significant. Moreover, both experienced house price measures become statistically insignificant after controlling for country-wave fixed effects (in columns (3) and (6)), while experienced consumer price inflation continues to be a significant predictor of homeownership.

Why are the results on house-price experiences less robust than those of experienced inflation? After all prior exposure to price changes in the housing market seems most relevant to beliefs about future home prices. And we have ruled out that the more limited house-price data explains the lack of robustness: Even though we limit the variation and statistical power of our analyses by restricting the analysis to a subset of countries or a shorter time horizon, general-inflation experiences continue to strongly predict homeownership. The results are also unaffected when we use nominal rather than real house-price changes. The estimated house-price experience effect continues to be smaller than that of general inflation (slightly larger effects than with real prices), and the results are mixed when including country or country-wave fixed effects.

One possible explanation is that house-price changes may be less apparent to households than general inflation. Households may be more familiar with the prices of goods frequently purchased and pay relatively little attention to changes in house prices, cf. D'Acunto et al. (ming). If anything they may be aware of the price appreciation of their prior home or a parent's home, which might be an interesting avenue for future research.

Another possible explanation for the weaker explanatory power of consumers' past exposure to house-price increases is their direct impact on affordability: fewer people can afford to become homeowners if their wages did not grow by more than the inflation rate. In Appendix-Table B1, we investigate this hypothesis directly using price-to-rent ratios as our proxy for the relative affordability of home purchases relative to renting. Aggregating to the country-wave level, we regress price-to-rent ratios on average experiences. We find no statistically significant relationship between price-to-rent ratios and countries' general inflation experiences (taking the average of households' log experience), but a strong positive relationship with past house-price growth experiences. One standard deviation higher house-price growth predicts about one standard deviation higher price-to-rent ratios. To put it in the context of our model, if we allow rental prices to move with general inflation, the price-to-rent ratio will not respond to inflation, but will respond to house-price growth.

Both explanations for the weaker explanatory power of house-price experiences on housing decisions—the lack of consumer attention to past house-price movements (and lack of data on house-price movements consumers do pay attention to) and the impact on affordability—might be at work and would be interesting to explore if data becomes available. At the same time, inflation histories emerge as a strong influence on tenure decisions, even after taking house-price growth into account.

Robustness

We test the robustness of our main analysis to a number of alternative specifications.

Multiple Imputation Data. In Appendix-Table B2, we test the sensitivity of our main estimates to the use of the multiple-imputation data. In column (1), we report all of the coefficients from our benchmark estimation in Table 2.5. Using only the non-imputed data (column (2)), we limit the analysis to about 60% of the sample when we control for wealth and income. In column (3), we estimate the model on non-imputed data without including wealth and income controls. Across all three specifications, the magnitude of the coefficient on log experienced inflation is similar, with a one log-point change in experienced inflation corresponding to an increase in the likelihood of homeownership from 65% to between 80% and 84%. While the coefficient on experienced inflation remains stable, the wealth and income controls increase the explanatory power of the model alter the effect of some of the other demographic coefficients, including age, education, and unemployment. This may indicate that one mechanism through which age, education, and employment affect ownership is through wealth accumulation. Most importantly, all coefficient estimates and also the increased explanatory power of the estimation with wealth and income controls on the subsample of non-imputed data closely match those from the estimation on multiple-imputation data.

Alternative Measures of Inflation Experience. In Appendix B.4, we discuss in detail several alternative methods of capturing past inflation experiences, with the corresponding estimation results shown in Appendix-Table B3. First, we demonstrate the robustness of our main result to the treatment of households with high inflation experience. We estimate coefficients of similar magnitude under several alternate specifications: estimating a linear effect of experienced inflation in all countries, limiting to the subsample of countries without high inflation, and winsorizing experienced inflation either before or after averaging to calculate the lifetime experience measure. Second we test several conceptually different measures of experienced inflation. We find that experienced inflation volatility also predicts homeownership, but with a smaller magnitude than the level. We also implement and extend the AR(1) model as described in Malmendier and Nagel (2015) to estimate households' one-year and five-year inflation forecasts from their lifetime experienced inflation. Higher estimated forecasts also significantly predict higher likelihoods of homeownership, but with smaller magnitudes than our main specification.

Alternative Wealth Controls. We also probe the robustness of our results to alternative methods for controlling for household wealth, which are discussed in detail in Appendix B.5. We show that the predictive power of log experienced inflation is robust to controlling for measures of household wealth net of home equity or house-price appreciation (see columns (1) and (2) of Appendix-Table B4). The main results are also robust to using nominal, rather than real income and wealth (column 3) and to adjusting real income and wealth for purchasing power parity across countries in the Euro area (column 4). Finally, in column (5), we test the robustness to defining the wealth and income deciles within rather than across countries.

Accounting for Persistence in Homeownership. Our main analysis tests the hypothesis that macroeconomic experiences predict homeownership *at the time of the survey*. One potential concern we discussed in Section 2.3 is persistence in homeownership: While beliefs formed up to the moment of first becoming a homeowner matter for the purchase decision, homeowners might be unlikely to switch back to renting. While the ACS and HFCS data does not identify when an individual first assumed homeownership status, the retrospective data from the Survey of Health, Ageing and Retirement in Europe (SHARE) does. We perform a full analysis of the role of past inflation exposure on households' first homeownership decisions using the SHARE data. The data and analysis are described in detail in Appendix B.6 with summary statistics and results reported in Appendix-Tables B5 and B6. We find that experienced inflation predicts if and when an individual first purchases a home.

Another approach to address persistence in homeownership could be to focus on recent movers as those individuals are forced to re-evaluate their tenure decision after moving. If moving was random, we might expect to estimate a stronger experience effect in the subsample of recent movers. Unfortunately, we are lacking quasi-random variation in moving. In our HFCS sample, for example, the 22% who have moved in the last 5 years are younger, more employed, and significantly more likely to be renters with 32% ownership vs. 71% in the sample that has not moved recently. The benefit of the retrospective SHARE data is that it allows us to address the persistence in homeownership without the selection issues of the cross-sectional data.

Additional Robustness. Our main result is robust to including age fixed effects or cohort (birth) year fixed effects instead of age. In our main analyses, we control for survey-wave fixed effects as most surveys occur over a concentrated period; however, our results are robust to including survey-year fixed effects. Our main result is also robust to using the HFCS replicate weights (bootstrap weights accounting for the sampling design). We also test the robustness of the results to clustering standard errors by country. Clustering increases standard errors on log experienced inflation by a factor of 7, but the effect remain statistically significant at the 1% level. Because we have only 22 countries, we use the score bootstrap approach (Kline et al. (2012)): the average p-value for the coefficient on log experienced inflation (across the 5 imputations) is 0.002. Again, inflation histories emerge as a robust determinant.

2.5 Conclusion

In this paper we present evidence that the macroeconomic histories individuals experience in their home countries have a long-lasting effect on the composition of and demand in the housing market. Heterogeneity in households exposure to past episodes of higher or lower inflation can explain differences in the likelihood of being a homeowner, both within and across countries. Thus, individual-level and country-level histories of past price increases emerge as an economically meaningful factor explaining large cross-country differences in housing markets as well as the variation in ownership within countries. We show that the relationship between prior inflation and tenure choices is not explained by housing market conditions. The effect of personal experiences appears to be powerful and long-lasting enough to influence even the homeownership decisions of immigrants who move to the same housing market (the U.S.) and still respond to the inflation exposure they experienced in their home countries. We also show that inflation experiences throughout life predict the hazard of first homeownership.

The results of this paper tie into the literature on the long-run effects of macroeconomic events such as high inflation and economic crises addressed in DeLong and Summers (2012), Giuliano and Spilimbergo (2009), and Oreopoulos et al. (2012) among others. In this paper we formulate and address a housing-market participation puzzle and show that a similar notion of long-run goes a long way explaining the puzzle.

Table 2.1: Summary of Household Characteristics

Variable	Mean	Median	SD
Panel A: ACS Data (N=165,387)			
Homeowner (Own Property where Surveyed)	0.69	1	0.46
Age	53.2	53	15.1
Male	0.54	1	0.50
Year of immigration	1977.4	1975	18.2
Percent of life in U.S.	65.6	69.9	28.2
Has child	0.37	0	0.48
Single	0.15	0	0.36
Married	0.56	1	0.50
Married to U.S. native	0.29	0	0.45
Educ: high school (U.S. system)	0.52	1	0.50
Educ: college (U.S. system)	0.39	0	0.49
Employed	0.65	1	0.48
Unemployed	0.03	0	0.17
Household income (2010 USD)	84,668	59,624	92,517
Panel B: HFCS Data (N=220,697)			
Homeowner (Own Main Residence)	0.62	1.00	0.48
Age	51.47	51	15.32
Male	0.55	1.00	0.50
Has child	0.45	0.00	0.50
Single	0.24	0.00	0.43
Married	0.54	1.00	0.50
Educ: high school (ISCED-97=3)	0.44	0.00	0.50
Educ: college (ISCED-97=5)	0.27	0.00	0.44
Employed	0.56	1.00	0.50
Unemployed	0.06	0.00	0.25
Household income (2010 Euros)	37,395	27,688	41,846
Household net wealth (2010 Euros)	204,546	87,805	633,343

Notes: Data in Panel A is from the 2008-2018 American Community Surveys (ACS), obtained from IPUMS (Ruggles et al. (2020)). Sample includes household heads who immigrated to the U.S. from one of the countries sampled in the Household Finance and Consumption Survey (HFCS), weighted to be representative of the U.S. population. Panel B provides the summary statistics for the HFCS sample from waves 1, 2, and 3 (fielded in 2008-2018), weighted to be representative of the population within and across countries.

Table 2.2: Summary of Housing Market Measures

Country	HFCS Homeownership	Tenant Protection	Rent Control	Tax Benefits	Buyer Trans. Cost	Price-to-Rent Ratio	PVR (% of Euros)	PVR (std.)
Lithuania	93%						98%	1.6
Slovakia	88%	-2.3			-2.5		59%	0.4
Croatia	85%						79%	1.0
Hungary	84%	-1.7	-0.6		-1.3		57%	0.3
Spain	82%	0.9	-0.8	0.5	0.0	0.7	89%	1.3
Malta	81%						64%	0.6
Poland	79%	-1.0	-1.3		-0.2		91%	1.4
Slovenia	77%	-1.9	-1.7		-2.5		81%	1.1
Estonia	76%				-2.2		88%	1.3
Portugal	75%	1.4	-0.3	-0.4	-0.9	-0.9	91%	1.4
Latvia	75%						94%	1.5
Cyprus	74%						64%	0.6
Greece	72%	1.9	-0.6	1.5	2.4	-0.4	54%	0.3
Belgium	70%	-2.0	-0.8	1.0	2.7	1.6	35%	-0.3
Ireland	68%	-2.0	-1.1	-0.2	-1.3	-0.4	83%	1.2
Italy	68%	-0.1	-0.8	-0.6	0.1	0.1	58%	0.4
Finland	68%	-1.3	-1.7	1.5	-1.5	-0.2		
Luxembourg	67%	-2.4	0.1	-0.5	0.6		73%	0.9
Netherlands	58%	-1.8	1.6	2.9	-1.0	0.9	84%	1.2
France	57%	1.0	-0.3	0.3	1.0	0.9	10%	-1.1
Austria	48%	1.1	0.2	-0.5	-0.7	0.3	64%	0.6
Germany	44%	-0.1	1.3	-0.9	-0.6	-1.2	14%	-1.0

Notes: Table is sorted by the homeownership rate (the percent of households who own their main residence) in the HFCS sample. The summary statistics are weighted to be representative of the population within and across countries. Housing market variables are constructed using data from the HFCS, Andrews et al. (2011), the World Bank, and the OECD. All housing market measures are normalized to have a mean of 0 and variance of 1 in the sample. Tenant protection is a comparative measure of tenant-landlord regulations. Rent control is a composite indicator increasing in the extent of controls of rents. Tax benefits is a comparative measure of the tax relief on debt financing of homeownership. Transaction costs measure the average cost associated with purchasing a home, including transfer taxes, real estate agent fees, notary fees, legal fees, and registration fees. Price-to-rent ratio is an index with a baseline for each country equal to the long-run average price-to-rent ratio within the country, where the long-run is defined as starting in 1980 or the average over all available data if the data begins after 1980. Prevalence of variable-rate mortgages (PVR) is the percent of main residence mortgage euros that carry an adjustable (vs. fixed) interest rate, calculated for each country-wave. Country averages for the table are the average across waves, weighted by the sum of household weights. In the last two columns, we display the averages of both the underlying measure (percent of mortgage euros) and the normalized measure.

Table 2.3: Summary of Experienced and Average Historical Inflation Rates by Country

		Average Past Inflation (%)			Experienced Past Inflation (%) ACS				Experienced Past Inflation (%) HFCS			
Country		Mean	Median	SD	Mean	Median	SD	N	Mean	Median	SD	N
AT	Austria	5.5	2.2	12.9	3.6	3.7	0.8	2,881	2.7	2.7	0.6	7,990
BE	Belgium	3.1	2.6	3.3	3.1	3.3	0.6	2,193	2.8	3.0	0.5	6,332
HR	Croatia	80.6	8.0	310.9	48.7	7.7	55.9	2,081	103.5	111.6	22.6	1,273
CY	Cyprus	3.7	2.8	5.2	3.4	3.6	0.5	261	3.3	3.4	0.7	3,704
EE	Estonia	31.5	2.2	124.2	36.3	35.5	30.3	170	47.9	51.4	14.0	4,596
FI	Finland	29.1	3.4	213.8	15.1	5.2	13.4	1,055	22.9	29.5	14.0	29,445
FR	France	7.5	3.2	11.8	3.4	3.5	0.9	11,248	3.3	3.3	1.1	37,720
DE	Germany	3.7	2.7	6.9	3.3	3.5	0.7	60,081	2.6	2.7	0.7	12,315
GR	Greece	25.4	5.2	85.5	5.4	4.2	2.3	8,552	7.7	7.7	2.6	8,521
HU	Hungary	371.4	4.7	2549.4	34.1	8.2	46.7	3,944	16.3	9.0	21.9	11,367
IE	Ireland	4.7	2.9	5.6	3.9	3.9	0.8	7,488	3.4	3.6	1.1	9,663
IT	Italy	11.3	3.5	38.7	4.2	4.0	1.0	20,776	4.9	5.0	1.3	20,914
LV	Latvia	31.5	1.7	118.6	24.2	4.2	25.1	1,089	47.7	50.8	12.9	2,328
LT	Lithuania	38.1	1.4	153.1	46.7	60.1	33.6	1,491	56.4	60.8	12.4	1,546
LU	Luxembourg	4.4	2.4	7.9					2.7	2.8	0.5	4,054
MT	Malta	2.7	2.1	3.5					2.8	2.8	0.3	2,665
NL	Netherlands	3.5	2.8	4.1	3.3	3.4	0.6	5,459	2.9	2.9	0.4	4,925
PL	Poland	19.8	4.7	68.3	17.0	14.9	12.3	19,952	27.2	29.2	5.8	8,804
PT	Portugal	6.2	3.1	9.2	4.4	3.9	1.4	10,024	6.8	7.1	1.6	15,321
SK	Slovakia	4.9	1.7	11.2	4.4	4.3	1.5	946	5.8	5.6	0.8	6,152
SI	Slovenia	57.0	10.8	207.4					61.6	67.6	19.9	4,615
ES	Spain	6.8	4.8	6.4	3.9	3.9	1.0	5,696	5.6	5.9	1.1	16,891
All					7.5	3.8	14.5	165,387	6.8	3.6	11.0	221,141
All (eq. wt.)					14.1	3.9	26.0	165,387	20.0	5.3	28.1	221,141
Across Countries												
Mean					7.5	6.0	4.8	19	6.8	7.1	2.1	22
Median					3.9	3.9	0.9	19	3.3	3.3	1.1	22
Mean (eq. wt.)					14.1	9.5	12.1	19	20.0	21.3	6.2	22
Median (eq. wt.)					4.4	4.0	1.4	19	5.7	5.7	1.2	22

Notes: The inflation data, ACS data, and HFCS non-imputed survey data are described in Section 2.3. Average Past Inflation is based on annual inflation rates from 1930 to 2018. The summary statistics of Experienced Past Inflation are weighted to be representative of the populations within and across countries unless otherwise noted. We indicate with “eq. wt.” that summary statistics are weighted to be representative within country, but equally weighted across countries. “Across Countries” statistics report the mean or median sample statistics across countries in the top panel.

Table 2.4: Inflation Exposure and Homeownership: Immigrants to the U.S. (ACS)

Sample:	Immigrants			All	
Dep. Var.: Homeowner	(1)	(2)	(3)	(4)	(5)
Log Experienced Inflation	1.22***	1.19***	1.09***	2.11***	2.53***
	(0.02)	(0.01)	(0.02)	(0.03)	(0.03)
Years in U.S. (% of Life)	1.01***	1.01***	1.01***	1.02***	1.02***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
State-Year Homeownership Rate		1.07***	1.07***		
		(0.00)	(0.00)		
U.S. Native				0.65***	
				(0.01)	
Demographic Controls	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Country of Birth FE	No	No	Yes	No	Yes
State & State x Year FE	No	No	No	Yes	Yes
Observations	165,387	165,387	165,387	10,964,255	10,964,255
Pseudo R^2	0.221	0.234	0.237	0.266	0.266

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Table reports exponentiated coefficients (odds ratios) from logit regressions on the 2008-2018 ACS data, using representative weights, with robust standard errors in parentheses. The dependent variable is an indicator for homeownership. Log Experienced Inflation is the log of the weighted average of inflation over the household head's lifetime, with linearly declining weights from the year before the survey to birth year, using inflation from the birth country from birth year to year of immigration to the U.S., and only U.S. inflation for natives. Demographic controls include household head age, gender, marital status, children in the home, education, employment status, and decile of total household income relative to the entire ACS population. All regressions also include survey-year fixed effects. State-year homeownership rate is the homeownership rate among other (non-immigrant) households calculated using the ACS from the same state and year. The sample in columns 1-3 are household heads born in one of the 22 ECB HFCS countries. In columns 4 and 5 we also include (and control for) households headed by a U.S. native.

Table 2.5: Inflation Exposure and Homeownership: Within and Across Countries (HFCS)

Dep. Var: Own Main Residence	(1)	(2)	(3)	(4)
Log Experienced Inflation	2.66***	1.36***	1.12***	1.14***
	(0.06)	(0.08)	(0.04)	(0.05)
Tenant Protection		1.07***		
		(0.02)		
Rent Control		0.62***		
		(0.02)		
Tax Benefits to Homeowners		1.70***		
		(0.04)		
Buyer Trans. Cost		0.83***		
		(0.01)		
Price-to-Rent Ratio		0.72***		
		(0.02)		
Demographic Controls	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	No
Country FE	No	No	Yes	No
Country-Wave FE	No	No	No	Yes
Observations	214,133	163,240	214,133	214,133
Countries	22	11	22	22
Pseudo R^2	0.511	0.530	0.535	0.536

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Table reports exponentiated coefficients (odds ratios) from logit regressions with robust standard errors in parentheses. Data is the HFCS multiple-imputation data, using representative weights. The number of observations is the maximum N across the 5 imputations. Pseudo R^2 is the average across the 5 imputations. The dependent variable is an indicator for owning the household main residence. Log Experienced Inflation is the log of weighted average of inflation over the household head's lifetime, with linearly declining weights from the year before the survey to birth year. Demographic controls include age, gender, marital status, children, education, employment status, and deciles of net wealth and household income. Tenant protection is a comparative measure of tenant-landlord regulations. Rent control is a composite indicator increasing in the extent of controls of rents. Tax benefits is a comparative measure of the tax relief on debt financing of homeownership. Transaction costs measure the average cost associated with purchasing a home. Price-to-rent ratio is a comparison of the cost of ownership and renting relative to the country's long-run average. Housing market variables obtained from Andrews et al. (2011) and OECD and are normalized to have a mean of 0 and variance of 1 across all available data.

