Simple Actions, Complex Habits: Lessons from Hospital Hand Hygiene

Alexandra Steiny Wellsjo*
University of California, Berkeley

January 27, 2022

Most Recent Version Available Here

Abstract

Across industries, routine tasks are often critical to performance but not consistently done. How should organizations encourage routine behaviors? The answer depends critically on the nature of habits. In an important health setting, I find that behavior is well-described by a model of automatic cue-based habit: habits are automatically triggered by cues, rather than a conscious process. Habits remain stable even as cognitive costs, like fatigue, vary. However, habits are cue-specific and can be disrupted over time and across settings. I document these features of habit in data that tracks whether 13,606 hospital healthcare workers wash their hands 123M times that they are expected to. Hand hygiene is critical to the prevention of healthcare-associated infections, which affect millions of patients and cost billions of dollars each year. Yet, workers wash their hands only half as often as guidelines suggest. The data reveals substantial heterogeneity in behavior consistent with automatic cue-based habit. Habitual washers fatigue 43% less than people who wash more consciously, but respond just as much to habit disruptions. Consistent with location-specific cue-based habit formation, healthcare workers are more likely to wash in rooms they have visited more often in the past. A better understanding of habit can provide guidance for motivating routine behaviors.

*Email: alexsteiny@berkeley.edu.
I would like to thank SwipeSense for providing access to their electronic monitoring data. Many have provided thoughtful feedback and helpful comments on this paper. I would especially like to thank Ulrike Malmendier, Stefano DellaVigna, Jonathan Kolstad, Ned Augenblick, David Birke, Colin Camerer, Kaveh Danesh, Paul Gertler, Peter Jones, Jonathan Holmes, David Laibson, Isabel Macdonald, Maxim Massenkoff, Ricardo Perez-Truglia, Andrei Shleifer, Avner Strulov-Shlain, Dmitry Taubinsky, Justin Sydnor, and participants in the UC Berkeley psychology and economics seminar, UC Berkeley finance seminar, the Behavioral Science and Health Symposium, and the Camerer Group.

1 Introduction

People rely on habits in many situations, when they are not making deliberate choices but instead acting on an automatic response. Studies estimate that 45% of daily activities are routine, repeated at the same time and place each day (Wood et al., 2002; Quinn and Wood, 2005). Within organizations, routine tasks can be critical to performance. In industries like manufacturing, aviation, and medicine, failure to follow safety routines can have fatal consequences.[1]

How should organizations encourage routine behaviors? By nature, these repeated actions are often difficult or costly to monitor. Standard agency theory tells us how to design contracts to motivate behavior in these settings.[2] However, organizations may avoid such contracts over concerns that extrinsic incentives crowd out intrinsic motivation (Gneezy and Rustichini, 2000; Benabou and Tirole, 2003), reduce effort on other tasks (Holmstrom and Milgrom, 1991), or make salient the cost-benefit trade-off (Von Thadden and Zhao, 2012).

If employees can develop habits for routine behaviors, organizations have an additional set of tools to motivate them. For example, if temporary incentives can build habit, changes should persist even after incentives are removed. There are some successful examples (e.g., Charness and Gneezy, 2009), but many short-run interventions have minimal long-run impacts (e.g., Acland and Levy, 2015; Carrera et al., 2018; Royer et al., 2015). A better understanding of how habits work can provide insights for alternative strategies to promote habit formation.

In this paper, I provide empirical evidence for two key features of habit: automaticity and cue-dependence. These features are often overlooked in the economics literature but central to the psychology definition of automatic cue-based habit (e.g., Verplanken, 2018; Wood and Neal, 2007; Orbell and Verplanken, 2010). Over time, repeated pairing of a cue with a behavior builds a habit; when faced with the cue in the future, the habitual behavior is executed automatically without the use of mental resources.[3] The automaticity of habit offers organizations a win-win solution: workers can achieve higher output without exerting more effort or crowding out other tasks. While temporary incentives may build habit, cue-management strategies may build habit more efficiently and avoid the potential drawbacks of traditional incentives.

[1] Consider for example, an assembly line worker operating heavy machinery in a routine production process, a pilot who does a preflight check before each flight, or a surgeon who performs the same surgery each day.

[2] See Gibbons and Roberts (2012) for an overview of the implications of principal-agent models for incentives in organizations.

[3] For examples of a cue-behavior pairs, consider walking into your kitchen in the morning (cue) and making coffee (behavior), or sitting down at your desk (cue) and checking your emails (behavior).

I study habits using data on whether 13,606 hospital healthcare workers wash their hands in 123M instances that they are supposed to.[4] Hand hygiene is a textbook example of a habitual behavior, in which hand washing can develop as an automatic response to a cue, such as walking into a room (Potthoff et al., 2018).

This routine behavior is highly important; proper hand hygiene by healthcare providers and hospital staff is one of the most powerful tools for reducing healthcare-associated infections (HAI) (Jarvis, 1994). HAI are a critical public health issue; in the U.S. alone, 1 in 25 hospital patients is affected at any point in time, resulting in approximately 100k deaths and $35 billion in hospital costs each year (ODPHP, 2019; CDC, 2009; Scott, 2009) and recent evidence suggests the COVID-19 pandemic has exacerbated the problem (Baker et al., 2021; Weiner-Lastinger et al., 2021). In addition to the medical benefits and direct costs of HAI, hospitals also have large financial and reputational incentives to keep HAI low and compliance with hand hygiene guidelines high.

Despite efforts by hospitals and public health organizations to achieve high compliance, studies find that healthcare workers wash their hands less than half the time they should (WHO, 2009). Almost 200 years since the benefits of hand hygiene were first discovered by Semmelweis (1861), behavior change remains elusive. Many interventions are ineffective or have short-lived results. Those that succeed 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 the demand for better hand hygiene, companies have developed technologies to electronically monitor workers' compliance with hospital guidelines to wash before and after each patient interaction. I use data from one such company, SwipeSense, to study behavior of healthcare workers across 32 hospitals. This data 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 a typical 12-hour shift, workers are supposed to wash almost 70 times. Across all situations and for all workers, the hospital guidelines are clear and consistent.[5] This standard allows for performance comparisons across people and situations in a way that is impossible in many other domains. This high-frequency measure provides the statistical power to precisely estimate not only the main effects of the factors that predict compliance but also heterogeneity in the effects, which are critical to understanding habit.

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

[5]The nature of these interactions may vary substantially across patients, workers, and situations. For example, a visit may include a nurse administering medication, a technician taking vital measurements, or a surgeon providing a post-operative update. Regardless, the guidelines are the same.

Second, the fact that hand washing is required immediately before and after each patient interaction allows for a tight link between the context (or cue) and whether the action was performed. This is more difficult in other settings, such as 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 the 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 they 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 healthcare workers in this paper are hand hygiene experts. Many have had formal training in nursing or medical school. With an average of 5.7 opportunities per hour, all have had extensive on-the-job experience. This differentiates the setting from others in which habit formation may occur in parallel with real learning about the costs and benefits of taking a new action or consuming a new good.[6]

This detailed data reveals significant heterogeneity in behavior that matches the predictions of a dual-system model of habit and attention based on the psychology descriptions of automatic cue-based habit. In the model, cues trigger an automatic and costless response (habit, System 1) that requires a costly, conscious process to override (attention, System 2) (Quinn et al., 2010).[7] People who do not have good habits rely on the conscious process and have to pay a lot of attention to remember to wash (System 2); their behavior is highly affected by changes in the costs of attention, like fatigue. Instead, people who wash habitually (System 1) don't need to pay attention and are less affected by attention costs. While automaticity buffers these habitual washers from the negative impacts of fatigue, their behavior can be disrupted by changes to their routines, such as taking time off work or switching departments. This model predicts a specific pattern of heterogeneity, with people responding more to changes in the system they rely on more heavily.

Like other models of habit, automatic cue-based habits deliver persistence in behavior. Unlike some other models of habit, persistence is only predicted for behavior repeated in the same context.[8] This implies a cue-dependent habit formation process; within a person, we

[6]See Volpp and Loewenstein (2020) for a discussion of this and other mechanisms often confused with habit.

[7]The general System 1 and System 2 framework is described in Kahneman (2011).

[8]This resulting context-specificity in behavior is similar to other economic models that incorporate cues; for example, cue-based consumption (Laibson, 2001), decision processes (Bernheim and Rangel, 2004), attention (Taubinsky, 2014), and memory (Bordalo et al., 2020, 2021; Enke et al., 2020; Wachter and Kahana, 2019), as well neuroeconomic models of habit (Camerer et al., 2018). The literature has not converged on a definition of context, or cues, and it likely varies across individuals and settings. I focus on a common candidate, the physical environment, though context could also be mental states, time, or norms (see, for example, Malmendier and Wachter, 2021, for a discussion of context in models of memory).

should observe many small habits forming with context-specific experience.

I document five empirical facts consistent with the predictions of the model.

Facts 1 and 2 document the aggregate effects of shocks to each system. Fact 1 demonstrates that increasing attention costs, specifically through fatigue, lowers compliance with hand hygiene guidelines. Similar to Dai et al. (2015), compliance falls by 5.9pp (s.e. 0.08), or about 9%, as workers fatigue over a 12-hour shift. Compliance is also lower at other times when fatigue is likely high; compliance is 1.4pp (s.e. 0.19) lower on night shifts (vs. day) and 0.7pp (s.e. 0.02) lower on consecutively-worked shifts with no day off in between. Fact 2 documents the negative impacts of two disruptions to habit. First, compliance is 1.4pp (s.e. 0.14) lower upon returning from a long break from work, consistent with habits that decay over time if not activated. Second, compliance is 3.0pp (s.e. 0.19) lower when workers cover a shift in a different department (a float day), consistent with habits triggered by location-specific cues.

These estimates are economically significant and precisely estimated. Even a 1pp change in compliance is associated with a 2% reduction in HAI. With the large amount of data, the statistical precision allows me to estimate heterogeneity in the response to shocks, which I document in Facts 3 and 4. Using compliance levels as a proxy for habit, I estimate heterogeneity in the effects across people with different levels of habit.

Fact 3 documents heterogeneity in fatigue that is consistent with the automaticity of habit. If the compliance of low compliers is bounded below by floor effects, we may expect the highest compliers to fatigue most throughout a shift. Instead, high compliers fatigue 43% less than low compliers. Consistent with automaticity of habit, high compliers are also less responsive to other sources of fatigue like working nights and consecutively-worked shifts.

Fact 4 shows a different pattern of heterogeneity in disruptions to habit. While high compliers benefit from automaticity in their usual settings, their behavior can be shocked by disruptions to their routines. Indeed, high compliers respond at least as much as low compliers to long breaks from work and department switches.

Fact 5 uses within-person heterogeneity over time and across locations to document a highly location-specific cue-based habit formation process. Healthcare workers are more likely to wash in rooms that they have visited more often in the past. These effects are significant; a counterfactual exercise estimates that compliance would be 2.7pp higher if an individual's work was concentrated in a single room.

This study differs from much of the existing literature in economics on habit formation in the field by focusing on the features of automaticity and cue-dependence 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)). A recent exception is Buyalskaya et al. (2021), who identify habit formation as the predictability of behavior from cues using field data on hand hygiene behavior and gym attendance. The authors use machine learning to identify an individual's time to habit formation as an increase in predictability over time. The paper takes a broad perspective on cues, interpreting predictability from any context variable as a cue-based habit (e.g., a habit to wash at the start of the shift, but not the end). In this paper, I apply a more narrow interpretation of context to distinguish features that drive automatic habit, affect costly attention, or change the benefits of hand hygiene.

This paper is not the first to suggest solutions for organizations who attempt to improve performance in routine behaviors. A literature on organizational routines highlights the benefits of coordination and standardization in processes and responses (Becker, 2004). This paper is more closely related to those that focus on management strategies for optimizing performance on routine tasks. For example, models of knowledge-hierarchies suggest a management-by-exception approach, wherein production workers do routine tasks and escalate complex problems to managers (Garicano, 2000). As a more concrete example, the use of checklists in routine tasks have been shown to improve surgical outcomes and reduce airplane accidents (Haynes et al., 2009; Guwande, 2011). While the benefits of automaticity in routine behaviors is not a new concept in economics (for an early example, see Weiss and Ilgen, 1985), there has been a lack of empirical evidence documenting its importance in the workplace. This paper contributes to the literature by improving our understanding of habits and their role in routine behaviors with implications for motivating behavior change.

The paper proceeds as follows. Section 2 describes a simple model of automatic cue-based habit. Section 3 describes the electronic monitoring data. Section 4 provides empirical evidence consistent with the predictions in Section 2. Section 5 discusses the implications for organizations and concludes.

2 Theoretical Framework

In this section, I describe a simple theoretical framework of attention and habits in hand hygiene.

Drawing from the psychology literature, habit formation involves the creation of an automatic cue-response. Habits are formed through repeated pairing of a cue with a behavior. When habits have formed, the behavior is performed automatically in response to the cue. This habitual behavior is an unconscious response and distinct from the cognitive processes in which we make deliberate choices. In the Kahneman (2011) framework, habits are the

System 1 process, an automatic and unconscious response to a cue. Overriding habit requires activating System 2 to make a conscious choice. This dual-system perspective is consistent with evidence from neuroscience that as participants practice a task, brain activity moves away from areas of executive control (Tricomi et al., 2009; Kelly and Garavan, 2005).

Activating the System 2 process to override habit is costly. The choice to activate it will depend on incentives like the benefits of behavior change and competing drains on mental resources.

Habits, on the other hand, are free. As people develop good habits, they benefit from taking the desired behavior automatically without draining mental resources. This is consistent with psychology evidence from dual-task experiments which find that performance on one task improves as the second becomes habitual (e.g., Brown and Carr, 1989).

This feature of habit leads to a hallmark prediction from the psychology of habits: reward insensitivity. While attentive choices will respond to changes in incentives, habitual behavior is an automatic response to the environment that is unaffected by changes in current incentives. In one of the few experimental examples, Neal et al. (2011) find that moviegoers with a habit of eating popcorn in the cinema ate the same amount even if it was stale, while those without a habit did not. In evidence from neuroscience, Tricomi et al. (2009) find that with more practice on a task, participants become less responsive to changes in the rewards.

Instead of responding to changes in incentives, habitual behavior is a function of the current context and past behavior in similar settings. Correlational studies in psychology find that, for people with strong habits, behavior is better predicted by past behavior in the same context than by stated intentions (e.g., Ji and Wood, 2007; Danner et al., 2008). Because of this context-specificity, habits are affected by changes in the environment. Studies find, for example, that new movers can change behavior even if they had strong habits prior to their move (e.g., Wood et al., 2005; Verplanken and Roy, 2016).

This model makes predictions about heterogeneity in behavior which I test empirically in Section 4.

2.1 Benchmark without habits

To set-up the framework, I begin with a model without habits.

Consider a healthcare worker who should wash in an opportunity o, characterized by time t and location l. Washing in each opportunity has a fixed benefit $b_i > 0$, which may vary across individuals. This heterogeneity captures differences across individuals in the perceived health benefits of washing and other intrinsic motivations like altruism for patients.

Washing requires paying attention a. I assume $a \in [0, 1]$ is the probability of paying attention to washing in a given opportunity, which translates in this benchmark model directly to the probability of washing, $p(a) = a$.

Attention is costly, with convex cost of attention $c(a)$. The attention costs in this model can be thought of as a simplification of a complex constrained maximization problem in which the worker allocates finite cognitive resources to all of the tasks that she must complete. Thus, costs in this model can be thought of as the opportunity cost of paying attention to hand washing instead of another task.

I assume a quadratic attention cost, with $c(a) = a^2/2$. I use a quadratic cost function as it has the simplifying feature that the convexity is constant across values of a (i.e., a zero third derivative). In Appendix A, I show that the predictions of this model hold for any convex cost function that is not increasing in convexity too much.

Attention costs are scaled by the other drains on attention at the time of the interaction, $d_t > 0$. Thus, paying attention is more costly when there are distractions or competing drains on attention.

Attention, a, is the choice variable in this problem, selected ex-ante to maximize the individual's expected utility

$$\begin{aligned} E[U(a)] &= b_i p(a) - d_t c(a) \\ &= b_i a - d_t \frac{a^2}{2} \end{aligned}$$

which yields an optimal attention level $a^*(b_i, d_t) = \min(b_i/d_t, 1)$.

Attention, and therefore compliance ($p^*(b_i, d_t) = a^*(b_i, d_t)$), are increasing with the individual's perceived benefits of hand hygiene, b_i, and decreasing with drains on attention, d_t.

Heterogeneity in benefits, b, induces not only changes in the levels of compliance across individuals, but also the response to drains on attention. Intuitively, people with high b_i are operating at an attention level with a higher marginal cost. Because drains on attention scale the marginal cost, this will have a larger impact on individuals with high b_i.

Formally, away from the boundary conditions

$$\frac{\partial^2 a^*(b_i, d_t)}{\partial b_i \partial d_t} = -\frac{1}{d_t^2} < 0$$

and so the negative effects of drains on attention are larger (more negative) for people with higher b_i.

2.2 Automatic cue-based habit

The benchmark model can be extended to include habits in several ways. Drawing on the psychology literature, in this section I describe a model in which habits are an automatic response to a cue.

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 are 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. Combining these features, I assume habit evolves automatically according to

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

with $\gamma \in [0,1)$ and $h_0 \in [0,1]$. $w_{il't-1}$ 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 $t-1$.

Note that the choice of t for a cue-based habit is not obvious. I assume an exponential decay of habit over time, which implies that habits decay over periods with no hand hygiene opportunities. Alternatively, one could model opportunities as the unit of time, in which case habits would not decay over periods of no opportunities but only with missed opportunities. The reality may be somewhere in between with habits decaying over time and more severely with actively missed hand hygiene opportunities. For simplicity, I assume an exponential decay over time, which treats missed and no opportunities equivalently.

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 portability of habit across locations. If habits are not location-specific $\delta(l,l') = 1\ \forall l, l'$. If habits are location-specific, $\delta(l,l) = 1$ and $\delta(l,l') < 1\ \forall l \neq l'$. We may expect that locations with a higher degree of similarity in cues have a δ closer to 1. With $\gamma < 1$, habits (and the importance of past washes) decay over time.

I model automatic habits as costless hand washing; habits allow the worker to wash without paying attention. As before, a translates to a probability of paying attention. With full attention $a = 1$, the worker will always wash. With no attention $a = 0$, the worker washes according to her habit level h_{ilt}. At attention level a, the probability of washing is given by

$$p(a, h_{ilt}) = a + (1-a)h_{ilt}. \tag{2}$$

Intuitively, the worker can comply at her habit level at not cost, but must pay costly attention if she wants to wash more. While the worker knows her habit level h_{ilt}, or likelihood of washing out of habit, identifying whether she actually washed out of habit in a particular opportunity requires paying attention. This representation is consistent with the self-report measures that psychologists use to identify habit: *[The behavior is something...] "I do automatically", "I do without having to consciously remember", "I do without thinking",* and *"I start doing before I realize I'm doing it"*.[9]

I assume the worker is myopic when selecting her attention level and does not anticipate the impact of today's behavior on future habit levels.[10]

The worker maximizes expected utility

$$\begin{aligned} E[U(a, h_{ilt})] &= b_i p(a, h_{ilt}) - d_t c(a) \\ &= b_i(a + (1-a)h_{ilt}) - d_t \frac{a^2}{2}. \end{aligned}$$

The optimal solution, a^*, satisfies the first order condition with

$$a^*(b_i, d_t, h_{ilt}) = \min\left(\frac{b_i}{d_t}(1 - h_{ilt}), 1\right). \tag{3}$$

which delivers an optimal probability of compliance

$$p^*(b_i, d_t, h_{ilt}) = p(a^*, h_{ilt}) = \min\left(\frac{b_i}{d_t}(1 - h_{ilt})^2 + h_{ilt}, 1\right). \tag{4}$$

This yields the first set of predictions.