Table 2.6: Heterogeneity in Within-Household Inflation Experiences and in Financing

	Married vs. Single		Variable vs. Fixed Rate	
Dep. Var: Own Main Residence	(1)	(2)	(3)	(4)
Log Experienced Inflation	3.26***	1.38***	2.88***	1.48***
	(0.12)	(0.07)	(0.11)	(0.11)
Log Experienced Inflation X Married	0.74***	0.73***		
	(0.03)	(0.03)		
Log Experienced Inflation X PVR			0.69***	0.85***
			(0.02)	(0.04)
Prevalence of Variable Rate (PVR)			2.27***	
			(0.09)	
Demographic Controls	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes
Country-Wave FE	No	Yes	No	Yes
Observations	170,512	170,512	185,097	185,097
Countries	22	22	21	21
Pseudo R^2	0.473	0.500	0.520	0.536

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Table reports exponentiated coefficients (odds ratios) from logit regressions with robust standard errors in parentheses. Data is the HFCS multiple-imputation data, using representative weights. The number of observations is the maximum N across the 5 imputations. Pseudo R^2 is the average across the 5 imputations. The dependent variable is an indicator for owning the household main residence. Log experienced inflation is the log of the weighted average of inflation over the household head's lifetime, with linearly declining weights from the year before the survey to birth year. Demographic controls include age, gender, marital status, children, education level, employment status, and deciles of net wealth and household income. Columns 1 and 2 exclude widowed and divorced household heads. PVR is a standardized measure of the prevalence of variable- relative to fixed-rate mortgages in each country-wave.

Table 2.7: Experiences of Real House-Price Growth and Homeownership

Dependent Var: Own Main Residence	(1)	(2)	(3)	(4)	(5)	(6)
Log Experienced Inflation (std.)	1.93***	1.74***	1.08**	1.76***	1.72***	1.07**
	(0.05)	(0.03)	(0.04)	(0.04)	(0.03)	(0.03)
Experienced Real House Price Growth (Full History, std.)		1.27***	0.87			
		(0.03)	(0.10)			
Experienced Real House Price Growth (Partial History from 1975, std.)					1.12***	1.14
					(0.02)	(0.15)
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	No	Yes	Yes	No
Country-Wave FE	No	No	Yes	No	No	Yes
Observations	100,931	100,931	100,931	136,728	136,728	136,728
Countries	6	6	6	10	10	10
Pseudo R^2	0.501	0.504	0.517	0.526	0.526	0.541

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Table reports exponentiated coefficients (odds ratios) from logit regressions with robust standard errors in parentheses. Data is the HFCS multiple-imputation data, using representative weights. The number of observations is the maximum N across the five imputations. Pseudo R^2 is the average across the 5 imputations. The dependent variable is an indicator for owning the household main residence. Experience measures are calculated as the weighted average over the household head's lifetime, with linearly declining weights from the year before the survey to birth year (re-scaling the weights for partial history). Full history of experienced real house-price growth obtained from Knoll et al. (2017) for Belgium, Finland, France, Germany, and the Netherlands. Full history of experienced real house-price growth obtained from Bordo and Landon-Lane (2014) for Spain. Partial history of experienced real house-price growth obtained from the Federal Reserve Bank of Dallas for Belgium, Croatia, Finland, France, Germany, Ireland, Italy, Luxembourg, the Netherlands, and Spain. All experience measures are standardized within the regression sample. Demographic controls include age, gender, marital status, children, education, employment status, log net wealth, and log income. All regressions also include fixed effects for survey wave or country-survey-wave.

Conclusion

This dissertation examines how individual choices are shaped by the past. In Chapter 1, I describe the role of habit formation in hand hygiene behavior of hospital workers. In Chapter 2, I explore the role of macroeconomic experiences in predicting homeownership choices.

Though in vastly different settings and through different mechanisms, both chapters highlight the importance of the past in predicting future choices. These findings have important implications for individuals, organizations, and the macroeconomy.

Bibliography

Acland, D. and Levy, M. R. (2015). Naiveté, projection bias, and habit formation in gym attendance. *Management Science*, 61(1):146–160.

Aiken, L. H., Clarke, S. P., Sloane, D. M., Sochalski, J., and Silber, J. H. (2002). Hospital nurse staffing and patient mortality, nurse burnout, and job dissatisfaction. *Jama*, 288(16):1987–1993.

Alesina, A. and Fuchs-Schündeln, N. (2007). Good-bye lenin (or not?): The effect of communism on people's preferences. *The American Economic Review*, 97(4):1507–1528.

Allcott, H. and Rogers, T. (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104(10):3003–37.

Ampudia, M. and Ehrmann, M. (2014). Macroeconomic experiences and risk taking of euro area households. *Bank of Canada Working Paper (No. 2014-10).*

Anari, A. and Kolari, J. (2002). House prices and inflation. *Real Estate Economics*, 30(1):67–84.

Andersen, H. S. (2011). Motives for tenure choice during the life cycle: the importance of non-economic factors and other housing preferences. *Housing, Theory and Society*, 28(2):183–207.

Andrews, D. and Caldera Sánchez, A. (2011). Drivers of homeownership rates in selected oecd countries. *OECD Publishing (No. 849).*

Andrews, D., Caldera Sánchez, A., and Johansson, Å. (2011). Housing markets and structural policies in oecd countries. *OECD Publishing (No. 836).*

Apostolides, A. (2011). The growth of two small economies in the great depression: Gdp estimation for cyprus and malta during the interwar period (1921-1938). *MPRA Paper No. 30276.*

Becker, G. S., Grossman, M., and Murphy, K. M. (1994). An empirical analysis of cigarette addiction. *The American Economic Review*, 84(3):396–418.

Becker, G. S. and Murphy, K. M. (1988). A theory of rational addiction. *Journal of political Economy*, 96(4):675–700.

Becker, M. C. (2004). Organizational routines: a review of the literature. *Industrial and corporate change*, 13(4):643–678.

Bell, C. M. and Redelmeier, D. A. (2001). Mortality among patients admitted to hospitals on weekends as compared with weekdays. *New England Journal of Medicine*, 345(9):663–668.

Bernheim, B. D. and Rangel, A. (2004). Addiction and cue-triggered decision processes. *American economic review*, 94(5):1558–1590.

Beshears, J., Lee, H. N., Milkman, K. L., Mislavsky, R., and Wisdom, J. (2020). Creating exercise habits using incentives: The trade-off between flexibility and routinization. *Management Science*.

Bordo, M. D. and Landon-Lane, J. (2014). What explains house price booms?: History and empirical evidence. *Macroeconomic Analysis and International Finance*.

Börsch-Supan, A., Brandt, M., Hunkler, C., Kneip, T., Korbmacher, J., Malter, F., Schaan, B., Stuck, S., and Zuber, S. (2013). Data resource profile: the survey of health, ageing and retirement in europe (share). *International journal of epidemiology*, 42(4):992–1001.

Botsch, M. J. and Malmendier, U. (2019). Inflation experiences and contract choice – evidence from residential mortgages. Working paper, UC-Berkeley.

Bracha, A. and Jamison, J. C. (2012). Shifting confidence in homeownership: The great recession. *The BE Journal of Macroeconomics*, 12(3).

Brounen, D., Eichholtz, P., Staetmans, S., and Theebe, M. (2014). Inflation protection from homeownership: Long-run evidence, 1814–2008. *Real Estate Economics*, 42(3):662–689.

Carrera, M., Royer, H., Stehr, M., and Sydnor, J. (2018). Can financial incentives help people trying to establish new habits? experimental evidence with new gym members. *Journal of health economics*, 58:202–214.

Case, B. and Wachter, S. M. (2011). Inflation and real estate investments. *U of Penn, Inst for Law & Econ Research Paper 11-33*.

Charness, G. and Gneezy, U. (2009). Incentives to exercise. *Econometrica*, 77(3):909–931.

Chiuri, M. C. and Jappelli, T. (2003). Financial market imperfections and home ownership: a comparative study. *European economic review*, 47(5):857–875.

Clark, W. A. and Dieleman, F. M. (1996). Households and housing. *Choice and outcomes in the housing market*.

Collin-Dufresne, P., Johannes, M., and Lochstoer, L. A. (2017). Asset pricing when this time is different. *The Review of Financial Studies*, 30(2):505–535.

Collins, J. M. and Choi, L. (2010). The effects of the real estate bust on renter perceptions of homeownership. *Available at SSRN 1569009*.

Correia, S. (2016). Linear models with high-dimensional fixed effects: An efficient and feasible estimator. Technical report. Working Paper.

D'Acunto, F., Malmendier, U., Ospina, J., and Weber, M. (forthcoming). Exposure to grocery prices and inflation expectations. *Journal of Political Economy*.

Dai, H., Milkman, K. L., Hofmann, D. A., and Staats, B. R. (2015). The impact of time at work and time off from work on rule compliance: The case of hand hygiene in health care. *Journal of Applied Psychology*, 100(3):846.

Dean, J. T. (2020). Noise, cognitive function, and worker productivity. Working paper.

DellaVigna, S. and Pollet, J. M. (2009). Investor inattention and friday earnings announcements. *The Journal of Finance*, 64(2):709–749.

DeLong, J. B. and Summers, L. (2012). Fiscal policy in a depressed economy. *Brookings Papers on Economic Activity*, 44(1 (Spring)):233–297.

Dobkin, C. (2003). Hospital staffing and inpatient mortality. *Unpublished Working Paper*.

Doling, J. (1973). A two-stage model of tenure choice in the housing market. *Urban Studies*, 10(2):199–211.

Drew, R. B. and Herbert, C. E. (2013). Postrecession drivers of preferences for homeownership. *Housing Policy Debate*, 23(4):666–687.

Earley, F. (2004). What explains the differences in homeownership rates in europe? *Housing Finance International*, 19(1):25.

Ehrmann, M. and Tzamourani, P. (2012). Memories of high inflation. *European Journal of Political Economy*, 28(2):174–191.

Eurosystem Household Finance and Consumption Network (2013). The eurosystem household finance and consumption survey-results from the first wave. Technical report, European Central Bank.

Fama, E. F. and Schwert, G. W. (1977). Asset returns and inflation. *Journal of financial economics*, 5(2):115–146.

Fisher, L. M. and Jaffe, A. J. (2003). Determinants of international home ownership rates. *Housing finance international*, 18(1):34–37.

Follain, J. R. and Ling, D. C. (1988). Another look at tenure choice, inflation, and taxes. *Real Estate Economics*, 16(3):207–229.

Giuliano, P. and Spilimbergo, A. (2009). Growing up in a recession: Beliefs and the macroeconomy. *NBER Working Paper 15321*.

Glaeser, E. L. and Nathanson, C. G. (2015). An extrapolative model of house price dynamics. *National Bureau of Economic Research No. w21037*.

Gordon, M. J. (1962). *The investment, financing, and valuation of the corporation*. RD Irwin.

Gould, D. J., Moralejo, D., Drey, N., Chudleigh, J. H., and Taljaard, M. (2017). Interventions to improve hand hygiene compliance in patient care. *Cochrane database of systematic reviews*, (9).

Greenlees, J. S. and McClelland, R. B. (2008). Addressing misconceptions about the consumer price index. *Monthly Lab. Rev.*, 131:3.

Gruber, J. and Köszegi, B. (2001). Is addiction rational? theory and evidence. *The Quarterly Journal of Economics*, 116(4):1261–1303.

Guwande, A. (2011). The checklist manifesto: How to get things right. *New York: Picadur*.

Gwin, C. R. and Ong, S.-E. (2008). Do we really understand homeownership rates? an international study. *International Journal of Housing Markets and Analysis*, 1(1):52–67.

Harris, M. C. and Kessler, L. M. (2019). Habit formation and activity persistence: Evidence from gym equipment. *Journal of Economic Behavior & Organization*, 166:688–708.

Haurin, D. R., Hendershott, P. H., and Wachter, S. M. (1996). Borrowing constraints and the tenure choice of young households. Technical report, National Bureau of Economic Research.

Haynes, A. B., Weiser, T. G., Berry, W. R., Lipsitz, S. R., Breizat, A.-H. S., Dellinger, E. P., Herbosa, T., Joseph, S., Kibatala, P. L., Lapitan, M. C. M., et al. (2009). A surgical safety checklist to reduce morbidity and mortality in a global population. *New England Journal of Medicine*, 360(5):491–499.

Henderson, J. V. and Ioannides, Y. M. (1987). Owner occupancy: investment vs consumption demand. *Journal of Urban Economics*, 21(2):228–241.

Hilber, C. (2007). The determinants of homeownership across europe: Panel data evidence. In *54th Annual North American Meetings of the Regional Science Association International Savannah*, volume 9.

Hussam, R. N., Rabbani, A., Reggiani, G., and Rigol, N. (2017). Habit formation and rational addiction: A field experiment in handwashing. *Harvard Business School working paper series# 18-030*.

International Monetary Fund (2014). World economic outlook database.

Ioannides, Y. M. (1987). Residential mobility and housing tenure choice. *Regional Science and Urban Economics*, 17(2):265–287.

Jarvis, W. (1994). Handwashing - the Semmelweis lesson forgotten? *The Lancet*, 344(8933):1311–1312.

King, D., Vlaev, I., Everett-Thomas, R., Fitzpatrick, M., Darzi, A., and Birnbach, D. J. (2016). Priming hand hygiene compliance in clinical environments. *Health Psychology*, 35(1):96.

Kingston, L., O'Connell, N., and Dunne, C. (2016). Hand hygiene-related clinical trials reported since 2010: a systematic review. *Journal of Hospital Infection*, 92(4):309–320.

Klevens, R. M., Edwards, J. R., Richards Jr, C. L., Horan, T. C., Gaynes, R. P., Pollock, D. A., and Cardo, D. M. (2007). Estimating health care-associated infections and deaths in US hospitals, 2002. *Public health reports*, 122(2):160–166.

Kline, P., Santos, A., et al. (2012). A score based approach to wild bootstrap inference. *Journal of Econometric Methods*, 1(1):23–41.

Knoll, K., Schularick, M., and Steger, T. (2017). No price like home: Global house prices, 1870-2012. *American Economic Review*, 107(2):331–53.

Laibson, D. (2001). A cue-theory of consumption. *The Quarterly Journal of Economics*, 116(1):81–119.

Liu, C. H., Hartzell, D. J., and Hoesli, M. E. (1997). International evidence on real estate securities as an inflation hedge. *Real Estate Economics*, 25(2):193–221.

Loewenstein, G., Price, J., and Volpp, K. (2016). Habit formation in children: Evidence from incentives for healthy eating. *Journal of health economics*, 45:47–54.

Luttmer, E. F. and Singhal, M. (2011). Culture, context, and the taste for redistribution. *American Economic Journal: Economic Policy*, 3(1):157–179.

Macfarlane, A. (1978). Variations in number of births and perinatal mortality by day of week in england and wales. *Br Med J*, 2(6153):1670–1673.

Mack, A., Martínez-García, E., et al. (2011). A cross-country quarterly database of real house prices: a methodological note. *Globalization and Monetary Policy Institute Working Paper*, 99.

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

Malmendier, U. and Nagel, S. (2015). Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1):53–87.

Malmendier, U., Nagel, S., and Yan, Z. (2016a). The making of hawks and doves: Inflation experiences and voting on the fomc. *Working Paper*.

Malmendier, U., Pouzo, D., and Vanasco, V. (2016b). A theory of experience effects. *Working Paper*.

Malmendier, U. and Shen, L. S. (2017). Scarred consumption. Working paper, UC-Berkeley.

Michal, J. M. (1960). *Central Planning in Czechoslovakia: Organization for growth in a mature economy*. Stanford University Press.

Morrison, W. E., Haas, E. C., Shaffner, D. H., Garrett, E. S., and Fackler, J. C. (2003). Noise, stress, and annoyance in a pediatric intensive care unit. *Critical care medicine*, 31(1):113–119.

Naikoba, S. and Hayward, A. (2001). The effectiveness of interventions aimed at increasing handwashing in healthcare workers-a systematic review. *Journal of Hospital infection*, 47(3):173–180.

Neal, D. T., Wood, W., and Quinn, J. M. (2006). Habitsa repeat performance. *Current directions in psychological science*, 15(4):198–202.

Office of Disease Prevention and Health Promotion (ODPHP) (2019). Health care-associated infections. `https://health.gov/hcq/prevent-hai.asp`. [Online; accessed 10-24-2019].

Oreopoulos, P., von Wachter, T., and Heisz, A. (2012). The short-and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics*, 4(1):1–29.

Painter, G., Gabriel, S., and Myers, D. (2001). Race, immigrant status, and housing tenure choice. *Journal of Urban Economics*, 49(1):150–167.

Pittet, D., Hugonnet, S., Harbarth, S., Mourouga, P., Sauvan, V., Touveneau, S., Perneger, T. V., et al. (2000). Effectiveness of a hospital-wide programme to improve compliance with hand hygiene. *The Lancet*, 356(9238):1307–1312.

Price, A. D. and Lu, J. (2013). Impact of hospital space standardization on patient health and safety. *Architectural Engineering and Design Management*, 9(1):49–61.

Purcell, S. R., Kutash, M., and Cobb, S. (2011). The relationship between nurses' stress and nurse staffing factors in a hospital setting. *Journal of Nursing Management*, 19(6):714–720.

Reinhart, C. M. and Rogoff, K. (2009). *This time is different: eight centuries of financial folly*. princeton university press.

Royer, H., Stehr, M., and Sydnor, J. (2015). Incentives, commitments, and habit formation in exercise: evidence from a field experiment with workers at a fortune-500 company. *American Economic Journal: Applied Economics*, 7(3):51–84.

Rubin, D. B. (2004). *Multiple imputation for nonresponse in surveys*, volume 81. John Wiley & Sons.

Ruggles, S., Flood, S., Goeken, R., Grover, J., Meyer, E., Pacas, J., and Sobek, M. (2020). Ipums usa version 10.0 [dataset]. *Minneapolis, MN*.

Sandler, D. H. and Sandler, R. (2014). Multiple event studies in public finance and labor economics: A simulation study with applications. *Journal of Economic and Social Measurement*, 39(1-2):31–57.

Schraeder, S. (2015). Information processing and non-bayesian learning in financial markets. *Review of Finance*, pages 1–31.

Scott, R. D. (2009). The direct medical costs of healthcare-associated infections in us hospitals and the benefits of prevention. *Centers for Disease Control and Prevention*.

Sinai, T. and Souleles, N. S. (2005). Owner-occupied housing as a hedge against rent risk. *The Quarterly Journal of Economics*, 120(2):763–789.