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

Proof. This follows directly from the derivative of optimal compliance with respect to distractions as

$$\frac{\partial p^*(b_i, d_t, h_{ilt})}{\partial d_t} = -\frac{b_i}{d_t^2}(1 - h_{ilt})^2 \leq 0.$$

[9]Statements from the Self-Report Habit Index (Verplanken and Orbell, 2003) and condensed Self-Report Behavioral Automaticity Index (Gardner et al., 2012).

[10]Models of rational addiction based on (Becker and Murphy, 1988) instead assume that agents anticipate their future habit formation when making choices today. In the domain of hand washing, Hussam et al. (2017) find evidence for the hallmark prediction of rational addiction, that anticipated future changes in benefits impact behavior today. On the other hand, evidence on digital addiction suggests that people behave myopically with respect to habit formation (Allcott et al., 2020). I suspect that the key insights from my model are largely orthogonal to whether agents anticipate future habit formation and thus, for simplicity, consider only the myopic case.

$\square$

To test this prediction empirically, I focus on fatigue as a proxy for the drain on attention, d. I use three proxies for fatigue in the data: time at work (i.e., hours on the shift), performance on the night shift (vs. the day shift), and consecutively-worked shifts with no day off in between.

The impact of habits on compliance is more nuanced. Habits affect compliance through two opposing channels. First, habits directly increase compliance through the $p(a, h_{ilt})$ function; fixing a away from the boundary conditions, an increase in habit of Δh increases compliance by $(1-a)\Delta h$. Second, there is an indirect effect of habits on compliance through a decrease in the optimal attention level. Attention is only relevant in the $1-h_{ilt}$ chance that the worker doesn't wash out of habit, resulting in a marginal benefit of attention $b_i(1-h_{ilt})$. As habit increases, this marginal benefit declines, putting downward pressure on attention.[11]

In the results that follow, I focus on the region of the parameter space where the direct effect of habits on compliance outweighs the indirect effect, as formalized in Prediction 2.

Prediction 2. *Compliance is increasing in habit, h_{ilt} when the direct effect on compliance outweighs the indirect effect through attention.*

Proof. Compliance is increasing in habit when

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

This condition holds in any region of the parameter space in which the benefits of hand hygiene are not too high, so that automatic habits make up a significant portion of washing. This is true in any region where $a^* \leq 0.5$. With very high benefits ($b_i \geq d_t$), compliance will be 100% regardless of habit and the condition holds vacuously. With $b_i \leq 2d_t$, attention is never more than 50%, so the condition holds for any habit level.

The only parameter space under which this condition does not hold is with high benefits that don't quite hit the boundary case, $d_t > b_i > 2d_t$. Compliance would be declining as habits increase from 0 to some threshold level of habit $\bar{h} \leq 0.5$, after which point compliance will rise with habits.

[11] Optimal attention does not decline by the full change in marginal benefit b, as convex attention costs also imply a lower marginal cost at the lower a.

The direct effect of habits outweighs the indirect effect when $b_i \leq 2d_t$, $b_i \geq d_t$, or when $d_t > b_i > 2d_t$ and habits are above the threshold level, $\bar{h}$. $\square$

In the data, I cannot directly observe habit levels. However, I can proxy for changes in habit within an individual using the two dimensions along which habit varies: location l and time t.

At the same time t, a worker can have high habit in one location l_1 and low habit in another location l_2. Prediction 2 implies that holding all else constant, moving from location l_1 to l_2 will lower compliance. In the empirical analysis, I consider exactly this type of event. Specifically, I examine changes in compliance on float days – days when a worker covers one or two shifts in a different department. If habits are highest in a worker's home department, float days should have lower compliance than regular shifts.

According to the evolution of habit in Equation (1), time between washing should lower habit. Consider, for example, habit h_{ilt} in location l at time t. If no hand hygiene is performed between time t and $T > t$, or $\{w_{il't}, ..., w_{il'T-1}\} = 0$, habit will fall with $h_{ilT} = \gamma^{T-t} h_{ilt} < h_{ilt}$. To test this empirically, I look at the impact of time off work on compliance. Time off work has two opposing effects. As documented by Dai et al. (2015), time off work can help workers recharge, lowering fatigue ($\downarrow d_t$). On the other hand, the longer the time off, the larger the corresponding decline in habit. In the data, I consider long breaks from work (at least two weeks off) as a proxy for a drop in habits.

Note that each of these events – float days and returns from long breaks – are unusual, likely paired with other distractions. Float days carry additional complications of trying to navigate in a new department like looking for supplies or figuring out who to ask for a consult. Returning from a vacation may involve more socializing with coworkers.

In addition to the within-person variation in habit, I also examine differences in habit across people. Heterogeneity in habit can come from several sources. First, job design and exogenous shocks can lead to heterogeneity in habit levels without any heterogeneity in the model parameters. Individuals' whose work is concentrated in more similar environments, who have more opportunities, and less time between shifts or patient interactions will experience less habit decay and benefit from more portable habits. Even with the same likelihood of paying attention, chance will cause some to realize a more positive sequence of washes than others. In the model, these sources of heterogeneity are captured by different sequences of $\{w_{ilt}\}$.

Second, individuals may differ in how easy it is to build and maintain a habit. For example, individuals with who have a smaller decay in habit (a γ_i close to 1) or more portability of habits across domains (a $\delta_i(l, l')$ close to 1) will have higher habit even with the same past washing experience. Third, heterogeneity in habit may come from variation in benefits, b_i, or in past drains on attention, $\{d_t\}$.

To simplify, in the remainder of this section I focus on heterogeneity driven by their past washing behavior, $\{w_{ilt}\}$, which vary by chance or features of the job that are exogenous to the other model parameters.

Compliance is increasing in the habit level, but the attention allocated to hand washing is declining in habit. To show this graphically, I plot the model solution in Figure 1 for a given b and d. The black line at the top of the area chart plots compliance p^* at the optimal attention level for each habit level, h_{ilt} (plotted along the x-axis). The dark region represents the share of compliance coming from effortful attention a^* while the light region is the 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.

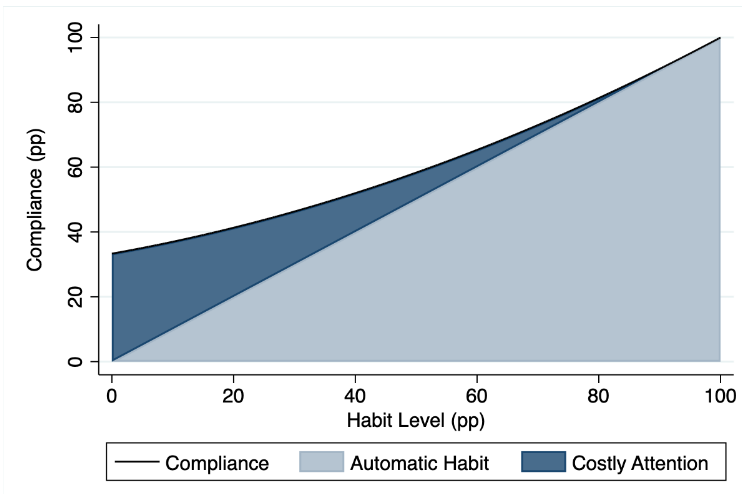

Figure 1. Model compliance and attention

Note: Plot shows model solution for compliance and attention by habit level with benefit $b = 1$ and drains on attention $d = 3$. Share of compliance coming from habit (h) is plotted in light blue with the share coming from attentive hand hygiene ($a^*(1 - h)$) in dark blue.

Distractions impact costly attention, but have no impact on habitual compliance. In Figure 1, 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 3 formalizes this idea.

Prediction 3. *Distractions ($\uparrow d_t$) will have a larger impact on compliance for people with low habit.*

Proof. Distractions have a negative effect on compliance. The prediction is true as $\frac{\partial p^*(b,d_t,h_{ilt})}{\partial d_t}$

is increasing (i.e., becoming less negative) in the habit level, with

$$\frac{\partial^2 p^*(b, d_t, h_{ilt})}{\partial d \partial h_{ilt}} = \frac{2b(1 - h_{ilt})}{d_t^2} \geq 0.$$

$\square$

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 4. *When compliance is increasing in habit, disruptions to habit ($\downarrow h_{ilt}$) will have a larger impact on compliance for people with high habit.*

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

$$\frac{\partial^2 p^*(b, d_t, h_{ilt})}{\partial h_{ilt}^2} = \frac{2b}{d_t} \geq 0.$$

$\square$

Figure 2 presents Predictions 3 and 4 graphically. From the baseline solution in Figure 1, I plot the change in optimal compliance for an increase in distraction (Figure 2a) or a decrease in habit level (Figure 2b) from different baseline compliance levels. The model predicts that 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. Figure 2c shows one example of a combination of a shock to habit paired with additional distractions, as we may expect to be the case in the empirical examples of float days and long vacations. Depending on the relative size of the shocks to habit and distractions, heterogeneity in the net result is ambiguous. Figure 2c shows an example where the heterogeneity cancels out, leaving a uniform response.

Predictions 3 and 4 provide testable implications for the heterogeneity in response to shocks by compliance level, driven by the key feature that habit and attention act as substitutes.

These predictions differ from those of the standard model with heterogeneity across individuals in the perceived benefits. In contrast to automatic cue-based habits, heterogeneity in b implies that high compliers should respond more to drains on attention. In a model without habit, vacations and switching departments may just reflect other drains on attention, like catching up with coworkers or needing to learn new tasks. In this case, the standard model predicts that the high types will also respond more negatively to disruptions.