Son, C., Chuck, T., Childers, T., Usiak, S., Dowling, M., Andiel, C., Backer, R., Eagan, J., and Sepkowitz, K. (2011). Practically speaking: rethinking hand hygiene improvement programs in health care settings. *American Journal of Infection Control*, 39(9):716–724.

Staats, B. R., Dai, H., Hofmann, D., and Milkman, K. L. (2017). Motivating process compliance through individual electronic monitoring: An empirical examination of hand hygiene in healthcare. *Management Science*, 63(5):1563–1585.

Suárez-Barraza, M. F., Bou, E., and Cataldo, C. (2008). Finding standards, routines and non-routines in toyota production system (tps): standardization without standardization? *Lean Manufacturing Journal-Reliable Plant Magazine*, 5:1–32.

Taubinsky, D. (2014). From intentions to actions: A model and experimental evidence of inattentive choice. Working paper.

Ulrich, R., Zimring, C., Quan, X., Joseph, A., and Choudhary, R. (2004). The role of the physical environment in the hospital of the 21st century: A once-in-a-lifetime opportunity. *Concord, CA: The Center for Health Design*, 311.

Wood, W., Labrecque, J. S., Lin, P.-Y., and Rünger, D. (2014). Habits in dual process models. *Dual process theories of the social mind*, pages 371–385.

World Health Organization (2009). *WHO guidelines on hand hygiene in health care: first global patient safety challenge clean care is safer care.*

Appendix A

Appendix to Chapter 1

A.1 Derivations

General Cost Function

In this section, I derive results for a more general version of the theoretical framework described in Section 1.2.

Rather than the functional form assumed in Section 1.2, the cost of attention is $c(a, d_t)$, where d_t is the cognitive load at time t. I assume $\partial c(a, d_t)/\partial a > 0$, $\partial^2 c(a, d_t)/\partial a^2 > 0$, $\partial c(a, d_t)/\partial d_t > 0$, and $\partial^2 c(a, d_t)/\partial a \partial d > 0$. Thus, costs are convex in attention and distractions increase the marginal cost of attention. All other aspects of the model are as described in Section 1.2.

The first order condition is given by

$$\frac{\partial c(a^*, d_t)}{\partial a} = b(1 - h_{lt}) \tag{A.1}$$

which yields interior solution

$$p(a^*, h_{lt}) = a^*(1 - h_{lt}) + h_{lt}. \tag{A.2}$$

Following directly from the first order condition A.1 and the conditions on $c(\cdot)$, optimal attention is increasing in b, and decreasing in d_t and h_{lt}.

$$\frac{\partial a^*}{\partial b} \geq 0 \qquad \frac{\partial a^*}{\partial d_t} \leq 0 \qquad \frac{\partial a^*}{\partial h_{lt}} \leq 0 \tag{A.3}$$

Compliance is increasing in b as

$$\frac{\partial p(a^*, h_{lt})}{\partial b} = (1 - h_{lt})\frac{\partial a^*}{\partial b} \geq 0. \tag{A.4}$$

Compliance is decreasing in d_t as

$$\frac{\partial p(a^*, h_{lt})}{\partial d_t} = (1 - h_{lt})\frac{\partial a^*}{\partial d_t} \leq 0. \tag{A.5}$$

The derivative of optimal compliance with respect to habit is ambiguous for a general cost function with

$$\frac{\partial p(a^*, h_{lt})}{\partial h_{lt}} = 1 - a^* + (1 - h_{lt})\frac{\partial a^*}{\partial h_{lt}}. \tag{A.6}$$

Compliance will be increasing in habit if

$$(1 - h_{lt})\frac{\partial a^*}{\partial h_{lt}} > a^* - 1 \tag{A.7}$$

Predictions 4 and 5 rely on the change in the response to shocks to cognitive load and habit across different habit levels. With the general cost function, these are given by the following equations.

$$\frac{\partial^2 p(a^*, h_{lt})}{\partial d_t \partial h_{lt}} = (1 - h_{lt})\frac{\partial^2 a^*}{\partial d \partial h_{lt}} - \frac{\partial a^*}{\partial d_t} \tag{A.8}$$

$$\frac{\partial^2 p(a^*, h_{lt})}{\partial h_{lt}^2} = (1 - h_{lt})\frac{\partial^2 a^*}{\partial h_{lt}^2} - 2\frac{\partial a^*}{\partial h_{lt}} \tag{A.9}$$

The predictions described in Section 1.2 hold under the general cost function whenever the following conditions are satisfied.

$$(1 - h_{lt})\frac{\partial a^*}{\partial h_{lt}} > a^* - 1 \tag{A.10}$$

$$(1 - h_{lt})\frac{\partial^2 a^*}{\partial d \partial h_{lt}} > \frac{\partial a^*}{\partial d_t} \tag{A.11}$$

$$(1 - h_{lt})\frac{\partial^2 a^*}{\partial h_{lt}^2} > 2\frac{\partial a^*}{\partial h_{lt}} \tag{A.12}$$

Specific Cost Function

Consider the cost of attention

$$c(a, d_t) = \frac{d_t a}{(1 - a)}$$

as described in Section 1.2.

The marginal cost of attention under cognitive load d_t is

$$\frac{\partial c(a, d_t)}{\partial a} = \frac{d_t(1 - a) + d_t a}{(1 - a)^2} = \frac{d_t}{(1 - a)^2}.$$

The first order condition is given by

$$\frac{\partial c(a^*, d_t)}{\partial a} = b(1 - h_{lt}) \tag{A.13}$$

$$a^* = 1 - \sqrt{\frac{d_t}{b(1 - h_{lt})}}. \tag{A.14}$$

which yields the interior solution

$$p(a^*, h_{lt}) = a^*(1 - h_{lt}) + h_{lt} \tag{A.15}$$

$$= \left(1 - \sqrt{\frac{d_t}{b(1 - h_{lt})}}\right)(1 - h_{lt}) + h_{lt} \tag{A.16}$$

$$= 1 - \sqrt{\frac{d_t(1 - h_{lt})}{b}}. \tag{A.17}$$

Following from the general framework, attention is increasing in b and decreasing in d_t and h_{lt} with

$$\frac{\partial a^*}{\partial b} = \frac{1}{2}\sqrt{\frac{d_t}{b^3(1 - h_{lt})}} \geq 0, \tag{A.18}$$

$$\frac{\partial a^*}{\partial d_t} = -\frac{1}{2}\sqrt{\frac{1}{d_t b(1 - h_{lt})}} < 0, \tag{A.19}$$

$$\frac{\partial a^*}{\partial h_{lt}} = -\frac{1}{2}\sqrt{\frac{d_t}{b(1 - h_{lt})^3}} \leq 0. \tag{A.20}$$

Compliance is increasing in b, decreasing in d_t, and increasing in h_{lt} as

$$\frac{\partial p(a^*, h_{lt})}{\partial b} = \frac{1}{2}\sqrt{\frac{d_t(1 - h_{lt})}{b^3}} \geq 0, \tag{A.21}$$

$$\frac{\partial p(a^*, h_{lt})}{\partial d_t} = -\frac{1}{2}\sqrt{\frac{(1 - h_{lt})}{d_t b}} \leq 0, \tag{A.22}$$

$$\frac{\partial p(a^*, h_{lt})}{\partial h_{lt}} = \frac{1}{2}\sqrt{\frac{d_t}{b(1 - h)}} \geq 0. \tag{A.23}$$

Predictions 4 and 5 follow from the following derivatives.

$$\frac{\partial^2 p(a^*, h_{lt})}{\partial h_{lt}^2} = \frac{1}{4}\sqrt{\frac{d_t}{b(1 - h)^3}} \geq 0 \tag{A.24}$$

$$\frac{\partial^2 p(a^*, h_{lt})}{\partial d \partial h_{lt}} = \frac{1}{4}\sqrt{\frac{1}{d_t b(1 - h)}} \geq 0 \tag{A.25}$$

As with changes in distraction, changes in benefits will be most impactful for low compliers as

$$\frac{\partial^2 p(a^*, h_{lt})}{\partial b \partial h_{lt}} = -\frac{1}{4}\sqrt{\frac{d_t}{b^3(1-h)}} \leq 0. \tag{A.26}$$

A.2 Sample Selection

Table A1: Sample Selection

	All Data	Main Sample	Compliance Subsamples: Low	Medium	High
Total					
Opportunities	166,557,886	123,465,032	32,546,968	32,740,042	32,349,146
Shifts	3,882,105	2,123,320	616,376	565,042	509,050
Workers	40,356	13,606	4,495	4,495	4,495
Networks	25	21	20	21	21
Average					
Compliance	59%	61%	45%	63%	79%
Shift length (hr)	7.4	9.7	9.5	9.8	10.0
Opportunities per hour	11.1	5.7	5.4	5.7	6.1
Tenure (mo)	11.1	17.5	19.0	17.5	15.9

Notes: Table shows the total number of opportunities, shifts, workers, and networks. Compliance is the average across opportunities, shift length is the average across shifts, opportunities per hour is the average number of opportunities per hour worked, and tenure is the average across workers. All data includes all opportunities in any shift (through February 2020). Main sample limits the data to opportunities occurring in a shift that lasts 5 to 13 hours and workers who average at least one opportunity per hour worked, at least one shift per week, and are in the data for at least six months. Workers from the main sample are assigned to the low, medium, or high compliance group based on their compliance during their first three months in the data. Data summarized for the compliance subsamples excludes the first three months used to categorize.

A.3 Hand Hygiene During the COVID-19 Pandemic

Compliance during the COVID-19 pandemic has followed an interesting trend. In Appendix Figure A1, I plot compliance and the number of opportunities per hour (residualizing both for individual levels in January and February).

Compliance in the first weeks of the pandemic increased by about 4pp. Surprisingly, this increase was very short lived. Peaking in mid-March, compliance quickly declined to levels below those in the beginning of the year. The drop in compliance may be due to mechanical reasons, such as personal protective equipment taking more than 30 seconds to put on and remove and thus hand hygiene occurring outside of the allotted time window.

However, the drop in compliance is also consistent with decaying habits as patient interactions became less frequent. The steep decline in compliance in mid-March follows shortly after a large drop in the number of hand hygiene opportunities, when hospitals began to reduce non-essential activity. The effects are large, with the average worker seeing a 1SD drop in opportunities per hour. As activity rebounded closer to pre-pandemic levels, compliance levels also rose with a slight lag.

Appendix Figure A2 provides further evidence of the link between opportunities and compliance; the drops in compliance were largest for departments who saw the most substantial drops in activity. These may also reflect differences in COVID protocols if, for example, large declines in number of opportunities reflect reserving beds for potential influx of COVID patients which correspond to more substantial changes in patient protocols.

Without more information, it is difficult to conclude too much from the response to COVID.

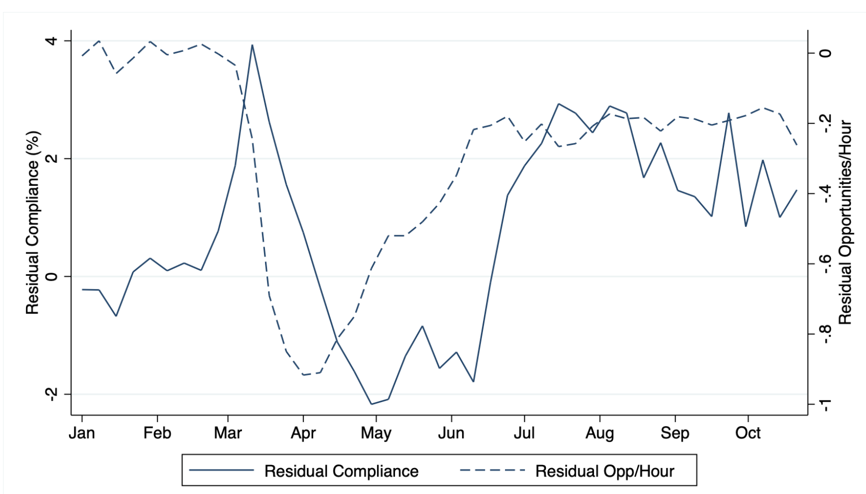

Figure A1: Compliance and number of opportunities in 2020

Graph plots residualized compliance and number of opportunities per hour (residualizing individual levels in Jan and Feb 2020).

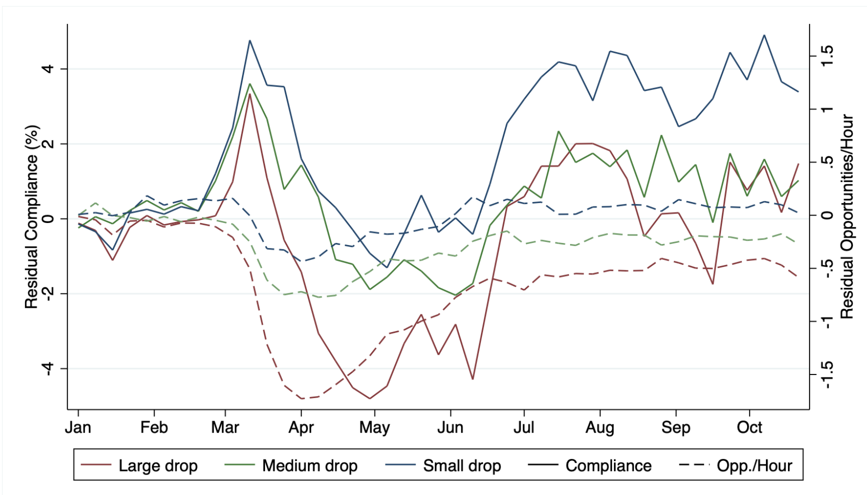

Figure A2: Compliance and number of opportunities in 2020, by change in opportunities

Graph plots residualized compliance and number of opportunities per department-hour (residualizing individual levels in Jan and Feb 2020). Departments are grouped within hospital by their reduction in average residual opportunities per hour from Jan/Feb to March/April/May.

A.4 Additional Event Study Plots for Disruptions to Habit

In Appendix Figure A3, I plot event time dummies for department switches that last 3 and 4+ days in addition to those shown in the main text.

In Appendix Figures A4 and A5, I plot the main event study plots estimated separately for low, medium, and high compliers (classified as described in Section 1.4). For float days, the responses across all types are very similar. For time off work, plots highlight the differential response to short and long breaks across high and low types. For low compliers, there is a benefit to time off that caps out after a few days off. For high compliers, there is a smaller benefit to having a short break, but a large drop in compliance if breaks exceed 2 weeks.

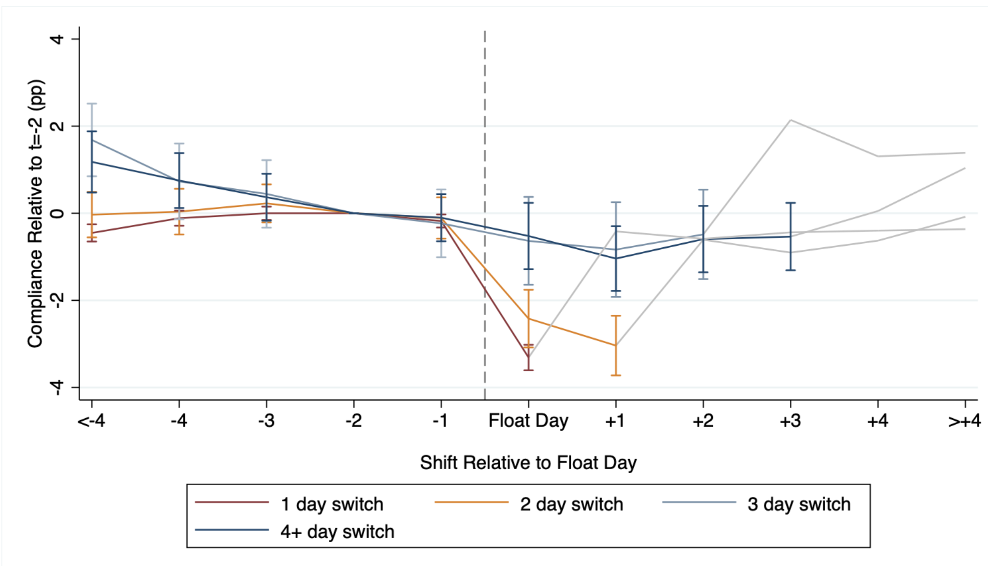

Figure A3: Event study graphs: float days (department switches)

(a) Low Average Compliance

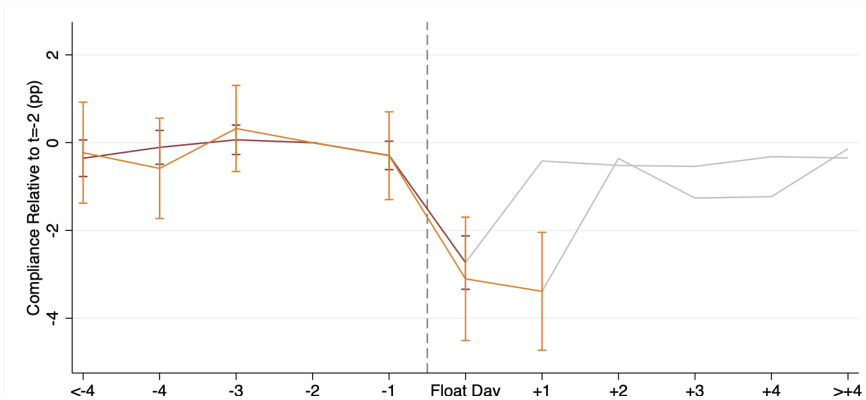

(b) Medium Average Compliance

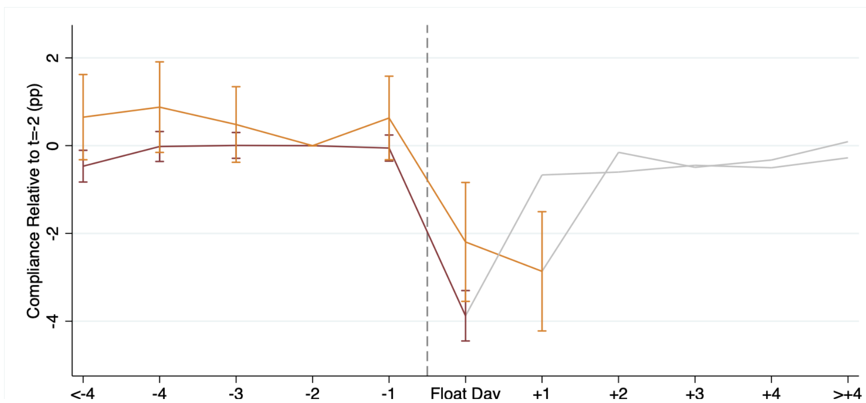

(c) High Average Compliance

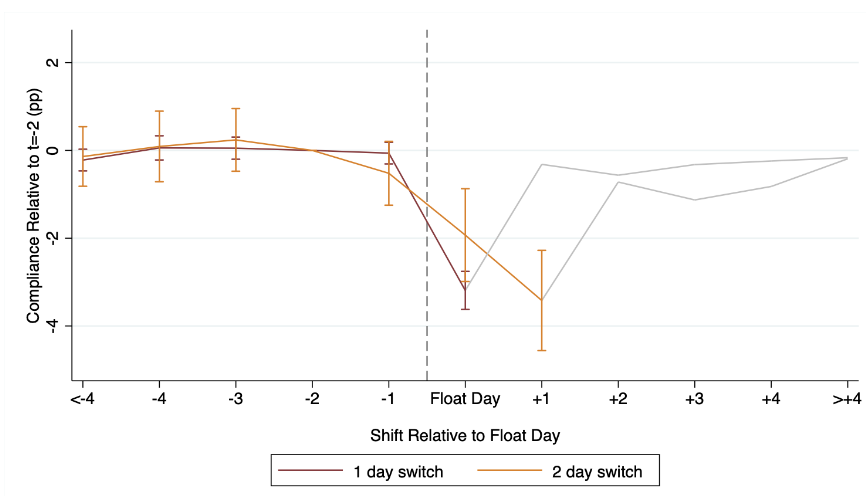

Figure A4: Event study graphs by compliance level: float days (department switches)