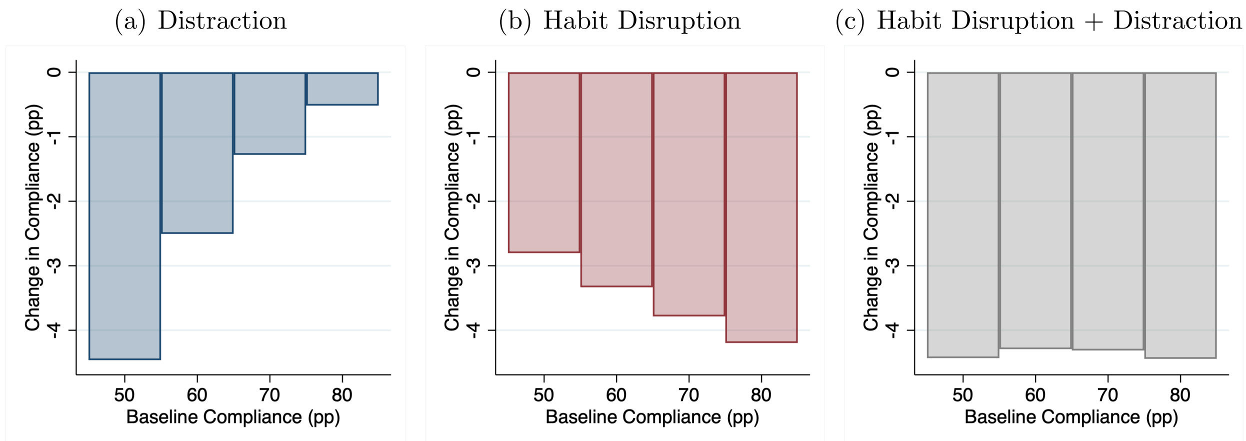

Figure 2. Model response to distraction and habit disruption

Note: Model response to an increase in distraction (a) or habit disruption (b) or both (c) in the cue-based habit model. Increase in distraction in increases d_t from 3 to 4.5 in (a) and from 3 to 3.35 in (c). Habit disruption modeled as a drop in habit by 5pp in both (b) and (c).

2.3 Habits as cravings

Rather than having a direct impact on behavior through automaticity, habits may instead alter preferences for hand hygiene. Through this channel, past hand washing increases the desire to wash in the future through cravings. This notion of past behavior increasing future cravings has been used as the motivating example in many descriptions of habit in rational addiction models (for an example with hand hygiene, see Hussam et al. (2017)).

Laibson (2001) provides one cue-based micro-foundation for this process; when cues have been paired with past consumption, the cue triggers a preparatory or compensatory process that raises marginal utility for consumption.[12] He gives examples such as a smoker feeling nicotine cravings upon seeing an open box of cigarettes or the smell of freshly baked bread initiating appetite arousal and salivation.

Laibson models this as an increase in the marginal utility of consumption in the presence of the cue. In this context, one could think of a craving for hand washing as a change in the benefits of hand washing (e.g., $b_{ilt} = b_i + \eta h_{ilt}$, with $\eta > 0$ a scaling parameter). This model will generate results similar to the benchmark model with heterogeneity in benefits. As people build habit, they want to wash more and will pay more attention to their hand hygiene behavior. Washing will be context-dependent and increase in places with more past washing experience. However, habit modeled this way will not deliver the same automaticity as the cue-based model; high habit types pay more attention to hand hygiene and thus will respond more to other drains on attention.

[12] Following Becker and Murphy (1988), Laibson (2001) models consumers as rational and forward-looking. As discussed above, for simplicity, I assume people are myopic with respect to their future habit formation.

In practice, past experiences may affect both automaticity and preferences. For example, with smoking, addiction may increase cravings for cigarettes while cue-based habit make the act of smoking automatic. The model highlights that these two channels are distinct, with different predictions about whether habits operate through deliberate or automatic processes. Effectively motivating behavior change through these distinct channels likely requires different approaches.

2.4 Links to other models

This model of automatic cue-based habit has many features similar to models of associative memory (for examples in economics, see Bordalo et al., 2020, 2021; Enke et al., 2020; Mullainathan, 2002; Wachter and Kahana, 2019). In fact, Verplanken (2018) defines habits as "memory-based propensities to respond automatically to specific cues, which are acquired by the repetition of cue-specific behaviors in stable contexts." One could model habits as retrieving a memory for context-specific default behavior. For example, when a nurse walks into a patient room, the context prompts her to remember what she does when entering a patient room. She will recall past instances of walking into patient rooms and will be more likely to recall events that happened in more similar contexts (e.g., the same room) and more recently. The more she retrieves memories of washing her hands in those past situations, the more likely she will be to wash. This process happens automatically when she walks into the room; deliberately choosing to wash requires paying attention to the outcome of this automatic process.

The neuroeconomic model of habit of Camerer et al. (2018) also has similar characteristics, but describes a distinct mechanism that centers around the reliability of rewards. People operate either in preference-mode, when they consider all options and make a deliberate choice, or in habit-mode, when they simply repeat their last action in the same context. The choice of which mode to operate in depends on the predicted value of each option and the reliability of those predictions, which are determined by past experiences. People act out of habit in familiar, predictable contexts and more deliberately in unfamiliar or uncertain settings.

Also motivated by neuroscience literature, Bernheim and Rangel (2004) model addiction with cues the trigger hot decision states. In hot states, the harmful good is always consumed, regardless of preferences in the cold state. I model costly attention as the tool to overcome such automatic behavior. In Bernheim and Rangel (2004), addicts can change their behavior through endogenous cue choice, deciding whether to put themselves at risk of cuing the hot state.

I focus on the role of cues as prompting an automatic habitual behavior, but cues can also trigger deliberate choices. For example, consider a poster of someone washing their hands that, upon being seen, reminds a nurse to wash her hands. Taubinsky (2014) models the impact of such attention-focusing cues which override habitual behavior to prompt an active choice.

3 Data

Hospitals in the data use a "gel-in, gel-out" policy, which requires hospital workers to wash or sanitize their hands upon 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 hand hygiene compliance. The technology uses three key components: RFID badges worn by workers, sensors on soap and sanitizer dispensers, and sensors in each patient room. Appendix Figure A1 shows the SwipeSense system. Badges are assigned to a unique worker, which allows for tracking of individuals. To 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 hospital), worker information (anonymized individual identifier, role, department), and compliance.

The dataset includes any worker at the hospitals who wears an RFID badge including physicians, registered nurses, nursing assistants, technicians, administrators, and environmental service workers, among others. 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 3, 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 13 hours, shown graphically as the region between the two red lines. These shifts make up 68% of shifts, but 90% of all interactions.

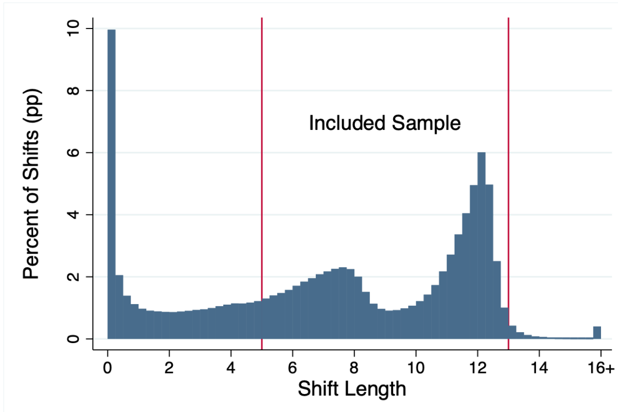

Figure 3. Distribution of shift length in the full sample

Note: Plot shows a histogram of shift length in hours across all shifts in the data. Length is capped at 16 hours for the figure. Shifts lasting 5 to 13 hours, as indicated by the red vertical lines, are included in the analysis sample.

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. 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 C summarizes the full data and the main sample. Compared to the full data, the main sample has more shifts per worker, longer shifts, more opportunities per shift (but fewer per hour). Despite differences in the shifts and workers, overall compliance is quite similar (61% in the main sample vs. 59% in the full data).

The main sample covers hand hygiene behavior of 13,606 workers in 123M opportunities across 2.1M shifts. These shifts take place in 276 departments, 32 hospitals, 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 during the pandemic in Appendix D. Though intriguing, potential changes in hospital procedures during this time make it difficult to draw conclusions without more information.

In Table 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 4 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.

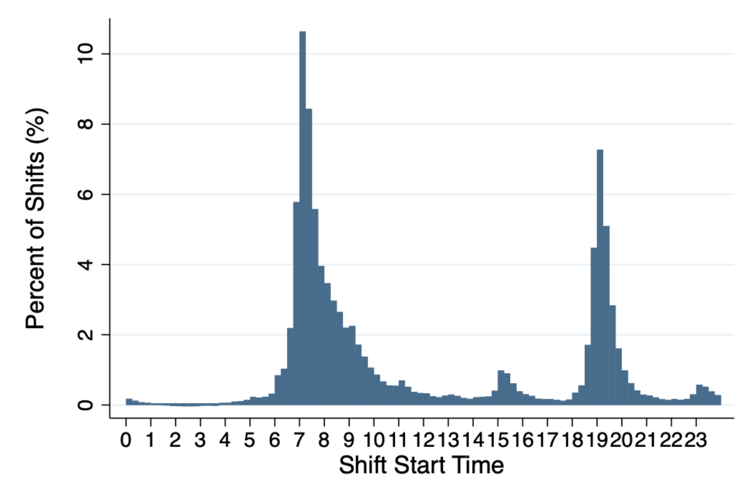

Figure 4. Distribution of shift start times

Note: Plots show histogram of shift start time across all shifts in the analysis sample.

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, weekdays have a slightly higher than proportional share of shifts: 77% of shifts are 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.[13] 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.[14] 55% of workers in the main sample work at least one float day.

The median shift is worked after only one day off. While workers are typically active in the data, I also observe long breaks, with 1.4% of shifts occurring after at least 2 weeks off. While rare events, they are distributed across workers, with 74% of the sample having at least one long break.

[13]85% of shifts are associated with a department.

[14]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).

Panel B of Table 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 5 graphically depicts the large amount of heterogeneity in the distribution of compliance across individuals. 17% of 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%.

The heterogeneity across individuals is not noise; these individual estimates are precise and stable. To quantify the precision, I use a split-sample technique to estimate correlations between individuals' average compliance estimated separately on two halves of the data. Splitting by days of the week (Mon/Wed/Fri/Sun vs. Tue/Thu/Sat), the within-person correlation is 0.99. The correlation is the same when splitting the data every other week. Even over time, individuals differences are quite stable. The within-person correlation between compliance in the first and second half of each worker's data is 0.84.

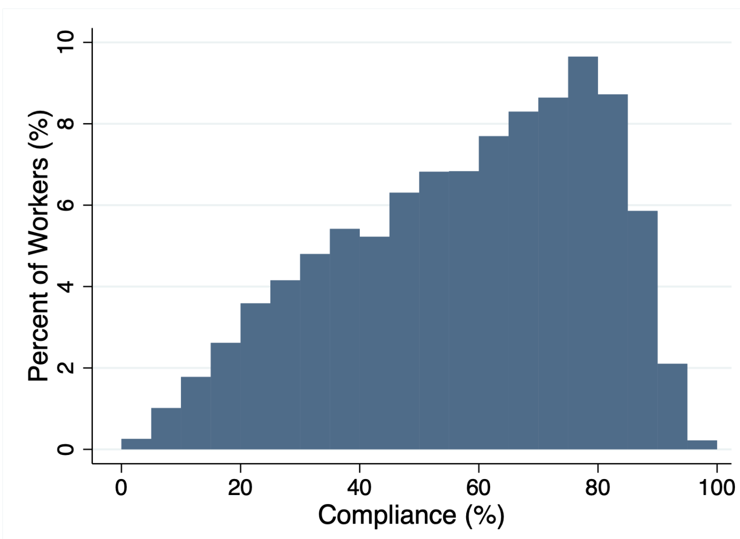

Figure 5. Heterogeneity in compliance

Note: Distribution of compliance rates across all individuals in the analysis sample.

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%).

4 Empirical Evidence for Cue-Based Habit

In this section, I provide evidence consistent with the predictions of the model of cue-based habits.

4.1 Fact 1: Compliance falls with fatigue

According to Prediction 1, other drains make paying attention to hand hygiene more costly. This will lower the optimal attention workers allocation to hand hygiene and, in turn, lower compliance rates. As described in Section 2, I test this prediction using three proxies for fatigue: hours at work, night shifts, and consecutively-worked shifts.

In Column 1 of Table 2, I report the results from an OLS regression of compliance on features of the opportunity, shift, and individual. Specifically, I estimate

$$c_{ilt} = \beta_F F_{it} + \beta_B B_{ilt} + \omega_i + \gamma_l + \eta_{lt} + \epsilon_{ilt} \tag{5}$$

where c_{ilt} is compliance by individual i at time t and location l. Compliance, c_{ilt}, is a binary outcome rescaled to 100. F_{it} is a vector of proxies for fatigue including 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 night and consecutively-worked shifts. B_{ilt} is a vector of potential proxies for benefits: the worker's total number of opportunities in the hour (normalized within person), an indicator for room entry (vs. exit), an indicator for weekend (vs. weekday) shifts, and a quadratic trend for the number of months the worker has been in the data to account for learning over time. I also include an indicator for the first shift each person has in the data, for which I cannot identify whether it was a consecutively-worked shift. I include individual fixed effects (ω_i) to account for individual differences, department fixed effects (γ_l) to control for variation in benefits across patient types (e.g., neonatal intensive care vs. emergency room), and network-month-year fixed effects (η_{lt}) 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). ϵ_{ilt} is an error term. To capture potential autocorrelation, I cluster standard errors by individual.[15]

Consistent with Prediction 1, all estimated coefficients on the proxies for fatigue are

[15]I estimate all regressions using the Stata command reghdfe (Correia, 2016).

negative. Figure 6 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 5.9pp (s.e. 0.08) from the beginning to the end of a 12-hour shift, or about 9% from the start of the shift. Compliance on the night shift is 1.4pp (s.e. 0.19) or 2% lower than on the day shift and 0.7pp (s.e. 0.02) or 1% lower on consecutively-worked shifts relative to those with at least one day off in between.

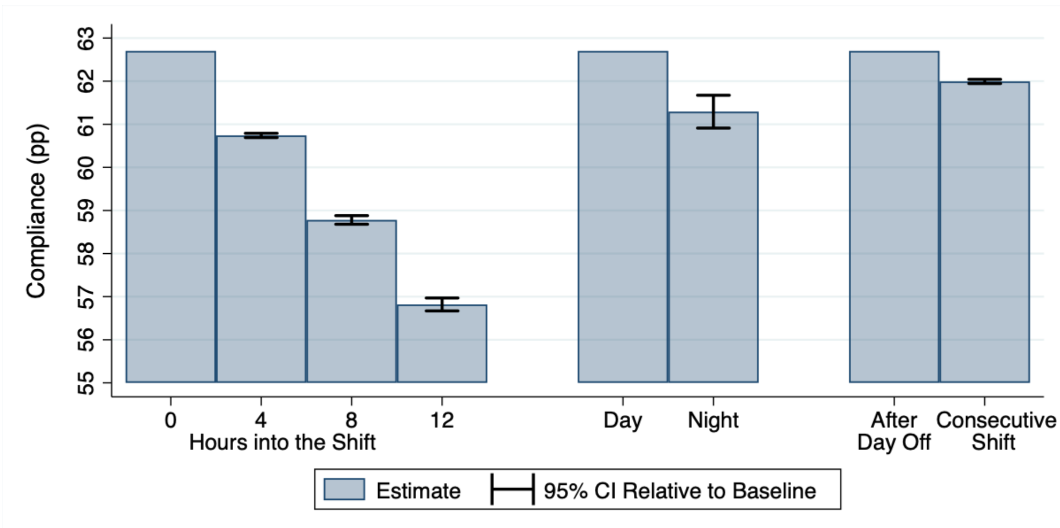

Figure 6. Estimated predictors of compliance

Note: Graph plots estimated coefficients from Column 1 of Table 2 relative to the baseline level of compliance. Error bars indicate 95% confidence interval on the difference relative to the omitted category. Y-axis re-scaled so that the baseline is average compliance across the first hour of non-consecutive day shifts (the omitted categories).

4.2 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 2, I consider two disruptions to habit: float days and time off work.

I begin by examining disruptions with an event study design. Specifically, I estimate Equation (5), adding a set of event time indicators for each event e

$$c_{ilt} = \sum_{e \in E} \sum_{j \in J_e} 1\{t = j\} + \beta_F F_{it} + \beta_B B_{ilt} + \omega_i + \gamma_l + \eta_{lt} + \epsilon_{ilt} \tag{6}$$

where the set J runs from 5 shifts prior to the event to 5 shifts after the event. I also include indicators for all shifts more than 5 shifts prior to and after the event.[16] For each event, the

[16] Following Sandler and Sandler (2014), to account for the fact that there may be multiple events, rather

reference category is 2 shifts prior to the event.

Beginning with float days, I estimate 6 where a set of events E includes float days that last either 1 or 2 shifts as well as "permanent" department switches, which I define as lasting 3 or more shifts.

Figure 7 plots the event time coefficients for department switches that last 1 or 2 days.[17] Despite the potential endogenous assignment of float days, the estimated event time coefficients reveal no 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 typically go back to their usual department.[18]

Turning to time off as a disruption, I use the same event study methodology to look at variation in compliance around breaks of different length. Specifically, the set of events E are breaks of 1 day, 2 days, 3 to 4 days, 5 to 7 days, 1-2 weeks, 2-4 weeks, and 4 or more weeks.

Figure 8 plots the event time coefficients. The first dashed line at -1 indicates the shift prior to the break and the second dashed line indicates two shifts after returning from a break. The key coefficients of interest are the shifts plotted between the two dashed lines, which are the first shift after returning from a break. These coefficients, all with event time of 0, are spread out to graphically illustrate the amount of time off work (or the number of days between -1 and 0).

As discussed in Section 2, we expect time off work to have two opposing 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 the two shifts prior, compliance is 0.55pp (s.e. 0.04) higher after one day off and 0.68pp (s.e. 0.04) higher after 2 days off. This positive effect of time off maxes out after a break of 3 to 4 days, when compliance upon return is 0.91pp (s.e. 0.04) higher.

Time off work can be restorative, however it can also have negative effects if habits decay over time. After 4 days off, there are no additional gains to longer breaks. Instead the trend appears to reverse, consistent with longer breaks resulting in more habit decay. For breaks that last 2-4 weeks, there is so significant difference in compliance relative to the reference

than binary indicators, the binned categories take on a value equal to the number of events that are more than +/-5 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 5 shifts after a float day would be equal to 2.

[17] Of the 39,151 switches that occur, 80% last only one shift, 10% last two shifts, and 10% last 3 or more shifts. Appendix Figure A4 also plots the event study coefficients for switches that last more than 2 shifts.

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

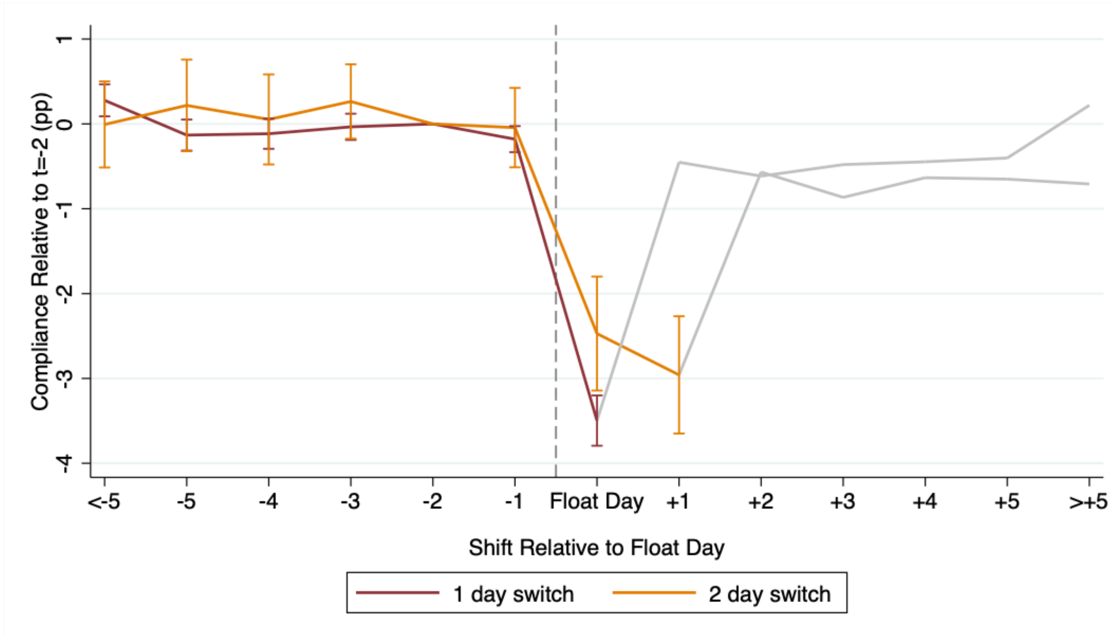

Figure 7. Event study graphs: float days (department switches)