(a) Low Average Compliance

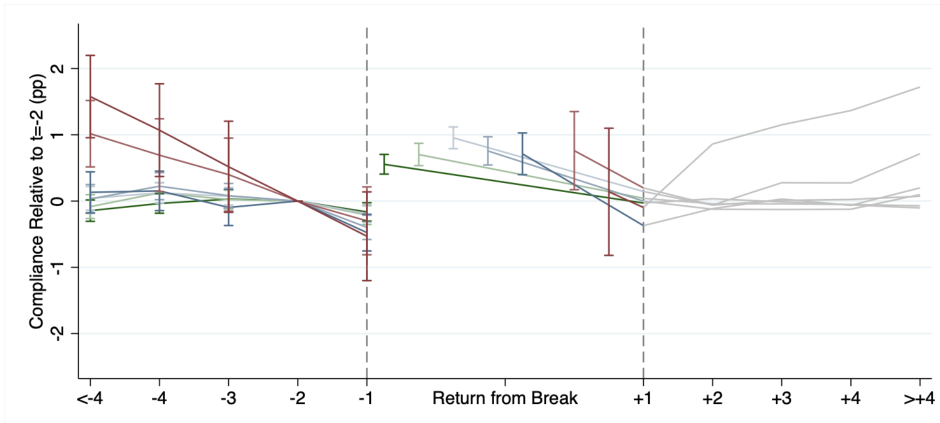

(b) Medium Average Compliance

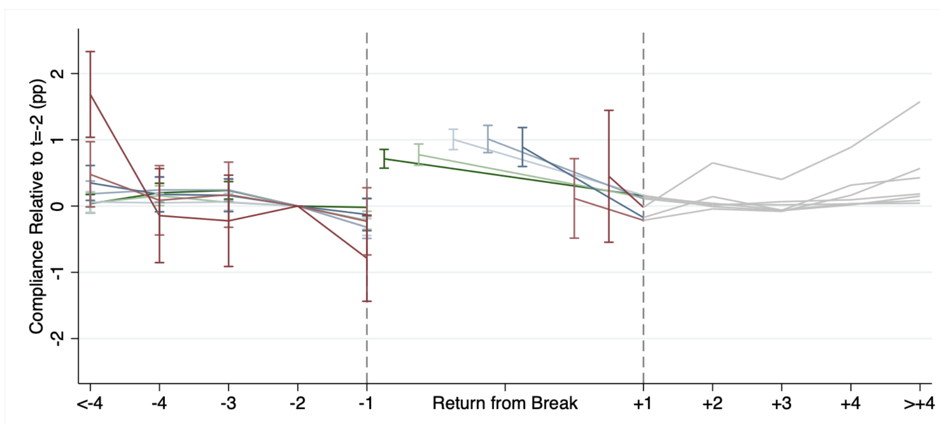

(c) High Average Compliance

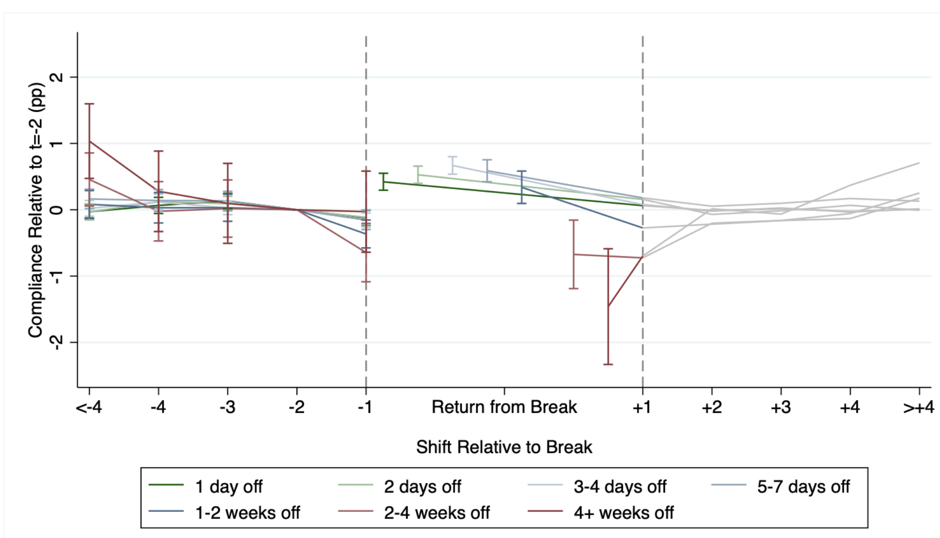

Figure A5: Event study graphs by compliance level: days off between shifts

A.5 Alternative Categorization of Compliance Levels

The key piece of evidence in support of the cue theory of habits is the differential response to cognitive load and habit disruption between low and high compliers. However, the choice of how to identify low vs. high compliers is somewhat arbitrary. In this section, I show that the results in Table 1.3 are robust to alternative methods for categorizing low vs. high compliers.

In the main analyses, I compare based on average compliance in the first three months in the data.[1] This treats compliance across facilities and over time as equal from a habit perspective (e.g., everyone with 80% compliance has the same level of habit). Given hand hygiene is the same action with similar consequences across places and over time, this is a reasonable benchmark. However, we can also consider controlling for differences either in the setting or in the type of worker in the definition of high vs. low compliers.

In Appendix Figure A6, I plot variants of Figure 1.10 under alternative categorizations. In (a) and (b), I show the results under two different ways of accounting for the variation across settings and individuals. In (a), rather than binning into thirds across all individuals, I bin within each network and role-type (e.g., top, middle, and bottom third of registered nurses within the same network). This ensures the bins are balanced across networks and roles. In (b), I take a more direct approach to controlling for differences. I sort on an individual fixed effects estimated from a regression of compliance in the first three months on fixed effects for the individual, network, year, and month. This accounts for differences in compliance levels across networks, aggregate trends over the years, and seasonality (e.g., flu season). Neither of these alternative categorizations substantially changes the results.

In each of these approaches, I use compliance in the first three months to categorize, and compliance in the remainder of the data to estimate the effects. However, if habits build over time, the first three months may be less informative for whether habits have formed than the last three months. In (c), I bin by compliance in the last three months of the pre-COVID data (December 2019 to February 2020) and estimate effects on the data prior to December 2019. I limit the data to the subset of workers who have at least 4 shifts in both the binning and estimation period. Again, the results are very similar to the baseline categorization. If anything, the differential responses to cognitive load and shocks to habit are more pronounced under this alternative categorization.

[1]The results are robust to using 2 or 4 months to categorize instead of 3.

(a) Relative to Network-Role

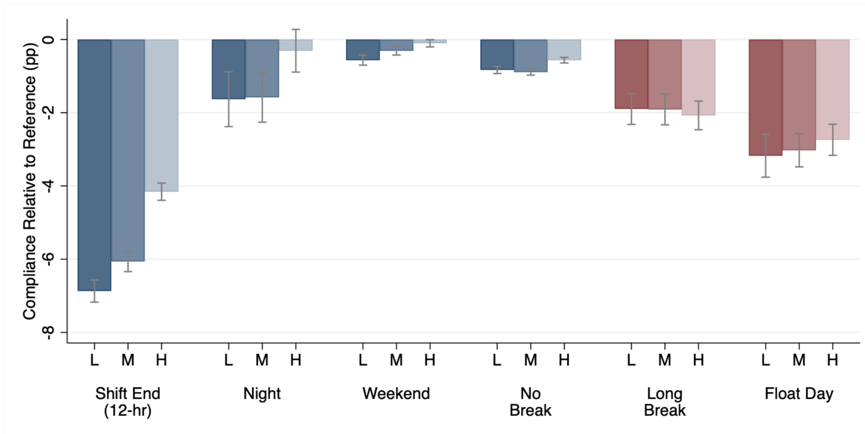

(b) Regression-Adjusted Start Compliance

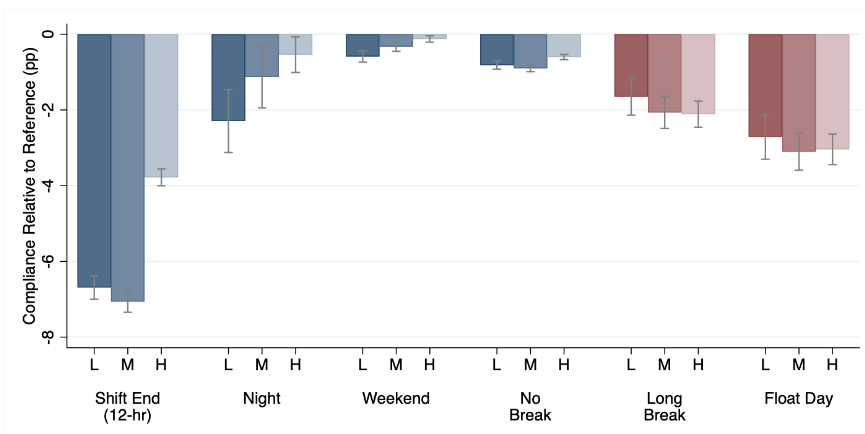

(c) Last 3 Months in the Data

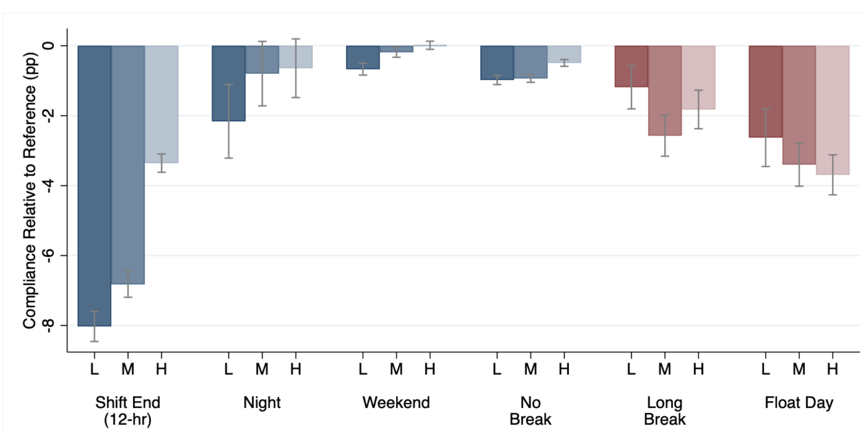

Figure A6: Effect of fatigue, distractions, and habit disruptions on compliance: alternative categorizations

Appendix B

Appendix to Chapter 2

B.1 Additional Figures and Tables

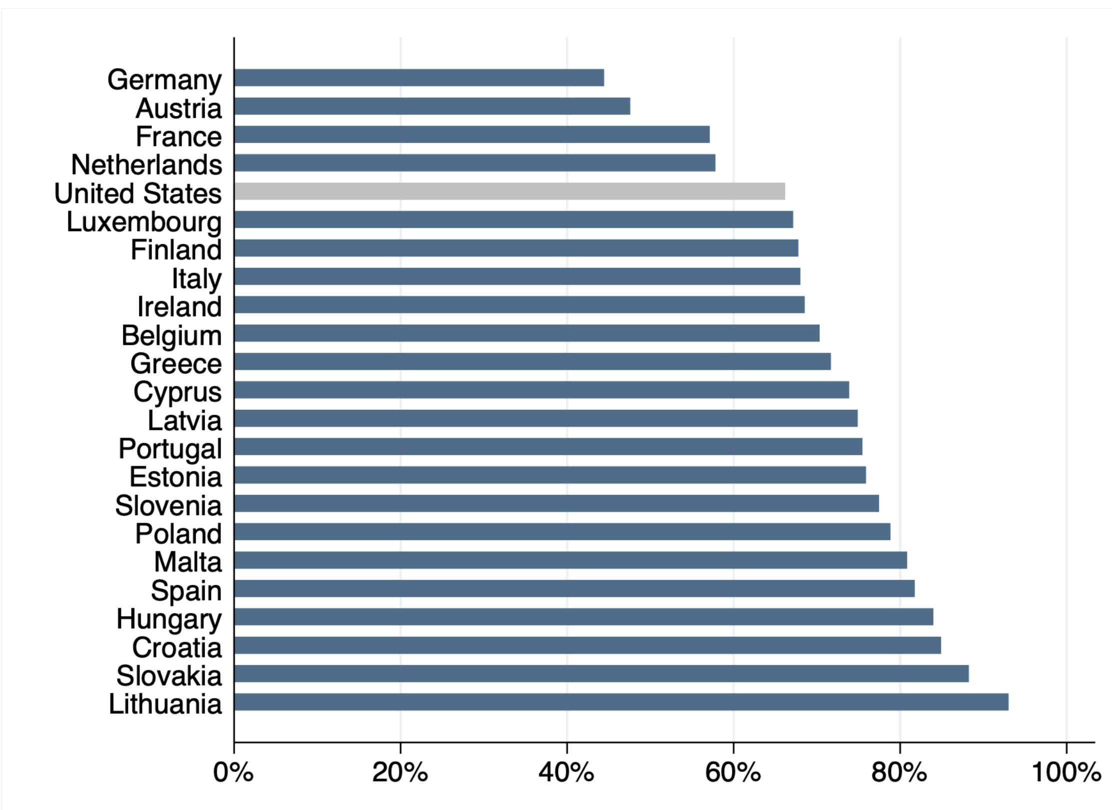

Figure B1: Homeownership rates in Europe and the United States (2008-2015)

European data is from the ECB Household Finance and Consumption Survey. U.S. homeownership is the average homeownership rate among U.S. natives from 2008-2018 American Community Survey.

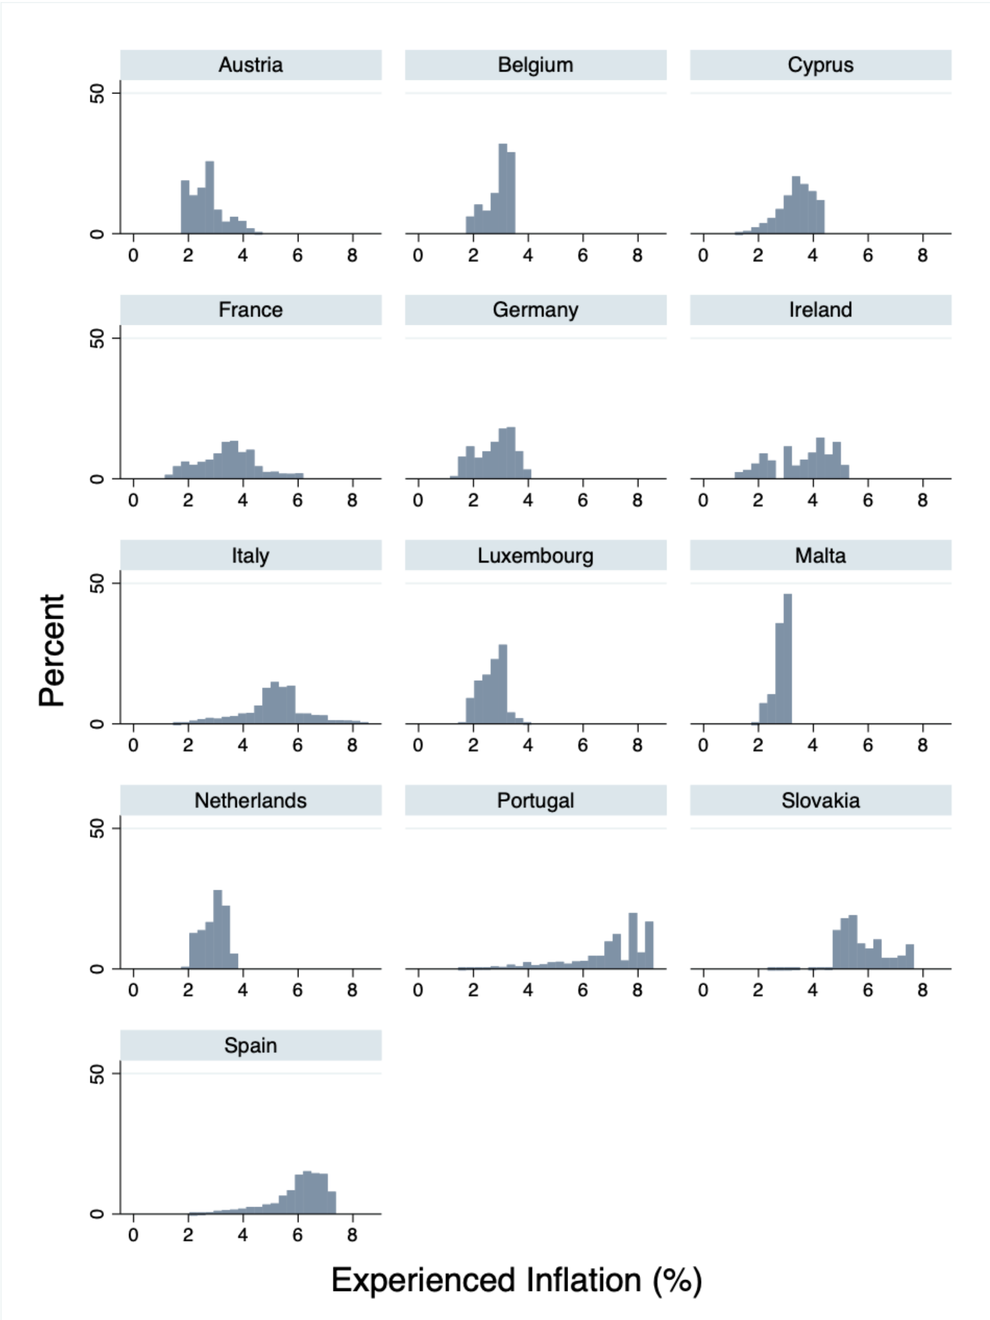

Figure B2: Distribution of experienced inflation, low- and moderate-inflation countries
Note: Inflation data sources described in the text.