Note: Coefficients and 95% confidence intervals estimated from equation 6 with event indicators E indicating department switches that last 1, 2, or 3+ shifts. See text for more detail on the plotted coefficients. Coefficients for event indicators surrounding permanent switches shown in Appendix Figure A4. The x-axis indicates shifts relative to the switch. To the left of the dashed line are shifts prior to the switch, float day indicates the first day in the new department, followed by subsequent shifts (+1, +2, etc.). The red (darker) line plots the event-time coefficients for a one-shift float day while the orange (lighter) line plots the event-time coefficients for a float that lasts two shifts. The gray line plots the coefficients after the float day, when the worker leaves the new department.

shift, suggesting that the benefit of time off has been completely washed out by the decay in habit. [19]

I synthesize these effects in Table 2 by adding indicators for simplified versions of the events to the baseline regression Equation (5) estimated in Column 1. Specifically, I estimate

$$c_{ilt} = \beta_F F_{it} + \beta_c consec_{it} + \beta_l long_{it} + \beta_f float_{it} + \beta_B B_{ilt} + \omega_i + \gamma_l + \eta_{lt} + \epsilon_{ilt} \quad (7)$$

where $long_{it}$ is an indicator for a shift following a long break (14+ days off), and $float_{it}$ is an indicator for a float day that lasts 1 or 2 shifts. The reference categories are shifts worked after a short break (1-13 days) and non-float days.

Column 2 of Table 2 shows the results. Relative to a short break (1-13 days off), shifts

[19] In Appendix Figure A5, I also include event time coefficients for the longest breaks (more than 4 weeks). Returns from these breaks make up only 0.5% of shifts, consistent with these being unusual events. Declining pre-trends indicate that the timing of these breaks may occur endogenously at low-compliance times, such as taking a leave of absence following a period of burnout. For this reason, I treat these very long breaks separately from more reasonable breaks of 2-4 weeks.

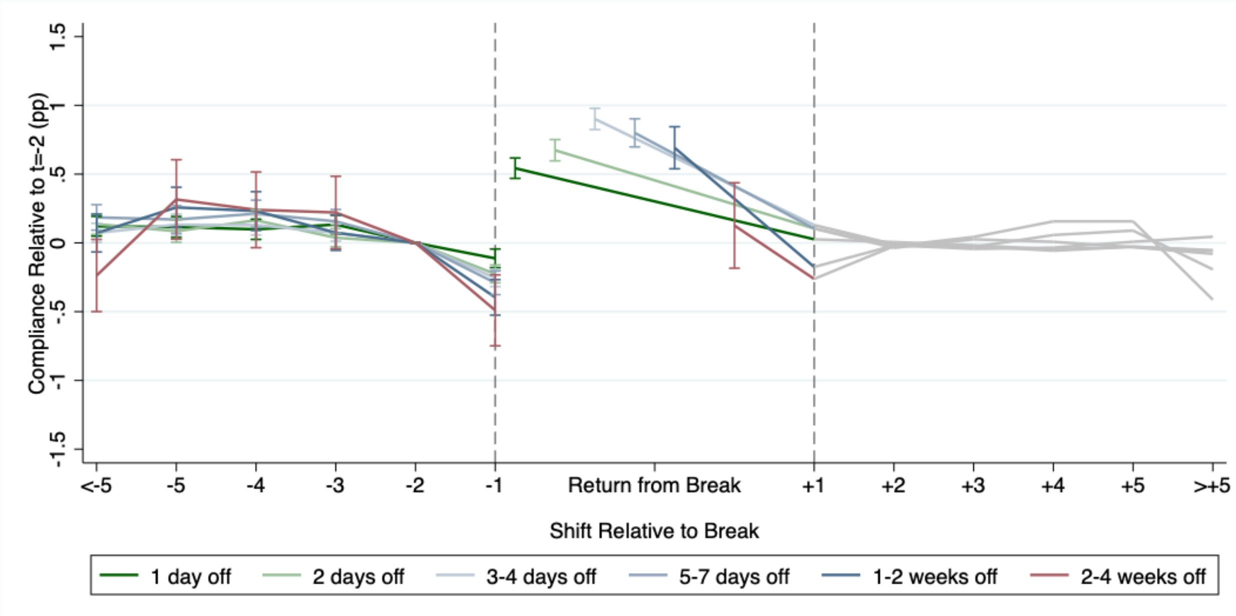

Figure 8. Event study graphs: days off between shifts

Note: Coefficients and 95% confidence intervals estimated from equation 6 with event indicators E indicating breaks between shifts of different lengths. Event time indicators for all breaks less than 4 weeks are shown in the figure; very long breaks are included in Appendix Figure A5. The x-axis indicates shifts relative to the break. To the left of the first dashed line are shifts prior to the break. Return from break is the estimated coefficient on the shift after a break of the specified length. The second dashed line at +1 indicates the second day after a break of the specified length, followed by subsequent shifts (+2, +3, etc.). The gray lines identify these shifts after the day of return.

worked on consecutive days have 0.77pp (s.e. 0.02) lower compliance. Shifts worked after returning from a long break (2-4 weeks) have 1.4pp (s.e. 0.14) or about 2% lower compliance. As before, these findings are consistent with the two opposing impacts of time off: rest to combat fatigue and decay of habit. For short breaks, the former dominates and for long breaks, the latter. On float days, compliance is 3.04pp (s.e. 0.19) or about 5% lower than on non-float days. This is consistent with location-specific habits that are disrupted when moving to a different department.

4.3 Fact 3: Fatigue has a larger impact on low compliers

According the cue-based habit model, fatigue should impact *attentive* hand washing but not *habitual* hand washing. Therefore, fatigue should have a larger effect when habit is low and hand washing requires more attention (Prediction 3).

Assuming that compliance is increasing in habit (Prediction 2), I use compliance levels as a proxy for habit and estimate the effects separately for people with low, medium, and compliance. Empirically, this heterogeneity analysis is only possible because of the precision of the estimates in Table 2. Even after splitting the sample, I have enough statistical power to distinguish effect sizes between subgroups.

To identify types, I split the data into a categorization sample, on which I identify types, and a separate estimation sample, in which I estimate the responses to fatigue and habit disruptions for each type.

To categorize workers into low, medium, and high compliers, I use compliance in each individual's first three months of the data, controlling for the fatigue, disruptions, and benefits. Specifically, I estimate Equation (7) including only each worker's first three months in the data (the categorization sample). From this regression estimate, I obtain individual fixed effects which I split into terciles.[20]

I then estimate the baseline regression in Equation (7) separately for each of the three sets of workers using only data from months four and beyond (the estimation sample). Splitting the data into separate categorization and estimation samples eliminates many of the possible mechanical relationships between the categorization and estimated effects. As an extreme example, someone with 0 or 100% compliance mechanically must have 0 effect sizes during that period. The separation between categorization and estimation breaks that link.

I report the results in Table 3 and show the results graphically in Figure 9. There are 4,495 workers in each bin with average compliance of 45%, 63%, and 78%.[21]

Consistent with Prediction 3, the negative impacts of fatigue are concentrated among low compliance workers. The magnitudes of the differences across types are large; high compliers decline 43% less than low compliers over the course of a shift (6.71pp vs. 3.82pp). The patterns are similar for the other fatigue measures; high compliers respond 76% less to working the night shift (0.5pp vs. 2.3pp) and 26% less to working a consecutive shift (0.81pp vs. 0.60pp).

4.4 Fact 4: Disruptions impact both low and high compliers

According to Prediction 4, disruptions to habit should affect exactly the opposite people from shocks to fatigue. Those who rely more on habit to wash will be more impacted by a shock to habit than those who need to pay costly attention to perform hand hygiene. As discussed in Section 2, the predicted heterogeneity across types is ambiguous if shocks to habit occur with distractions.

Using the same analysis as Fact 3, 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 3,

[20] The main results are robust to alternative ways of categorizing compliance levels, as described in Section 4.4.1.

[21] Appendix Table C displays summary statistics separately for the three groups. High types visit the data