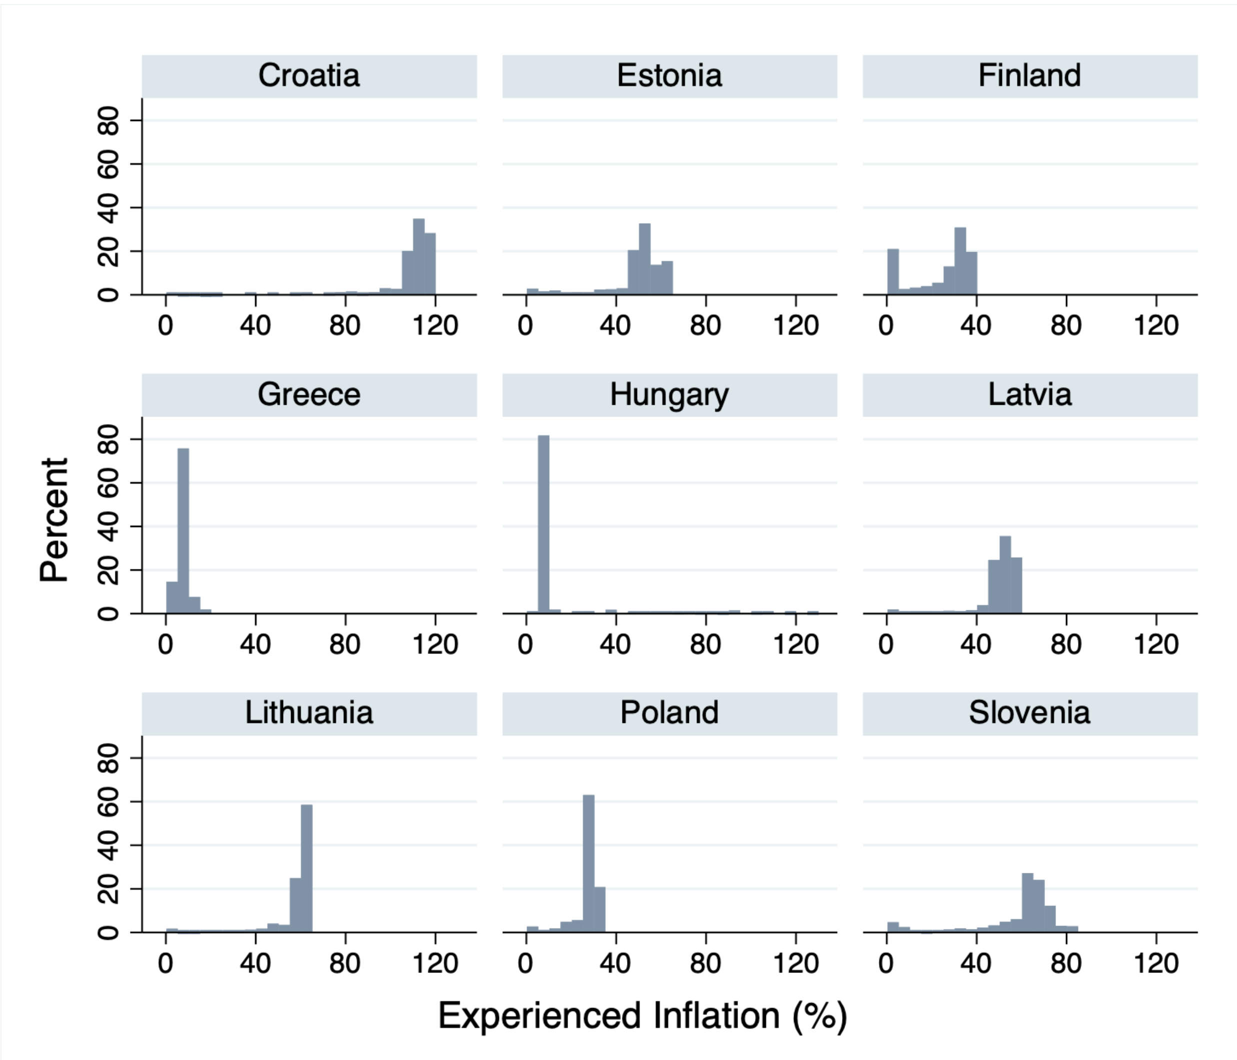

Figure B3: Distribution of experienced inflation, high-inflation countries

Note: Inflation data sources described in the text.

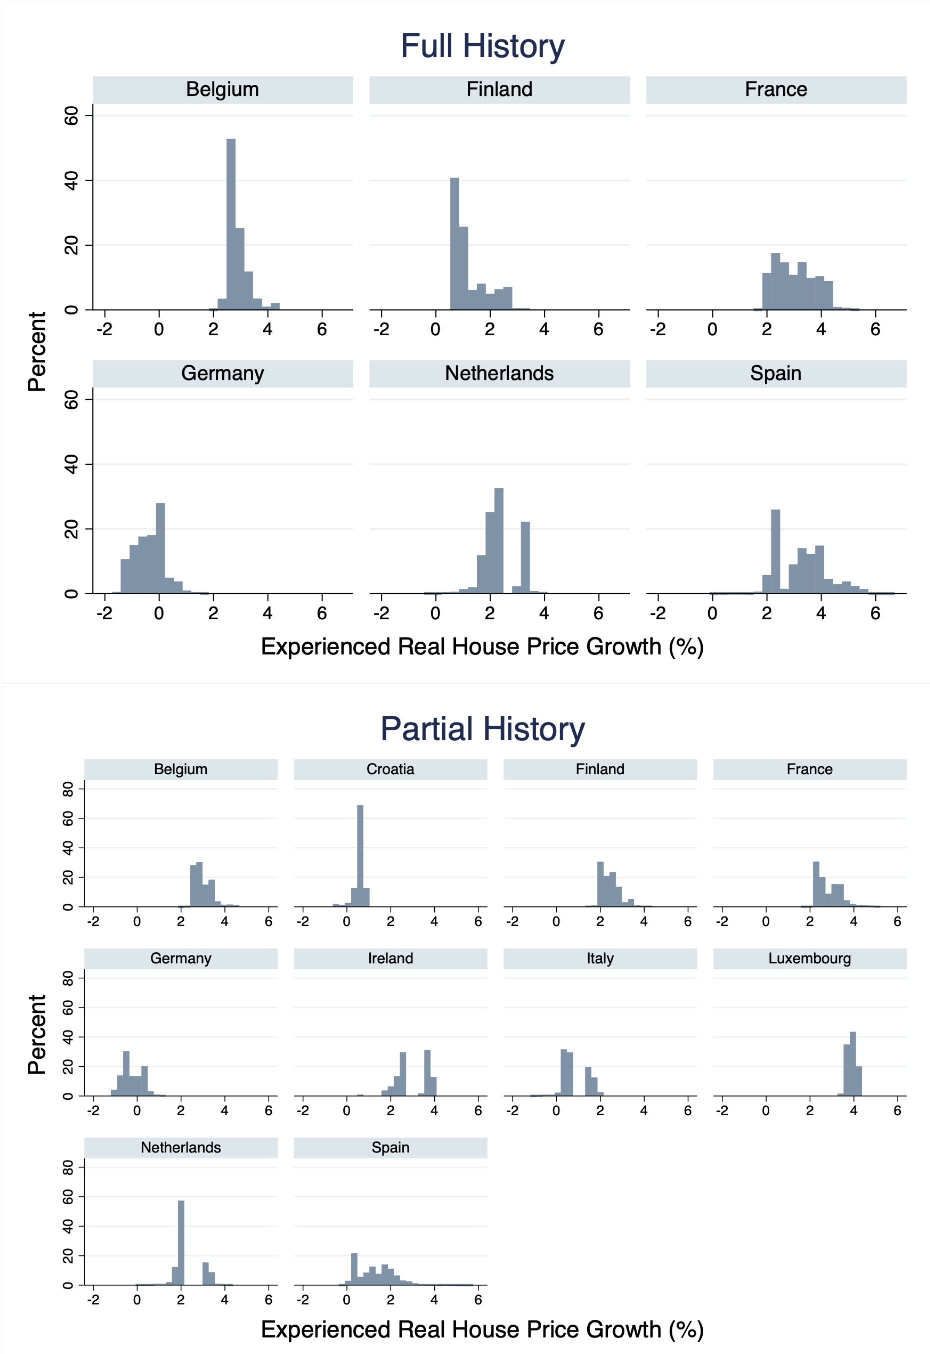

Figure B4: Distribution of experienced house price growth, by country

Note: House price data sources described in the text.

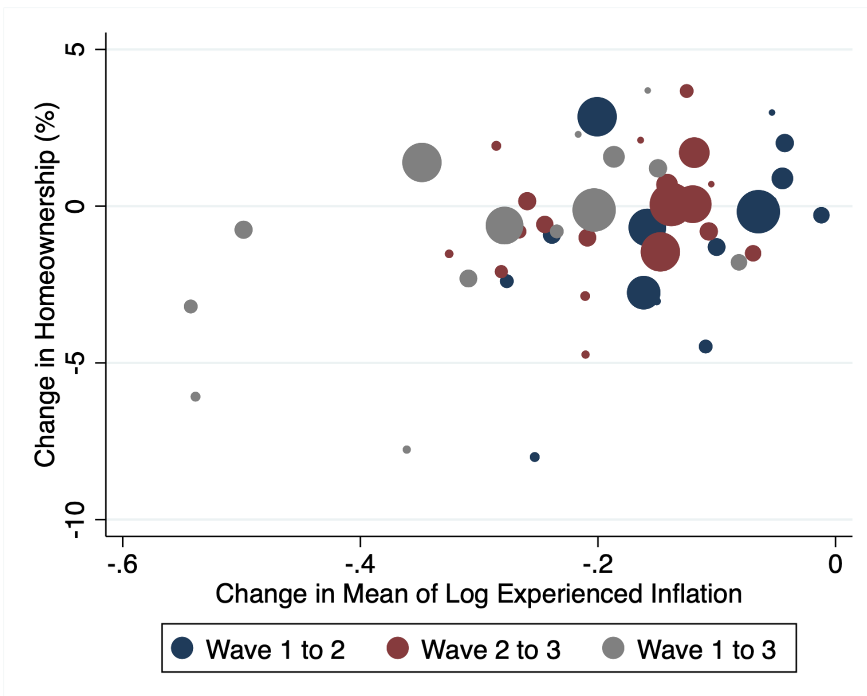

Figure B5: Changes in homeownership and experienced inflation across survey waves

Note: Figure plots the change in homeownership for each country across survey waves by the corresponding change in the average of households' log experienced inflation. Size indicates relative population. Homeownership from the HFCS. Inflation data sources and measures of experienced inflation described in the text.

Table B1: Using Experiences to Predict Price-to-Rent Ratio

Dependent Var: Price-to-Rent Ratio (Std.)	(1)	(2)	(3)	(4)	(5)
Avg. of Log Experienced Inflation (Std.)	0.29		-0.10		0.17
	(0.58)		(0.07)		(0.23)
Avg. Experienced House Price Growth (Full History, Std.)		1.08***	1.12***		
		(0.08)	(0.09)		
Avg. Experienced House Price Growth (Partial History, Std.)				0.90***	0.87***
				(0.15)	(0.11)
Country-Wave Observations	31	17	17	22	22
R^2	0.024	0.913	0.920	0.761	0.780

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: OLS regression coefficients. with standard errors clustered by country in parentheses. The data is the HFCS data, averaged using representative weights. Countries weighted by population from the World Bank. The dependent variable is the price-to-rent ratio relative to the country long-run average across survey years. Experience measures are calculated as the weighted average over the household head's lifetime, with linearly declining weights from the year before the survey to birth year (re-scaling the weights for partial history). Full history of experienced real house-price growth obtained from Knoll et al. (2017) for Belgium, Finland, France, Germany, and the Netherlands. Full history of experienced real house-price growth obtained from Bordo and Landon-Lane (2014) for Spain. Partial history of experienced real house-price growth obtained from the Federal Reserve Bank of Dallas for Belgium, Croatia, Finland, France, Germany, Ireland, Italy, Luxembourg, the Netherlands, and Spain. Dependent variable and all experience measures are standardized.

Table B2: Imputed versus Non-Imputed Data (HFCS)

Dep. Var: Own Main Residence	(1)	(2)	(3)
Log Exp. Inflation	2.661***	2.858***	2.120***
	(0.056)	(0.084)	(0.035)
Age	0.990***	0.986***	1.034***
	(0.002)	(0.002)	(0.001)
Male	1.017	1.009	1.142***
	(0.035)	(0.046)	(0.027)
Married	1.905***	1.905***	2.469***
	(0.089)	(0.116)	(0.074)
Widow	1.620***	1.578***	1.237***
	(0.112)	(0.137)	(0.060)
Divorced	1.139**	1.142*	0.793***
	(0.067)	(0.088)	(0.032)
Educ:2 (middle school)	0.793***	0.735***	0.933*
	(0.041)	(0.046)	(0.033)
Educ:3 (high school)	0.762***	0.730***	1.204***
	(0.037)	(0.044)	(0.037)
Educ:5 (college)	0.620***	0.610***	1.863***
	(0.035)	(0.044)	(0.066)
Has Child	1.459***	1.401***	1.468***
	(0.054)	(0.068)	(0.035)
Employed	1.373***	1.323***	1.641***
	(0.074)	(0.086)	(0.059)
Unemployed	1.288***	1.301***	0.845***
	(0.088)	(0.114)	(0.042)
Retired	1.374***	1.421***	1.094**
	(0.078)	(0.100)	(0.045)
Second Wave	1.417***	1.386***	0.980
	(0.060)	(0.083)	(0.027)
Third Wave	1.431***	1.364***	0.976
	(0.059)	(0.083)	(0.027)
Wealth and Income Deciles	Yes	Yes	No
Imputed Data?	Yes	No	No
Observations	214,133	128,431	219,857
Countries	22	22	22
Pseudo R^2	0.511	0.549	0.153

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Table reports exponentiated coefficients (odds ratios) from logit regressions with robust standard errors in parentheses on the HFCS multiple-imputation data in column (1) and the non-imputed data in columns (2) and (3). With the imputed data, the number of observations is the maximum N across the five imputations and the Pseudo R^2 is the average across the five imputations. Observations are weighted using the HFCS representative weights.

B.2 Survey of Homeowners

Recruitment We conducted a survey of 700 homeowners in six HFCS countries: Austria, Germany, Ireland, Italy, Portugal, and Spain. We recruited 100 participants from each country to our survey from Dynata's market research panel. We also recruited 100 participants through Amazon's Mechanical Turk (MTurk).[1] The results are similar across samples, so we combine them in the results below.

Participants recruited though Dynata were compensated for completing the survey using a combination of incentives including cash, gift cards, airline points, sweepstakes entries, and charity donations. Participants recruited from MTurk were paid $1 to complete our 3-minute survey.[2]

The survey was initially written in English, translated using translation services, and then edited by native speakers. The survey took place from March 5th to May 12th, 2020 for MTurk participants and from May 13 to May 15, 2020 for Dynata.

Survey Questions and Results After providing informed consent, participants were asked the following questions. Below we provide the exact question text and summary of responses from those who completed our survey.

1. In which country do you currently reside?
   →*Screened out 9 participants not from the target countries.*

Country	N	Percent
Austria	100	14
Germany	116	17
Ireland	105	15
Italy	150	21
Portugal	105	15
Spain	124	18
Total	700	100

2. Do you rent or own your home?
   →*Screened out 55 participants who did not select "Own."*

Response	N	Percent
Rent	0	0
Own	700	100
Other	0	0
Total	700	100

[1] We initially intended to recruit 100 participants from each country through MTurk, but were unable to recruit a sufficient sample during the COVID-19 crisis.

[2] Several participants were paid $0.50 before we increased the fee in an attempt to recruit more participants.

3. Why did you decide to buy rather than rent your home?
   Text box, free fill in.

4. What do you think are good reasons for buying a home? Please select all that apply.
   Order of options was randomized across participants, with "None of the Above" at the end.

Response	Percent Selected
Ownership provides peace of mind.	65%
Better selection of homes to buy than to rent.	19%
More flexibility to redecorate or remodel.	41%
House prices are likely to increase over time.	37%
Rent prices are likely to increase over time.	47%
Real estate is a good investment if there is inflation.	50%
Mortgage rates are low.	27%
Ownership provides tax benefits.	18%
Mortgage payments force me to save money.	15%
Mortgage payments are more predictable than rent prices.	31%
None of the above.	3%

5. Did concerns about inflation impact your decision to buy a home?[3]

Response	N	Percent
Yes	228	34
No	380	56
I am not sure	66	10
Total who know what inflation is	674	100

6. Have you personally experienced high inflation? *[Asked only if know what inflation is.]*

Response	N	Percent
Yes	283	42
No	391	58
Total who know what inflation is	674	100

7. Do you worry about inflation in the future? *[Asked only if know what inflation is.]*

Response	N	Percent
Yes	460	68
No	214	32
Total who know what inflation is	674	100

8. What do you think inflation will be next year? *[Asked only if know what inflation is.]*

[3] "I don't know what inflation is." was not available to the MTurk sample.

Mean	38
25th percentile	2
Median	3
75th percentile	10
SD	770
N who know what inflation is	674

9. Were you born in *[Country from Q1]*?

Response	N	Percent
Yes	634	91
No	66	9
Total	700	100

10. In which country were you born?*[Asked only if Q9=no.]*

11. What is your age?

Mean	44
25th percentile	34
Median	43
75th percentile	54
SD	13
N	700

Results by Experience In addition to the results shown in the main section of the paper, we analyze the key results by those who indicated that they have vs. have not personally experienced high inflation.

We find no difference between the two groups in their evaluation of real estate as an inflation hedge. Figure B6 shows that about half of respondents indicated that real estate is a good investment if there is inflation regardless of whether they personally lived through high inflation. However, the figure also shows that those with high inflation experience were more likely to say that their own homeownership decisions were impacted by inflation (45% vs. 26%) and more likely to be worried about inflation in the future (76% vs. 63%).

We also find that respondents who personally experienced high inflation have significantly higher expectations of next year's inflation. Excluding one outlier at 20,000%, expected inflation is 6.9% among those that did not experience high inflation and 11.3% among those who did. If we instead winsorize expected inflation at 20%, the 90th percentile, those with high inflation experiences have expected inflation about 1 pp higher than those who did not (means of 5.7% vs. 6.7%).

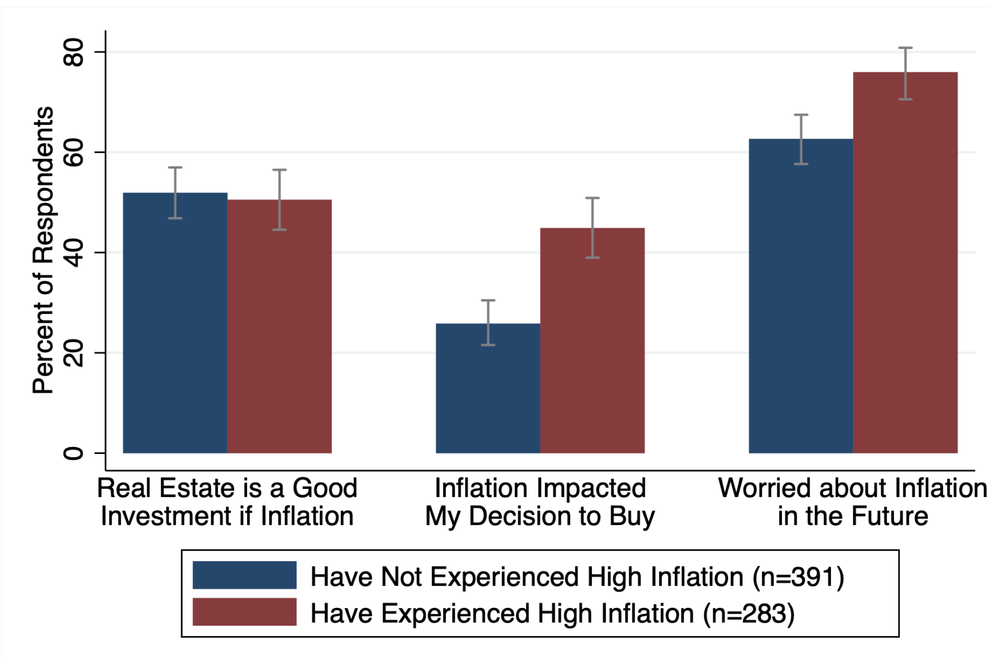

Figure B6: Inflation results by inflation experience.

Figure shows means and 95% confidence intervals, separately for respondents who reported that they have (or have not) personally experience high inflation. Total sample includes 674 respondents who know what inflation is.

B.3 Theoretical Framework

Generalized Theoretical Framework

In this section, we consider a more general version of our model. We relax some of the parameterizations and we allow for alternative inflation-protected assets (other than housing). To distinguish between the markets with and without an alternative inflation hedge, we introduce additional notation: $U_{t+1}(\cdot,\cdot,n_t)$ indicates utility when the alternative asset pays a nominal rate of n_t, and $U_{t+1}(\cdot,\cdot,r_t)$ indicates utility when the alternative asset pays a real rate of r_t.

In addition, we expand the model to allow for costs of ownership such as maintenance costs, property taxes, and costs of being a landlord (e.g., tenant protection, regulations, and rent control). We model the cost c as proportional to the value of housing and payable at $t+1$, amounting to cM_{t+1},[4] and assume that initial wealth is sufficiently high relative to housing costs to be positive under any realization, to accommodate the log utility specification.

We derive the utilities of renting and owning under fixed- and variable-rate financing in market with and without an alternative inflation hedge.

Housing as the only inflation hedge In this scenario, the alternative asset pays a nominal rate n_t between t and $t+1$, known to households at time t. Under this assumption, the household's expected utility conditional on renting is

$$\begin{aligned} E_t\left[U_{t+1}(R,h_t,n_t)\right] &= E_t\left[u\left(\frac{w_{t+1}(R,h_t,n_t)}{1+\pi_{t+1}}\right)\right] \\ &= \log\left((w_t-h_t)(1+n_t)\right) - E_t[\log(1+\pi_{t+1})], \end{aligned} \tag{B.1}$$

where $w_{t+1}(R,h_t,n_t)$ is the expected nominal wealth in $t+1$ conditional on renting at the prevailing prices.

Households' expected utility conditional on buying with a fixed-rate mortgage is

$$\begin{aligned} E_t\left[U_{t+1}(FR,m_t,n_t)\right] &= E_t\left[u\left(\frac{w_{t+1}(FR,m_t,n_t)}{1+\pi_{t+1}}\right)\right] \\ &= E_t[\log(M_t(1+\pi_{t+1})(1+g_{t+1})(1-c) - m_t(1+n_t^f) \\ &\quad + (w_t-(M_t-m_t))(1+n_t)) - \log(1+\pi_{t+1})], \end{aligned} \tag{B.2}$$

where $w_{t+1}(FR,m_t,n_t)$ is expected nominal wealth in $t+1$ conditional on buying and financing with a fixed-rate mortgage at the prevailing prices.

[4] The results are qualitatively unchanged if c grows with inflation instead of house prices.

Similarly, buying with a variable-rate mortgage m_t yields

$$\begin{aligned} E_t\left[U_{t+1}(VR, m_t, n_t)\right] &= E_t\left[u\left(\frac{w_{t+1}(VR, m_t, n_t)}{1+\pi_{t+1}}\right)\right] \\ &= E_t[\log(M_t(1+\pi_{t+1})(1+g_{t+1})(1-c) - m_t(1+r_t^v)(1+\pi_{t+1}) \\ &\quad + (w_t - (M_t - m_t))(1+n_t)) - \log(1+\pi_{t+1})], \end{aligned} \tag{B.3}$$

where $w_{t+1}(VR, m_t, n_t)$ is expected nominal wealth in $t+1$ conditional on buying and financing with a variable-rate mortgage at the prevailing prices.

Housing with alternative inflation hedge. In the second scenario, the alternative asset is inflation-protected and pays a *real* rate r_t between t and $t+1$, known to households at time t. Here, the expected utility conditional on renting is

$$\begin{aligned} E_t\left[U_{t+1}(R, h_t, r_t)\right] &= E_t\left[u\left(\frac{w_{t+1}(R, h_t, r_t)}{1+\pi_{t+1}}\right)\right] \\ &= E_t\left[\log\left((w_t - h_t)(1+r_t)(1+\pi_{t+1})\right) - \log(1+\pi_{t+1})\right] \\ &= E_t\left[\log\left((w_t - h_t)(1+r_t)\right)\right], \end{aligned} \tag{B.4}$$

where $w_{t+1}(R, h_t, r_t)$ is the expected nominal wealth in $t+1$ conditional on renting at prevailing prices.

The expected utility conditional on buying with a fixed-rate mortgage of value m_t is

$$\begin{aligned} E_t\left[U_{t+1}(FR, m_t, r_t)\right] &= E_t\left[u\left(\frac{w_{t+1}(FR, m_t, r_t)}{1+\pi_{t+1}}\right)\right] \\ &= E_t[\log(M_t(1+\pi_{t+1})(1+g_{t+1})(1-c) - m_t(1+n_t^f) \\ &\quad + (w_t - (M_t - m_t))(1+r_t)(1+\pi_{t+1})) - \log(1+\pi_{t+1})] \\ &= E_t\left[\log(M_t(1+g_{t+1})(1-c) - \frac{m_t(1+n_t^f)}{1+\pi_{t+1}} + (w_t - (M_t - m_t))(1+r_t))\right], \end{aligned} \tag{B.5}$$

where $w_{t+1}(FR, m_t, r_t)$ is the expected nominal expected wealth in $t+1$ conditional on buying with a fixed-rate mortgage at prevailing prices.

The expected utility conditional on buying with a variable-rate mortgage of value m_t is

$$\begin{aligned} E_t\left[U_{t+1}(VR, m_t, r_t)\right] &= E_t\left[u\left(\frac{w_{t+1}(VR, m_t, r_t)}{1+\pi_{t+1}}\right)\right] \\ &= E_t[\log(M_t(1+\pi_{t+1})(1+g_{t+1})(1-c) - m_t(1+r_t^v)(1+\pi_{t+1}) \\ &\quad + (w_t - (M_t - m_t))(1+r_t)(1+\pi_{t+1})) - \log(1+\pi_{t+1})] \\ &= E_t[\log(M_t(1+g_{t+1})(1-c) - m_t(1+r_t^v) + (w_t - (M_t - m_t))(1+r_t))], \end{aligned} \tag{B.6}$$

where $w(VR, m_t, r_t)$ is the nominal expected wealth in $t+1$ conditional on buying with a variable-rate mortgage at prevailing prices.

The generalized model generates four equations capturing the sensitivity of utility to experiences. The pointwise derivatives of the utility difference between buying and renting with respect to inflation for each of the two types of mortgages and alternative assets are:

$$\frac{\partial}{\partial \pi_{t+1}} \left[U_{t+1}(FR, m_t, n_t) - U_{t+1}(R, h_t, n_t)\right]\Bigg|_{\pi,g} = \frac{M_t(1+g)(1-c)}{w_{t+1}(FR, m_t, n_t|\pi, g)} > 0 \tag{B.7}$$

$$\frac{\partial}{\partial \pi_{t+1}} \left[U_{t+1}(VR, m_t, n_t) - U_{t+1}(R, h_t, n_t)\right]\Bigg|_{\pi,g} = \frac{M_t(1+g)(1-c) - m_t(1+r_t^v)}{w_{t+1}(VR, m_t, n_t|\pi, g)} > 0 \tag{B.8}$$

$$\frac{\partial}{\partial \pi_{t+1}} \left[U_{t+1}(FR, m_t, r_t) - U_{t+1}(R, h_t, r_t)\right]\Bigg|_{\pi,g} = \frac{m_t(1+n_t^f)}{w_{t+1}(FR, m_t, r_t|\pi, g)(1+\pi)} > 0 \tag{B.9}$$

$$\frac{\partial}{\partial \pi_{t+1}} \left[U_{t+1}(VR, m_t, r_t) - U_{t+1}(R, h_t, r_t)\right]\Bigg|_{\pi,g} = 0 \tag{B.10}$$

Because the partial derivatives are weakly positive in all four cases, our model predicts that homeownership will be increasing in experienced inflation in any market with a mix of funding opportunities and access to inflation hedges. Equations (B.7) and (B.8) confirm Prediction 1 from the main text, generalized in allowing for housing costs.

With equation (B.9) > 0, we confirm Prediction 1 under fixed-rate financing in a market with alternative inflation hedges. Here, the benefit of homeownership among households who have experienced higher inflation is that they can borrow at what they perceive to be a low real rate. In equation (B.10), we find no response of homeownership to experienced inflation in a market with an alternative inflation hedge (and thus no real-asset motivation) and variable-rate financing (and thus no cheap borrowing motivation).

We also find that Prediction 2 is robust to the existence of alternative inflation hedges. Equations (B.7) and (B.8) mirror the derivation in the main text. Turning to the scenario with alternative inflation hedges, (B.9)-(B.10)=(B.9)> 0 implies that the effect of experienced inflation continues to be stronger with fixed-rate financing.

Finally, we confirm that the results of Prediction 3 are robust to the availability of inflation hedges. The pointwise derivatives of the utility difference between buying and renting with respect to house-price growth are:

$$\frac{\partial}{\partial g_{t+1}} \left[U_{t+1}(FR, m_t, n_t) - U_{t+1}(R, h_t, n_t)\right]\Bigg|_{\pi,g} = \frac{M_t(1+\pi)(1-c)}{w_{t+1}(FR, m_t, n_t|\pi, g)} > 0 \tag{B.11}$$

$$\frac{\partial}{\partial g_{t+1}} \left[U_{t+1}(VR, m_t, n_t) - U_{t+1}(R, h_t, n_t)\right]\Bigg|_{\pi,g} = \frac{M_t(1+\pi)(1-c)}{w_{t+1}(VR, m_t, n_t|\pi, g)} > 0 \tag{B.12}$$

$$\frac{\partial}{\partial g_{t+1}} \left[U_{t+1}(FR, m_t, r_t) - U_{t+1}(R, h_t, r_t)\right]\Bigg|_{\pi,g} = \frac{M_t(1+\pi)(1-c)}{w_{t+1}(FR, m_t, r_t|\pi, g)} > 0 \tag{B.13}$$

$$\frac{\partial}{\partial g_{t+1}} \left[U_{t+1}(VR, m_t, r_t) - U_{t+1}(R, h_t, r_t)\right]\Bigg|_{\pi,g} = \frac{M_t(1+\pi)(1-c)}{w_{t+1}(VR, m_t, r_t|\pi, g)} > 0 \tag{B.14}$$

Equations (B.11) and (B.12) mirror the derivation in the main text. Because the relationship between house price growth and the benefit of homeownership is independent of the type of financing and availability of inflation hedges, equations (B.13) and (B.14) are also positive.

Simulations of the Model

To further illustrate the influence of experience-based learning in all four settings, we simulate tenure decisions under different plausible inflation-exposure scenarios. We will also use the simulations to consider a wider parameter space than in our baseline setting from Section 2.2 and under alternative assumptions.

Baseline. To simulate the model, we parametrize beliefs of agents who are influenced by past macro histories and, for comparison, of agents with rational beliefs. We start with the most simplistic version, by assuming that past macro histories induce deterministic beliefs that are exactly the same as what they observed in the past. For example, a household who sees 5% inflation in t would expect 5% inflation in $t+1$.

We explore the influence of past realizations of inflation on agents' tenure decisions under this parameterization in Figure B7(a). For each historical inflation level, we plot the rental price (as a percent of the house price, h_t/M_t) at which the agent is indifferent between renting and owning, separately for each the four markets: fixed- vs. variable-rate mortgage and with an alternative asset that pays a known nominal or real return. A lower h_t/M_t indicates a higher value of ownership relative to renting.[5]

Figure B7(a) shows that, in all four markets, the slope is (weakly) negative, indicating that all else equal, higher past inflation increases the willingness to pay for ownership. Second, the effect of past inflation experiences on ownership is stronger when households have access to fixed- rather than variable-rate mortgages, evidenced by the steeper slope of the blue (darker) relative to the corresponding red (lighter) lines. Third, the graph shows that the effect of experienced inflation will be stronger in a market without alternative inflation hedges as the solid lines (for markets without inflation hedges) are steeper than the corresponding dashed lines (for markets with alternative hedging opportunities).

For comparison, in Figure B7(b), we plot the corresponding graph for a household who has rational beliefs. In this case, past realizations of inflation have no bearing on inflation expectations and therefore do not impact the relative value of ownership. All lines overlap.

Note that there is a level of experienced inflation (in this case, 4%), at which the experienced-based household has the same beliefs as the rational household. If the experienced-

[5] We also assume the household expects real house-price growth $g_{t+1} = 2\%$, has log utility over real wealth, initial wealth $w_t = 200,000$, house price $M_t = 100,000$, loan-to-value ratio $m_t/M_t = 0.8$, ownership costs $c = 2\%$, the alternative asset offers either a real return $r_t = 2\%$ or a nominal return $n_t = 6.1\%$ (corresponding to 4% anticipated inflation), and we assume mortgage rates carry a 1% premium relative to the alternative asset (i.e.,$n_t^f = 7.1\%$ and $r_t^v = 3\%$).

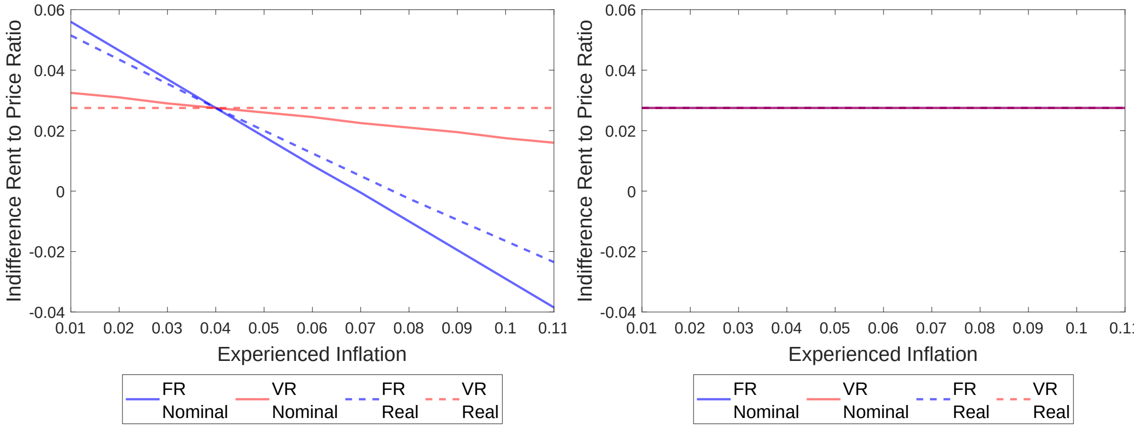

(a) Experienced-Based Beliefs (b) Rational Beliefs

Figure B7: Simple simulation of the model

based household lives through higher inflation, she is willing to pay more than the rational household for ownership. If she lives through lower inflation, she is willing to pay less.

In Figure B8, we present results under less simplistic parameterizations of experiences, namely, assuming instead that experienced-based households are uncertain about future inflation and real house-price growth. Specifically, we model households as having log-normal, uniform, or normally distributed beliefs about inflation and house-price growth. Along the x-axis we vary the mean of the experienced-based inflation belief distribution, fixing the standard deviation of beliefs about inflation and beliefs about house-price growth. Roughly consistent with the actual data, we assume the standard deviation of inflation beliefs is 6% and that real house-price growth is distributed with a mean of 2% and a standard deviation of 7%. Under all three distributional assumptions, the theoretical predictions hold.

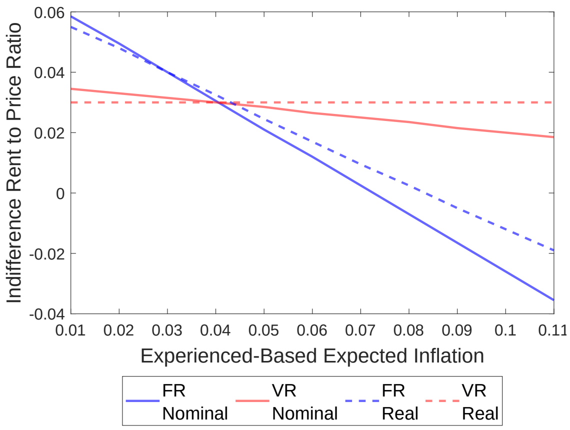

(a) Lognormal Distributions

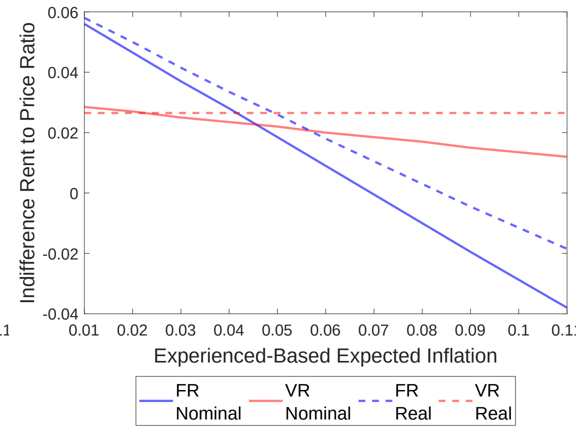

(b) Uniform Distributions

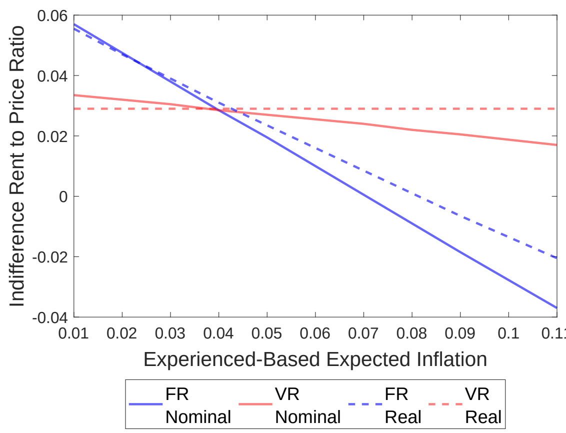

(c) Normal Distributions

Figure B8: Simulation with alternative distributions of beliefs

Robustness of Prediction 1. In the main text, we restrict the parameter space by requiring $M_t(1+g_{t+1})(1-c) > m_t(1+r_t^v)$. This condition fails when expected house-price growth is low, costs are high, LTV is high, and variable mortgage rates are high. Most predictions hold more generally, but, as we show in Section 2.2, the positive influence of past inflation on the value of ownership under variable-rate financing depends on this restriction in the scenario without an alternative inflation hedge. Assuming beliefs are normally distributed, in Figure B9(a) we show that Prediction 1 is robust to low beliefs about future house-price growth ($g_{t+1} \sim N(-2\%, 1\%)$), high costs of ownership ($c = 10\%$), and high variable mortgage rates ($r_t^v = 5\%$ compared to 3% in the benchmark simulations). In Figure B9(b), we increase LTV all the way to 90% and find a slightly upward slope. That is, experiencing higher inflation predicts *lower* value of ownership for households who can finance with a variable-rate mortgage in a market with no alternative inflation hedges. However, the response remains strong in the predicted direction for households with access to fixed-rate financing. Assuming a mix of financing opportunities, the simulations imply that Prediction 1 still holds in the aggregate.