(a) Fatigue

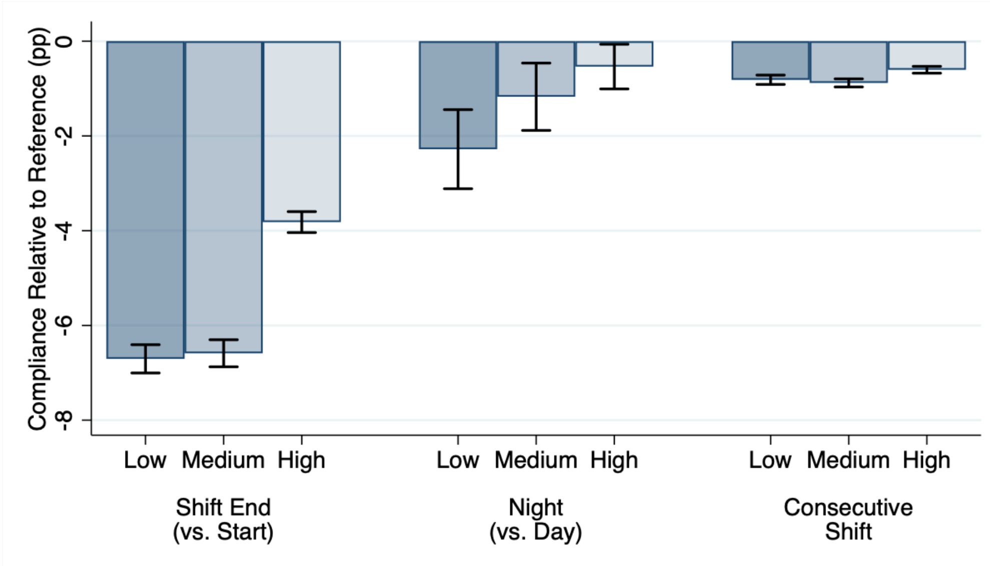

(b) Habit Disruptions

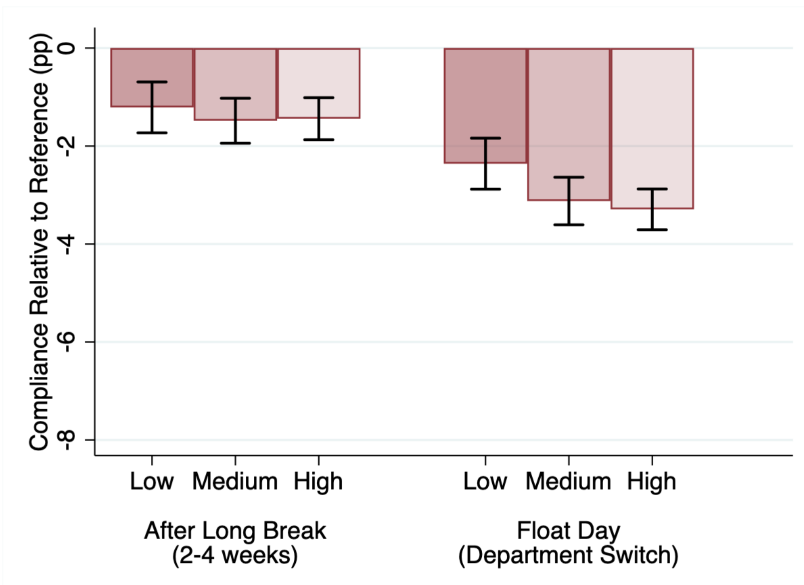

Figure 9. Effect of fatigue and habit disruptions on compliance

Note: Graph plots estimated coefficients and 95% confidence intervals for the impacts of fatigue (a, blue bars) and habit disruptions (b, red bars) relative to the omitted categories. For each effect, the three bars indicate the estimated estimates for low (L), medium (M), and high (H) compliers. Coefficient estimates are from an OLS regression of compliance on fatigue, habit disruptions, and proxies for the benefits of hand hygiene, as reported in Table 3.

shown graphically in Figure 9.

Contrary to the effects of fatigue, disruptions to habit affect high compliers at least as much as low compliers. After returning from a long break, high compliers experience a drop in compliance of about 1.4pp (s.e. 0.22). Low compliers experience a similar drop of 1.2pp (s.e., 0.26). The point estimates imply that high compliers respond 19% more than low compliers, but the difference is not statistically significant. On float days, high compliers experience a drop in compliance 39% larger than their low compliance peers. This difference is only marginally significant; compliance of high compliers falls by 3.3pp (s.e. 0.21) compared to 2.4pp (s.e. 0.27) for low compliers.

Consistent with the relatively large standard errors on these rare events, this heterogeneity across types is less robust than previous results. Across alternative methods for categorizing Low and High compliers as shown in F, the results vary. Generally speaking, there are either no significant differences across the types, or the high types respond more, as predicted by the cue-based habit model.

4.4.1 Robustness of Facts 1-4

The results above are robust to a number of alternative specifications and modeling choices, which I describe in more detail below.

High vs. Low Compliers A key piece of evidence in support of the cue-based habit model is the differential response to fatigue and habit disruption between low and high compliers. However, the choice of how to identify low vs. high compliers is somewhat arbitrary. The results in Table 3 are robust to alternative methods for categorizing low vs. high compliers. In Appendix Figure A6, I plot variants of Figure 9 under alternative categorizations described below.

In the main analyses, I divide workers into terciles based on individual fixed effects estimated from a regression of compliance over an individual's first three months in the data on distractions, disruptions, and benefit proxies as described in Equation (7). This is a parametric approach that attempts to controls for differences in compliance across individuals and settings that do not reflect differences in habit.

A more non-parametric approach is to bin by average compliance in the first 3 months of the data, as shown in Appendix Figure A6 (a). Comparing average compliance levels in the categorization period assumes that compliance across hospitals and over time is the same, or in the model, assumes the same benefits b and distractions d. Therefore, individuals with the same compliance level have the same habit. Given hand hygiene is the same action with similar consequences across places and over time, this is a reasonable benchmark. The results are qualitatively very similar, with the response to habit disruptions statistically indistinguishable across types.

As a middle ground to the full regression-adjusted estimates used in the main analyses, in (b), I account for the variation in b and d across networks and roles. Rather than binning into thirds across all individuals, I bin average compliance over the first three months 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. The results are qualitatively very similar, with the response to habit disruptions statistically indistinguishable across types.

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 ever form than the last three months. In (c), I bin by average compliance in the last three months of the 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 fatigue and shocks to habit are more pronounced under this alternative categorization.

Fatigue Within-Shift The majority of distraction and habit disruptions proxies vary at the shift level. The linearity of the model makes an implicit assumptions that these effects are additive. More generally, the model assumes that after controlling for benefits, distractions, and disruptions, an individual's shifts are otherwise the same.

While estimates of shift-level effects necessitate some assumptions about cross-shift variation, fatigue *within a shift* can be estimated without any such assumptions. In Appendix Table A2, I estimate variants of the regressions in Table 3 that replace the shift-level controls with individual-shift fixed effects. In these regressions, the fatigue effect is identified entirely within-shift. Under this much more flexible specification, the estimated effects are nearly identical.

Definition of Compliance Hospital guidelines require workers to wash or sanitize before entering and after exiting a patient room. The SwipeSense system records compliance with this guideline as using a sensored soap or sanitizer dispenser within 30s of entry and exit. If there is more than 30s between patient visits, this requires two washes between patients: one when leaving the first patient room and one when entering the next. Are the observed effects driven by providers thinking that one wash between patients is enough?

In Appendix Table A3, I show the results under a more lenient measure of compliance that gives providers credit for washing at least once between patient rooms. I include only room entries as observations and define compliance as washing or sanitizing either within 30 seconds of entry or within 30 seconds of the prior room exit (if not the first opportunity of the shift).[22] I estimate this variant of Equation (5) separately for Low, Medium, and High compliers as defined in the main analyses.

By construction, compliance under this alternative measure is higher (58% vs. 45% for low types, 78% vs. 63% for medium types, and 90% vs. 79% for high types). However the estimated impacts of fatigue and habit disruptions are very similar, both in absolute levels and in the patterns of responses across types. If there are diminishing returns to washes between patients, these results indicate that the drops in compliance are not concentrated among the marginally less beneficial opportunities.

4.5 Fact 5: Location-specific habit formation

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

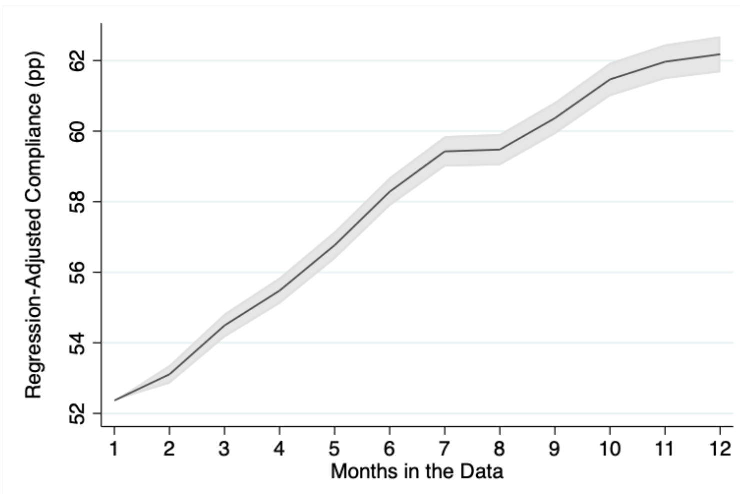

Figure 10. Compliance over the first year in the data

Note: Plot shows estimated coefficients and 95% confidence interval from a regression of compliance on months in the data and individual fixed effects, as described by Equation (8). Y-axis re-scaled to begin at average first-month compliance. Sample includes only the first year in the data and is limited to workers with at least 12 months in the data. Standard errors clustered by individual.

First, I show how compliance evolves over the first year in the data for the 8,076 workers with tenure of at least one year. In Figure 10, I plot compliance by month, controlling for

[22]In the data, I can only observe compliance that occurs within 30s of room entry or exit, so I cannot generalize the measure to any washing between patient visits.

individual differences.[23] Compliance increases substantially over the first year in the data. From a first month average compliance of 52pp, compliance increases by 10pp over the first year.[24]

This increase in compliance over a worker's tenure could be explained by many models 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 simple 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 alternative orientation may 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.

To analyze the effect of room-specific experience, I zoom in on a subset of workers who are plausibly new hires, and thus for whom we can fully capture their experiences over their first year at the hospital. Specifically, I focus on the 3,995 workers who first appear in the data at least 8 weeks after the start of the data for each hospital and who remain in the data for at least one year.

I then examine how compliance varies with experience in a room, relative to that worker's first time visiting the room. To compare across similar rooms, I focus on individual-room pairs that reach at least 100 opportunities in the first year (28% of individual-room pairs) and estimate changes in compliance over these first 100 opportunities.

In Figure 11, I plot the evolution of compliance in a room as function of location-specific

[23] Specifically, I estimate the following regression equation using the first 12 months of data for each worker

$$c_{it} = \sum_{k=1}^{12} \beta_k 1\{t = k\} + \omega_i + \epsilon_{ilt} \tag{8}$$

where c_{it} is compliance by individual i at time t, $\{\beta_k\}$ are the set of indicators for each month in the data, ω_i are individual fixed effects, and ϵ_{it} is an error term. Standard errors are clustered by individual.