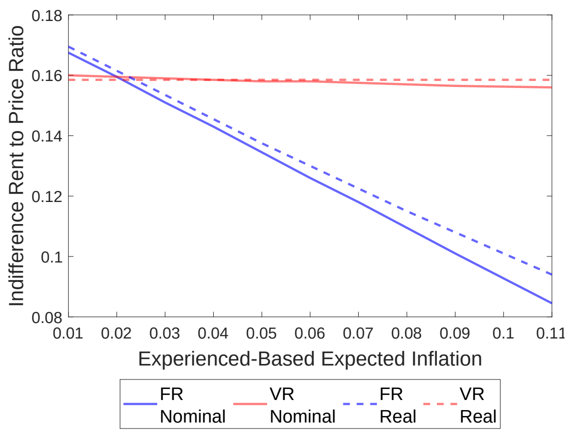

(a) Low g, high c, high variable-rate

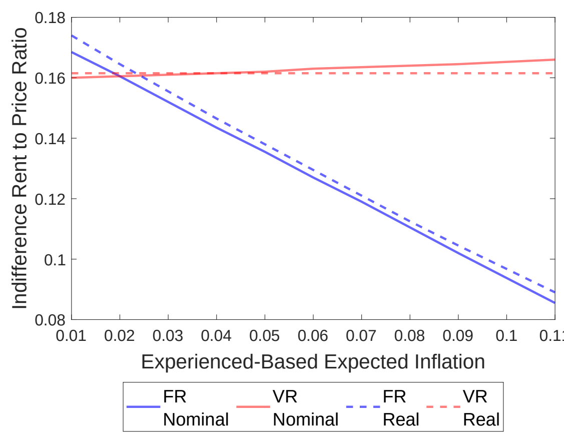

(b) Low g, high c, high variable-rate, high LTV

Figure B9: Stress test of Prediction 1

Robustness of Prediction 2. In a market with alternative inflation hedges, our model predicts an unambiguously stronger response to experienced inflation for households with access to fixed-rate compared to variable-rate financing. We argued in Section 2.2 that this is likely also the case in a market without alternative inflation hedges. In the simulations thus far, we have seen this evidenced by the fact that the solid blue line is steeper than the solid red line. In Figure B10, we test the robustness by simulating conditions least favorable to Prediction 2. Specifically, this prediction may fail when 1) $m_t(1+r_t^v)$ is small and 2)

$$\frac{M_t(1+g_{t+1})(1-c)}{w_{t+1}(FR, m_t, n_t)} << \frac{M_t(1+g_{t+1})(1-c)}{w_{t+1}(VR, m_t, n_t)}.$$

In Figure B10(a) we show that, although the magnitude drops, the prediction holds with low real rates relative to the nominal ($r_t = r_t^v = 1\%$, $n_t = n_t^f = 7\%$), a higher expected real house-price growth of 6%, and a 0% cost of ownership.[6] Lowering LTV to 20% (Figure B10(b)) greatly reduces the magnitude, however Prediction 2 still holds.

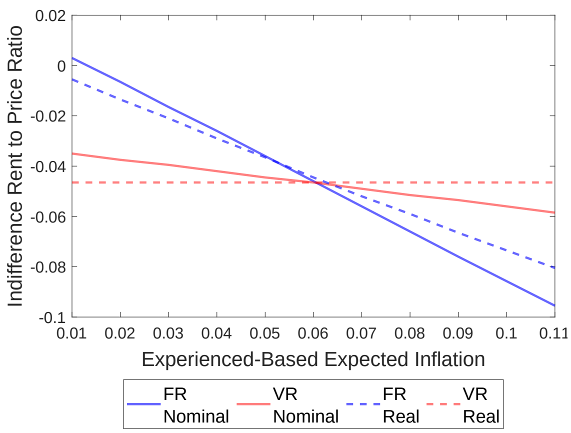

(a) Low real rates, high g, and low c

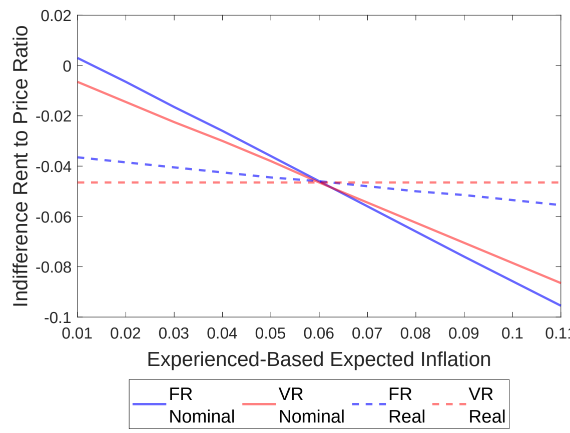

(b) Low real rates, high g, low c, and low LTV

Figure B10: Stress tests of Prediction 2

[6]We assume beliefs are log-normally distributed but results are similar for other distributions.

Loan-to-Value. In the baseline simulation, we assume the mortgage value is 80% of the value of the home. As discussed above, our predictions appear to be sensitive to loan-to-value ratios. Maintaining the benchmark parameters and varying only loan-to-value ratios, we find that the key predictions of our model hold except at LTVs above 90%, as demonstrated in Figure B11.

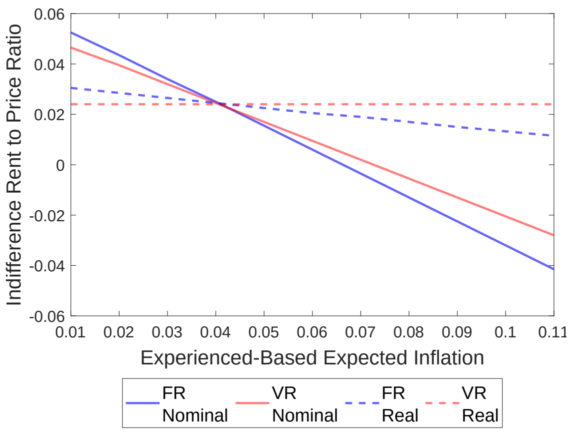

(a) 20% LTV.

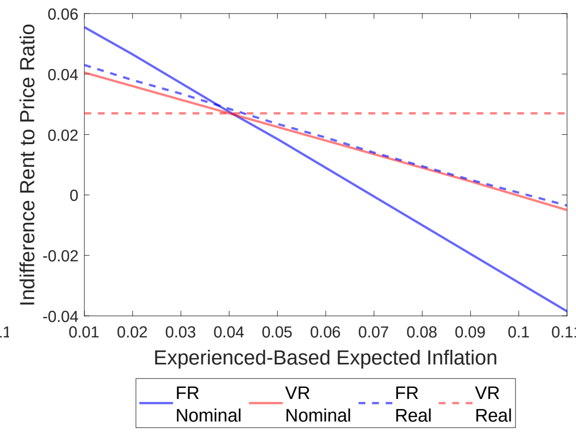

(b) 50% LTV.

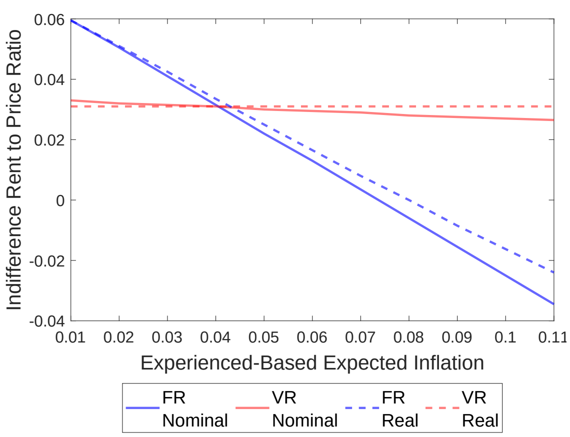

(c) 90% LTV.

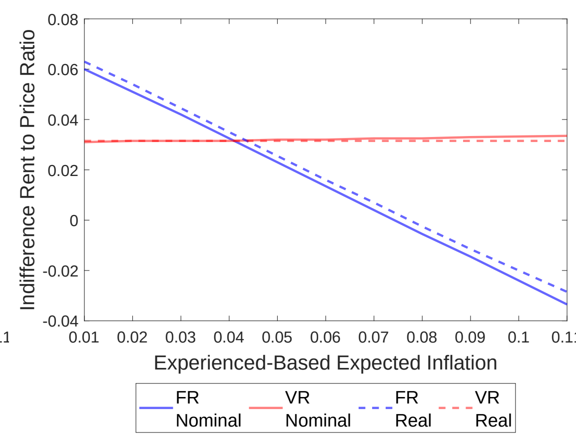

(d) 99% LTV.

Figure B11: Simulation with alternative loan-to-value ratios

Housing Booms and Crises. We now explore the robustness of our predictions to more extreme changes in real house-price growth, as they may occur during housing booms or crises. To do this, we vary the assumptions about the mean real house-price growth, assuming beliefs about future inflation and house-price growth are normally distributed. Consistent with Prediction 3, we see in Figures B12(a) and (b) that higher mean g (i. e., a housing boom) increases the valuation of ownership overall, but does not meaningfully change Predictions 1 and 2. Similarly, a low mean $g = -2\%$ (i. e., a housing crisis) lowers overall ownership but does not affect our predictions, as demonstrated in Figure B12(c). Even in the case of an extreme housing crisis with mean $g = -20\%$ (Figure B12(d)), when mortgages would be underwater in the majority of the parameter space, Predictions 1 and 2 appear largely robust. At this very low g, we do see a reversal of Prediction 1 (though small in magnitude) in markets with variable-rate mortgages and no alternative inflation hedges.

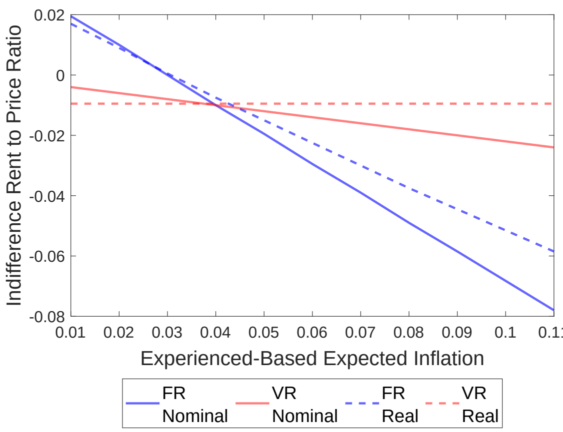

(a) Mean g 6%.

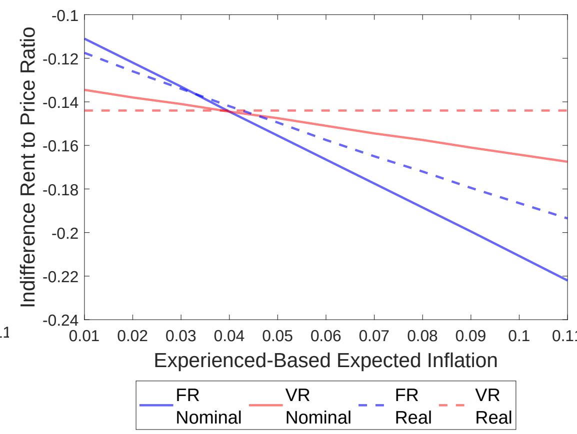

(b) Mean g 20%.

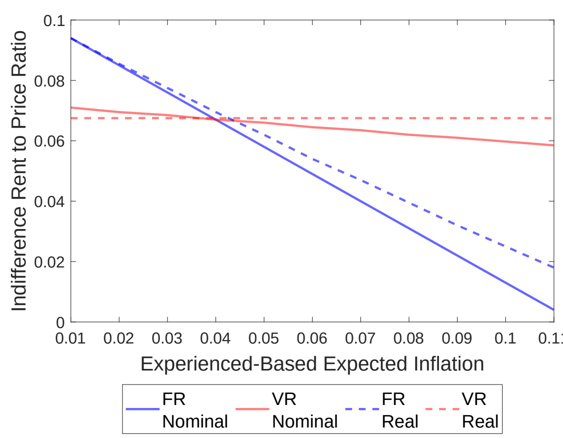

(c) Mean g -2%.

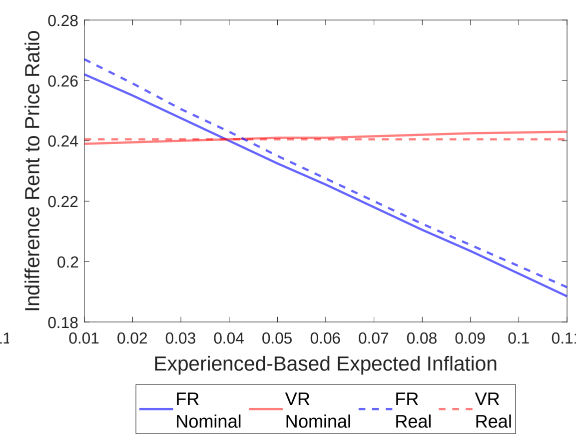

(d) Mean g -20%.

Figure B12: Simulation of extreme real house price growth

Risk Aversion. The theoretical model assumes log utility. Here, we show that the results are robust to agents having more or less risk-averse preferences. We assume constant relative risk aversion and show that the predictions hold for a range of possible risk aversions in Figure B13.

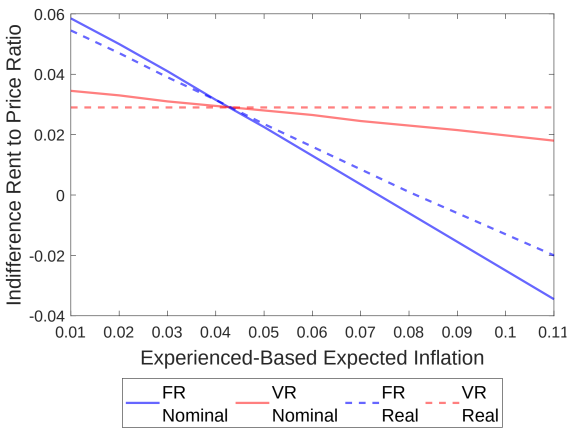

(a) Risk Neutral (CRRA $\theta = 0$)

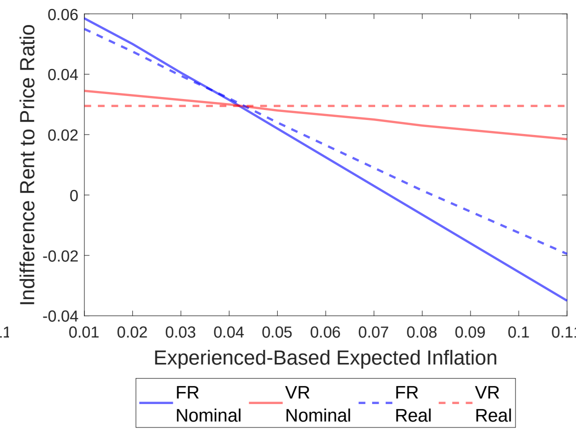

(b) Low Risk Aversion (CRRA $\theta = 0.25$)

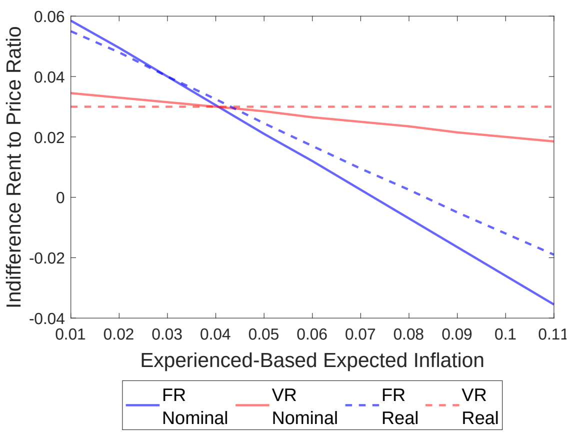

(c) Log Utility (CRRA $\theta = 1$), benchmark

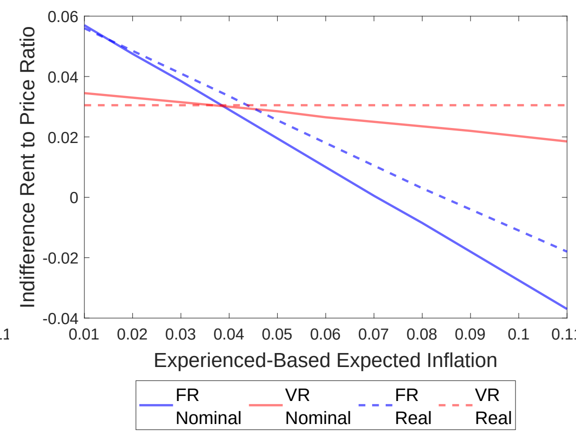

(d) High Risk Aversion (CRRA $\theta = 2$

Figure B13: Simulation with alternative levels of risk aversion

Inflation Variance. In the theoretical framework and in our empirical analysis, we model the level of experienced inflation as affecting the level of beliefs about future inflation. However, we can also think of the variance in experiences as affecting the variance of the belief distribution. In Figure B14(a), we replicate the benchmark graph with normally distributed beliefs, varying the mean of the distribution and holding the standard deviation at 6%. In Figure B14(b), we instead hold the mean of inflation beliefs fixed at 4% and vary the standard deviation of inflation beliefs across the x-axis. Compared to changes in the means, we see little movement in the value of ownership as we vary the standard deviation of beliefs. The only detectable effect is a slight lowering in the value of ownership under fixed-rate financing in a market with alternative inflation hedges. By financing at a fixed-rate, the household gives up the inflation-hedging benefits of the alternative asset.

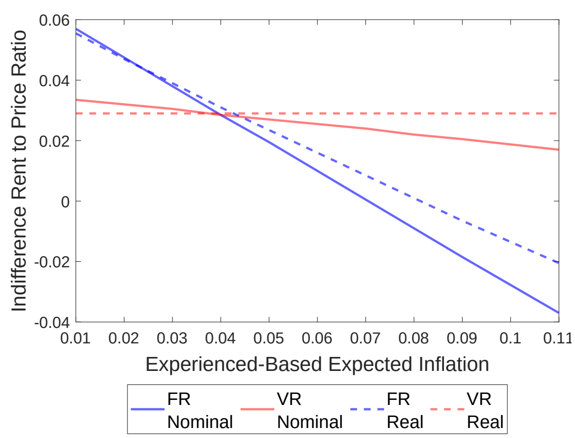

(a) Vary Mean Inflation

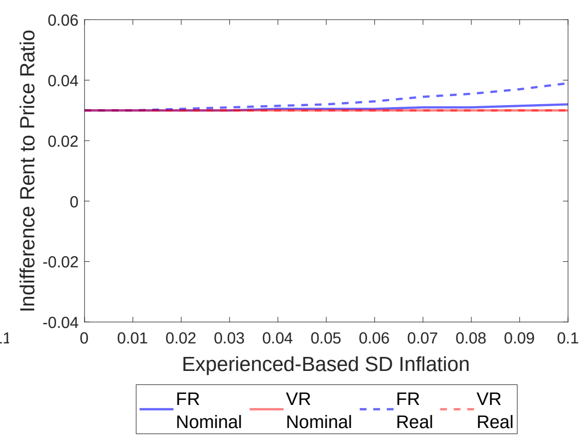

(b) Vary SD Inflation

Figure B14: Simulations of experiences affecting the mean and variance of beliefs, normally distributed beliefs

B.4 Alternative Measures of Inflation Experience

In Appendix-Table B3, we test several alternative methods of controlling for inflation experiences. In this table, we standardize all continuous experience measures to facilitate comparisons of the magnitudes.

First, we demonstrate the robustness of our main result to the treatment of households with high inflation experience. In our baseline analyses, we apply a log transform to average experienced inflation over the lifetime to account for non-linearity in the effects and to limit the impact of high-experience outliers. In column (1), we report the coefficient on the standardized measure. We find that a one-standard deviation increase in the log of experienced inflation predicts a 114% increase in the odds of homeownership, or an increase from 65% to 80%.

In column (2), we estimate a linear effect of experienced inflation. As in our main specification, we find a significant positive effect of experiences. A one-standard deviation increase in experienced inflation predicts a 73% increase in the odds of homeownership, or an increase in the probability of ownership from 65% to 76%.

In column (3), we estimate the linear effect of experienced inflation only for the subset of countries that do not have any high-inflation outliers (or households with lifetime experienced inflation above 10%). High-inflation countries include Croatia, Estonia, Finland, Greece, Hungary, Latvia, Lithuania, Poland, and Slovenia. Within this subsample, the predicted relationship between experienced inflation and homeownership is slightly stronger: a one-SD increase in experienced inflation predicts an 113% increase in the odds of homeownership, or an increase in the probability of ownership from 65% to 80%.

In column (4), we return to the full sample, but winsorize lifetime experienced inflation at 10%. We also include an indicator for any household above the threshold. This allows us to estimate effects on the entire sample, while accounting for the non-linearity we observed in the aggregate data. In this specification, we find that a one-SD increase in winsorized experienced inflation is associated with a 175% increase in the odds of homeownership, or an increase in the predicted probability of ownership from 65% to 84%. Interestingly, we estimate a negative effect of being above the 10% threshold, indicating that the predicted probability of homeownership is lower for high-inflation households compared to those at 10%. We chose a 10% threshold to winsorize the data as it is a clear break in the distribution of lifetime experiences and, coincidentally, corresponds to the visual trend break in the aggregate data. We find qualitatively similar results if we winsorize instead at 40% or 70%, which correspond to other natural breaks in the distribution.

In column (5), rather than winsorizing the lifetime average of experienced inflation, we cap each year's inflation at 25% before calculating a weighted average over the lifetime. In this way, we limit the effect that any given year's inflation has on lifetime experiences. We also include an indicator for whether the household ever lived through inflation above the threshold (i. e., whether any year in their experienced inflation measure was above 25%). We find that the measure of capped experienced inflation positively and significantly predicts higher homeownership, with a one-SD increase predicting an increase in the probability of ownership from 65% to 81%. There is no significant impact of having lived through inflation above 25%. Using alternative annual thresholds of 50% and 100%, we continue to find a positive significant effect of the winsorized experience measure, but mixed effects of ever experiencing inflation above the cap.

In columns (6) through (8), we test several conceptually different measures of experienced inflation. First, we test the hypothesis that inflation volatility predicts individual homeownership. We calculate individual experienced inflation volatility as the standard deviation of inflation over the lifetime. We find that a one-SD increase in inflation volatility predicts a 72% increase in the odds of homeownership, or an increase in the predicted probability from 65% to 76%. While still sizable, the magnitude of the effect is smaller than the level of experienced inflation. As we show in Appendix B.3, this weaker result is consistent with our model simulations.

We have also implemented an extended version of the AR(1) model as described in Malmendier and Nagel (2015) to estimate households' one-year inflation prediction from their lifetime experienced inflation. Extending the AR(1) model to our context is not straightforward as one-year inflation is unlikely to be relevant for homeownership decisions, which are long-term investments. Hence, we have to take a stance on the relevant forecast period for homeownership decisions as well as how individuals make long-term forecasts and iterate the one-year belief formation process forward.

Before we choose a set of assumptions, we start from simply relating homeownership to the original Malmendier-Nagel one-year forecast, despite the mismatch in horizon. We use their estimate of 3.044 for the gain parameter, and implement their AR(1) model to estimate households' one-year inflation prediction from their lifetime experienced inflation. In column (7), we find that the predicted inflation measure over the next-year significantly predicts the likelihood of being a homeowner.[7]

Turning to the more relevant long-term horizon, we take the approach to let individuals recursively estimate an AR(1) model of inflation up to the year before the survey. We then assume that they use the estimated coefficients (as of the survey year) to iterate the model forward T periods to make a projection of inflation in each subsequent year, T. As shown in column (8), we find that the five-year aggregate inflation forecast significantly predicts ownership. We also find significant, though smaller, relationships using the predicted ten- and twenty-year inflation forecasts.

In addition to the choice of timing, there are alternative ways of modeling long-term forecast formation, for example, assuming people anticipate future learning or assuming that people project their one-year forecast onto all future years. For these reasons, we choose to use the lifetime weighted average approach of measuring macroeconomic experiences in our main analyses.

[7]The results are robust to using alternative gain parameters ranging from 2 to 5.

Table B3: Alternative Measures of Inflation Experiences and Household-Level Homeownership, Standardized Coefficients

Dependent Var: Own Main Residence	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Experienced Inflation (std.)	2.14***							
	(0.04)							
Experienced Inflation (std.)		1.73***						
		(0.03)						
Experienced Inflation (Exclude High-Inflation Countries, std.)			2.13***					
			(0.06)					
Experienced Inflation (Wins. at 10, std.)				2.75***				
				(0.07)				
High-Inflation Experienced (Above 10)				0.44***				
				(0.03)				
Experienced Inflation (Wins. at 25 in Each Year, std.)					2.25***			
					(0.04)			
Any Year High-Inflation Experienced (Above 25)					0.97			
					(0.04)			
Standard Deviation of Experienced Inflation (std.)						1.72***		
						(0.09)		
Predicted AR(1) 1-Year Inflation Forecast (std.)							1.44***	
							(0.04)	
Predicted AR(1) 5-Year Inflation Forecast (std.)								1.13***
								(0.03)
Observations	214,133	214,133	142,201	214,133	214,133	214,133	214,133	214,133
Pseudo R^2	0.511	0.496	0.523	0.516	0.511	0.487	0.486	0.482

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports exponentiated coefficients (odds ratios) from logit regressions with robust standard errors in parentheses. Continuous experience measures are standardized within sample. The data is the HFCS multiple-imputation data, using representative weights. The number of observations is the maximum N across the five imputations. The pseudo R^2 is the average across the five imputations. The dependent variable is an indicator for owning the household main residence. Experienced Inflation is the weighted average of inflation over the household head's lifetime, with linearly declining weights from the year before the survey to birth year. High-inflation countries (where any household had a lifetime experienced inflation above 10%) include Croatia, Estonia, Finland, Greece, Hungary, Latvia, Lithuania, Poland, and Slovenia. In column (4), experienced inflation is winsorized at 10% and we include an indicator for having experienced inflation about 10%. In column (5), each year's experienced inflation is winsorized at 25% prior to averaging and we include an indicator for ever living through a year of inflation above 25%. Volatility is measured as the standard deviation of annual experienced inflation over the lifetime so far. Predicted inflation is predicted from experienced inflation using an AR(1) model. 5-year forecast calculated by iterating estimated AR(1) model forward, fixing coefficients as estimated in the survey. All regressions include demographic controls (age, gender, marital status, children, education, employment status, and deciles of wealth and income) and a fixed effect for survey wave.

B.5 Alternative Measures of Household Wealth

In our main analyses, we control for the decile of total household net wealth at the time of the survey. One concern with including wealth as an independent variable is that wealth may be endogenous if owning a home acts as a means of forced savings or asset accumulation.

In column (1) of Appendix-Table B4 we try to address this endogeneity by removing home equity from net wealth. We calculate a homeowner's current home equity as the current value of their main residence minus current mortgages with household main residence as collateral. Experienced inflation continues to predict higher odds of homeownership, at statistically significant levels. The explanatory power of this model over the baseline treatment of wealth is significantly lower, with a Pseudo R^2 of 0.22 compared to 0.51 in our baseline model, Table 2.5, column (1).

One concern with this analysis is that we might be over-correcting. With this definition of wealth, a household suffers a large drop in wealth immediately after purchasing a home, when instead we should view those households as having the same wealth. As a way to try to improve upon the measure of wealth, we use the current value of the household's main residence and its value at the time of purchase to calculate a real gain from homeownership due to house-price appreciation. We then subtract this gain from wealth to calculate wealth net the gain from owning the main residence. We can only calculate this measure for a subset of households who, if owners, reported the purchase price of their home, so the sample size in column (2) is substantially smaller. Using this alternative definition of wealth, the effect of experienced inflation remains large and statistically significant.

Measuring wealth net of the increase in home price is not ideal for several reasons. First, this is a noisy measure as we can at most observe the increase in the price of the current home and not any previously owned property. Inertial effects in homeownership are likely to be problematic if the household currently owns a home, they may be more likely to have owned a home in the past. Another problem with this variable is that it does not account for additional investment into the home. If the value of the home increases because the homeowner invested in adding a second floor, we would be subtracting more than just asset accumulation from being a homeowner. An additional concern is that for homeowners, this measure does not represent their counterfactual choice had they not purchased their home. For example, if a household purchased their home 20 years ago, we subtract 20 years of price increases but, presumably, the household would have invested their home equity elsewhere and would have received a return on their investment. For these reasons, we leave this as a robustness exercise.

Table B4: Inflation Experiences and Homeownership (Alternative Wealth Measures)

Dependent Variable: Own Main Residence	(1)	(2)	(3)	(4)	(5)
Log Experienced Inflation	2.80***	2.45***	2.71***	2.88***	2.35***
	(0.05)	(0.05)	(0.06)	(0.08)	(0.04)
Wealth and Income Deciles	Wealth net home equity	Wealth net HMR gain	Nominal	PPP-adj	Within-country
Demographic Controls	Yes	Yes	Yes	Yes	Yes
Observations	214,133	165,822	219,827	192,703	214,133
Countries	22	22	22	19	22
Pseudo R^2	0.217	0.410	0.513	0.516	0.474

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports exponentiated coefficients (odds ratios) from logit regressions with robust standard errors in parentheses. The data is the HFCS multiple-imputation data. With imputed data, the number of observations is the maximum N across the five imputations, and the pseudo R^2 is the average across the five imputations. Observations are weighted using the HFCS representative weights. The dependent variable is an indicator for owning the household main residence. Log experienced inflation is the log of weighted average of inflation over the household head's lifetime, with linearly declining weights from the year before the survey to birth year. Demographic controls include age, gender, marital status, children, education, employment, and deciles of wealth and income. All regressions also include a fixed effect for survey wave. In column 1, wealth is calculated as net home equity for owners with available price data. Column 2 excludes homeowners who do not report the purchase price of their home and uses Wealth net HMR gain, i. e., net wealth minus the gain from price appreciation of a homeowner's current home. Column 3 controls for nominal wealth and income. Column 4 adjusts wealth and income for purchasing power parity (limited to Euro Area countries). In Column 5, wealth and income deciles are defined within-country.

B.6 Accounting for Persistence in Homeownership Using SHARE Data

Our main analysis tests the hypothesis that macroeconomic experiences predict homeownership *at the time of the survey*. However, homeownership is persistent and therefore the relevant experience measure may be experiences at the point of first home-ownership. With retrospective data from the SHARE, we are able to zoom in on the first home purchase and ask whether macroeconomic experiences throughout life predict if and when an individual first purchases a home.

The SHARE microdata consists of a panel following elderly individuals (above age 50) in countries across Europe, starting with the first wave in 2004 to the most recent wave in 2015. We use data collected primarily in 2008-2009 from the SHARELIFE wave of the study for 14 countries in Europe.[8] In this wave, study participants were asked retrospective questions about several major aspects of their life, such as family structure, employment status, and homeownership. The data allows us to construct a yearly panel for each individual from age 20 to the year of the survey with indicators for whether the individual was married, had children under the age of 18, was employed, whether they had established their own household, and tenure status.

We also calculate a measure of experienced inflation for each of these individual-year observations using the individual's country and age as described in Section 2.3. In addition to the data used in our main analyses, we also use historical inflation from Reinhart and Rogoff for Denmark, Sweden, and Switzerland. We obtain historical inflation for the Czech Republic from GFD and Michal (1960).[9]

We drop about 6% of individuals with incomplete homeownership histories or who never established their own household. The final sample includes 26,691 individuals in 17,959 households from 14 countries. Appendix-Table B5 displays the summary statistics.

Using this data, we estimate a Cox proportional hazard model, defining a failure as the first year in which the individual was a homeowner after establishing their own household. We allow for a flexible baseline hazard over age. The key independent variable is log experienced inflation, which we adjust for individual-years who have experienced the German

[8]This paper uses data from SHARE Wave 3 (DOI: 10.6103/SHARE.w3.700, see Börsch-Supan et al. (2013) for methodological details). The SHARE data collection has been funded by the European Commission through FP5 (QLK6-CT-2001-00360), FP6 (SHARE-I3: RII-CT-2006-062193, COMPARE: CIT5-CT-2005-028857, SHARELIFE: CIT4-CT-2006-028812), FP7 (SHARE-PREP: GA N211909, SHARE-LEAP: GA N227822, SHARE M4: GA N261982) and Horizon 2020 (SHARE-DEV3: GA N676536, SERISS: GA N654221) and by DG Employment, Social Affairs & Inclusion. Additional funding from the German Ministry of Education and Research, the Max Planck Society for the Advancement of Science, the U.S. National Institute on Aging (U01_AG09740-13S2, P01_AG005842, P01_AG08291, P30_AG12815, R21_AG025169, Y1-AG-4553-01, IAG_BSR06-11, OGHA_04-064, HHSN271201300071C) and from various national funding sources is gratefully acknowledged (see www.share-project.org).

[9]For Ireland and the Czech Republic, we are missing early-life inflation experiences data for 22 individuals born before 1922. For these individuals, we re-normalize the weights to sum to 1 over available inflation years.

hyperinflation (thus having an experienced inflation measure in the millions) or have a negative lifetime average. For each of these two groups, we include separate indicators and set log experienced inflation to 0.

In all analyses, we control for the year, gender, and several time-varying demographics: whether the individual is married, has a child under the age of 18, and is employed. The results are in Appendix-Table B6. In columns (1) and (2), we limit the analysis to the 65% of individuals with complete demographic data over the relevant time frame. In columns (3) and (4) we use all available data, filling covariates with 0 when missing and including indicators for missing demographics. In columns (2) and (4), we also add country fixed effects. Standard errors are clustered by household, as defined at the time of the survey.[10]

The estimated hazard ratios in columns (1) and (3) indicate that a one log-point increase in experienced inflation predicts an 7-13% increase in the hazard of becoming a homeowner. The results are robust to controlling for country fixed effects in columns (2) and (4), where a 1 log-point increase in experienced inflation predicts a 6-7% increase in the hazard of homeownership. Hence, we confirm a significant role of past exposure to inflation on the decision to become a first-time homeowner and its timing.

[10]Our main results are unweighted as it is not clear that the SHARE survey weights are appropriate for the retrospective data. The results are mixed if we instead use the calibrated cross-sectional individual weights.

Table B5: Summary of SHARE Data

Country	Homeownership			Experienced Inflation (%)			
	Ever Own	Average Age First Own	Ind. Obs.	Mean	Median	SD	Ind.-Year Obs.
Austria	69%	30.3	909	8.19	6.55	6.02	21,294
Belgium	86%	30.6	2,731	3.74	3.56	1.1	46,074
Czech Republic	63%	28.6	1,778	3.43	2.85	2.66	41,636
Denmark	89%	28.4	1,919	5.36	5.05	1.53	26,567
France	81%	33.7	2,254	9.54	8.09	4.77	47,308
Germany	65%	32.8	1,802	5.77	5.56	1.16	44,144
Greece	90%	31.5	2,935	26.74	13.22	29.11	46,101
Ireland	90%	29.9	792	5.86	5.22	2.43	11,635
Italy	78%	33.7	2,417	12.64	10.22	8.7	53,018
Netherlands	74%	31.1	2,135	4.6	4.43	0.87	46,181
Poland	69%	27.9	1,882	21.67	10.96	23.63	36,980
Spain	87%	32.2	2,122	8.31	7.84	1.84	37,825
Sweden	87%	31.4	1,781	5.24	4.86	1.37	30,814
Switzerland	65%	36.4	1,234	3.21	3.31	0.71	34,978
Total	79%	31.4	26,691	9.36	5.7	13.35	524,555

Notes: Summary statistics of microdata obtained from Wave 3 of the SHARE. Homeownership variables are on the individual level and describe the percent of individuals who ever own their home and the average age at first ownership for individuals who ever own. For summary statistics of experienced inflation, each observation is an individual-age, for ages 20 to the minimum of age of first ownership, age at survey year, and 80. Experienced inflation excluded for 3% of Germans who lived through the hyperinflation. Experienced inflation is the weighted average of inflation over the household head's lifetime, with linearly declining weights from year before the observation year to birth year.

Table B6: Inflation Experiences and First Year of Homeownership (SHARE)

	(1)	(2)	(3)	(4)
Log Experienced Inflation	1.07***	1.07***	1.13***	1.06***
	(0.01)	(0.02)	(0.01)	(0.02)
Experienced German Hyperinflation	1.33	1.89***	0.93	1.36
	(0.25)	(0.38)	(0.17)	(0.26)
Negative Experienced Inflation	0.94	0.92	1.02	1.21
	(0.16)	(0.16)	(0.17)	(0.21)
Male	0.97**	0.96**	1.04***	1.04***
	(0.01)	(0.02)	(0.01)	(0.01)
Married	11.26***	11.63***	9.51***	10.06***
	(0.37)	(0.38)	(0.30)	(0.32)
Has Child under 18	0.49***	0.49***	0.53***	0.54***
	(0.01)	(0.01)	(0.01)	(0.01)
Employed	1.59***	1.61***	1.10***	1.20***
	(0.04)	(0.05)	(0.03)	(0.03)
Sample	Complete Covariates		All Available Data	
Indicators for Missing Covariates			Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
Country Fixed Effects	No	Yes	No	Yes
Observations	237,291	237,291	522,200	522,200
Individuals	17,412	17,412	26,691	26,691
Countries	14	14	14	14
Pseudo R^2	0.040	0.043	0.028	0.032

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Hazard ratios estimated from Cox proportional hazards model with failure defined as the first year of homeownership after establishing own household. Standard errors are clustered at the household level. Data are unweighted individual responses from the SHARE Wave 3 retrospective survey. We include time-varying indicators for being married, having children under the age of 18, and being employed. Columns 1 and 2 include only individuals with complete demographic data from age 20 to the first year of homeownership or survey year if never a homeowner. In columns 3 and 4, demographic indicators are filled with 0's for approximately 50% of observations with at least one missing covariate. Log experienced inflation is the log of weighted average of inflation over the household head's lifetime, with linearly declining weights from year before the observation year to birth year. This variable is 0 for households who lived through the German hyperinflation and for those with negative experienced inflation, with corresponding indicators.