[24] 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. 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.

experience. Specifically, I estimate the following regression equation

$$c_{ilt} = RoomOpp_{ilt} + r_{il} + mo_{it} + \beta_F F_{it} + \beta_c consec_{it} + \beta_l long_{it} + \beta_B B_{ilt} + \epsilon_{ilt} \tag{9}$$

where $RoomOpp_{ilt}$ are fixed effects for bins of past opportunities in the room, r_{il} are individual-room pair fixed effects, and mo_{it} are individual-month fixed effects. These mo_{it} capture individual-level changes in total experience across months as well as the network-level time trends controlled for in the Equation (7). The other variables are proxies for distractions, disruptions, and benefits as described in Equation (7). As room history is more carefully accounted for, I exclude the indicator for float days from this regression. Standard errors are clustered by individual.

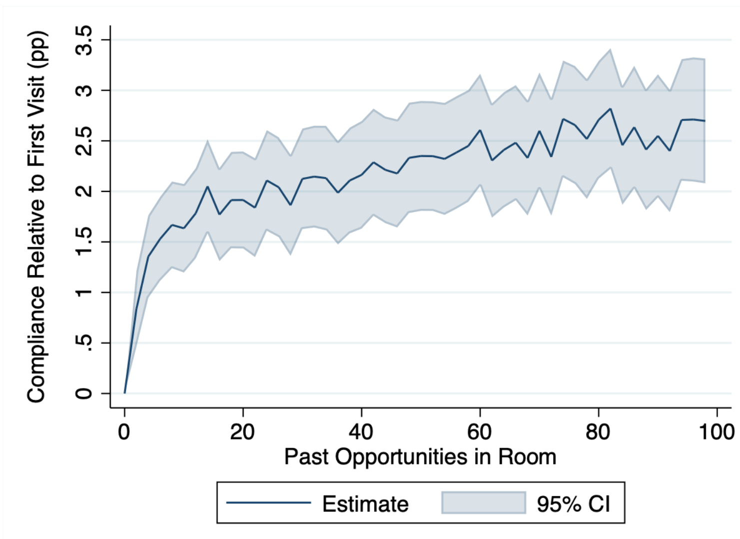

Figure 11. Initial Opportunities in a Room

Note: Figure plots coefficient estimates and 95% confidence intervals for compliance relative to an individual's first visit to a room as a function of the number of past opportunities in that room. Estimates from a regression of compliance on binned values of the number of prior opportunities in a room, individual-room fixed effects, individual-month fixed effects and controls for distractions, disruptions, and benefits as described in Equation (9). Sample includes only the first year in the data and is limited to workers with at least 12 months in the data who first appear in the data at least 8 weeks after the start of the hospital's data. For each individual, the sample is further restricted to individual-room pairs that reach 100 opportunities.

Figure 11 plots the coefficients and 95% confidence intervals for the key coefficients of interest, $RoomOpp_{ilt}$. These fixed effects show compliance relative to that individual's first visit (2 opportunities) to that room, controlling for an individual's total experience through individual-month fixed effects. The Figure shows that compliance increases substantially with room-specific experience. After 10 opportunities, compliance is about 1.5pp higher than in the first visit. After this large initial increase, compliance continues to grow with experience, albeit at a slower rate. After 30 opportunities, compliance is another 0.5pp higher than at 10 opportunities. Even after 50 opportunities, compliance increases another 0.5pp over the next 50 opportunities.

As these experience effects are largest for the initial opportunities, a natural question is to what extent room-specific learning matters over the longer run. In Appendix Figure A8, I plot the share of total opportunities that occur in rooms with various levels of past experience. In the first month of the data, almost half of opportunities occur in rooms with less than 10 prior opportunities. While declining over time as experience mechanically increases, the share of opportunities that occur in very infrequently visited rooms (< 10 opportunities) levels out at a non-trivial share of 5%, indicating that these room-specific experience effects continue to have an effect on compliance even after a year at the hospital.

4.5.1 Robustness of Fact 5

In this Section, I explore the robustness of the finding on room-specific experience and discuss some of the additional assumptions of the habit formation process as described in Section 2.

Of the 317,499 individual-room pairs in this sample, only 28% ever reach 100 opportunities. As these frequently-visited rooms are outliers, one may worry that behavior in these rooms is the exception rather than the norm. In Appendix Figure A7, I show that the pattern is similar using lower thresholds of 30 (51% of individual-room pairs), 50, or 70 opportunities. In all samples, the pattern is well-approximated by a log relationship. In Column 1 of Appendix Table A4, I estimate a variant of Equation (9) including all opportunities in all rooms visited in the first year by this new hire sample. The estimates imply slightly larger benefits to experience, with compliance in 100th opportunity 3.7pp higher than the first.

The habit model assumes the relevant room-experience measure is the number of past washes, rather than opportunities. However, a regression of past washes on compliance is confounded if past washes and compliance are both driven by temporary shocks to the benefits of washing. The analysis in Figure 11 is like an intent-to-treat analysis with past opportunities as a plausibly more exogenous proxy for past hand washing. Though potentially confounded, in Column 2 of Appendix Table A4, I show that the experience effects are driven by compliant opportunities.

A potential concern is that, with so many different jobs and departments, the number of opportunities may not be the right measure of experience. In Appendix Figure A9, I plot the relationship between compliance and experience measured in relative terms as the share of past opportunities which occurred in that room. Excluding highly visited outliers, the results are robust.

Finally, in Column 3 of Appendix Table A4, I explore the role of decay in habits over time. Consistent with the model, I find that more recent experiences are more correlated with behavior, but that experiences in the distant past (2+ months ago) still matter.

4.6 Revisiting the models

To summarize the empirical evidence, Table 4 lists the empirical findings along with a summary of whether heterogeneity in the benefits of hand washing and cue-based habit can explain them.

The first two facts, that distractions and disruptions lower compliance, can be explained by either model. In the standard model, both shocks are a drain on attention. In the habit model, these shocks are different.

The third and fourth findings describe the heterogeneity across types and are the key pieces of evidence distinguishing the models.

The cue-based habit model, built on the feature that attention and habits are substitutes, predicts a specific pattern of heterogeneity in responses to distractions and disruptions to habit across low and high habit types. Low types respond more to fatigue and high types respond more to habit. Finding 3 confirms that the low types respond more to fatigue. On the disruptions to habit (Finding 4), both low and high types respond, with some evidence for a stronger response by high compliers. As discussed above, this can be rationalized in the cue-based habit model if disruptions to habit also increase fatigue.

The standard model, with heterogeneity in b, predicts that high compliers will respond more to both distractions and disruptions. This is clearly rejected by empirical Finding 3.

It is interesting to note that the cue-based habit predictions, supported by Findings 3 and 4, can be differentiated more generally from any model in which some types (e.g., high or low compliers) are more elastic to shocks. For example, a logit model predicts those close to 50% compliance should be the most responsive to shocks. This could fit Finding 3 (as low compliers are closest to 50%) but would imply that those types should also respond more to disruptions to habit. This is rejected by Finding 4.

An alternative model that fits these empirical findings would need to account not only for the heterogeneity in responsiveness across types, but also the differential heterogeneity across types of shocks.

Finally, Finding 5 provides evidence that compliance is increasing in location-based experience. This is a natural consequence of the cue-based habit, but is less obvious under alternative models. Generating this location specificity in the standard model requires room-specific changes in perceived costs or benefits of hand washing that correlate with experience.

5 Implications for Organizations & Conclusion

In this paper, I document evidence of the automaticity of habit in hand hygiene behavior. The empirical findings are consistent with a model of cue-based habit, in which habits are built through repeated pairing of a cue with performing hand hygiene. Attention and habits in this model are substitutes; as habits build, workers have to pay less attention and are less affected by drains on attention like fatigue and distraction.

Many habit interventions incentivize behavior for a short period of time, aimed at encouraging habit formation. With cue-dependent habits, this type of intervention could work in theory. However, it requires incentivizing behavior across a wide variety of cues to effectively build a set of cue-specific habits. This context-specificity of habit points to why persistent behavior change is difficult to implement (for examples, see Acland and Levy, 2015; Carrera et al., 2018; Royer et al., 2015).

Moreover, events that disrupt habit are ubiquitous in the workplace. Industries like agriculture, hospitality, and retail are inherently seasonal, with long breaks between performing the same tasks. Many jobs involve changes in physical location, from consultant or sales jobs to construction work on different sites. The findings in this paper suggest that these aspects of the job may impede habit formation and lower productivity.

Cue-based habits can also explain why successful interventions are multi-dimensional, encouraging behavior change through several approaches and likely covering a variety of contexts (for examples on hand hygiene, see Naikoba and Hayward, 2001; Kingston et al., 2016; Gould et al., 2017). As one 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 involved convening multidisciplinary hand hygiene teams in each department who identified barriers to compliance, evaluated workflows, developed and provided training to peers, developed and implemented a system of peer-observation, data collection, and feedback. Luckily, habits also point to a strategy for simplifying such intensive multi-dimensional approaches. The model makes specific predictions about who are not benefiting from automatic habits and when they are most susceptible. Targeting interventions to those individuals at those times can be more efficient than these intensive, broad-based approaches.

Instead of changing incentives, cue-based habit implies that altering cues could be a more direct approach to facilitate habit formation. The location-specificity of habits (both across departments and rooms) suggest that cues can be tied to the physical environment. Standardizing cues across environments, for example through room layouts, has the potential to increase compliance by making habits more portable across rooms. In hospitals, 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)).

Alternatively, job design that incorporates these ideas could focus on standardizing tasks and cues for a particular person. For hand hygiene in hospitals, this could mean concentrating work in a smaller set of rooms. Common features of the hospital work environment, like rotational resident programs or float days, put workers in departments where they have had little experience. Even within a department, assigning workers to the same rooms across shifts may facilitate faster habit accumulation and limit habit decay over time. While only 7 rooms are visited in the typical shift, the average worker in the data visits a total of 85 different rooms across the hospital.

To quantify the potential impact of standardization across rooms or the concentration of work in rooms with similar layouts, I estimate the predicted change in compliance if all of a worker's experience were concentrated in a single room using the log-experience model estimated in Column 1 of Appendix Table A4. Under this model, if all of an individual's opportunities were concentrated in a single room (or if all rooms had identical layouts), compliance would be 2.7pp higher. Estimates on the benefits of compliance suggest that an increase of this size would result in 38k fewer infections, almost 2,227 fewer deaths, and $0.8B in cost savings in the US alone.[25]

Under the cue-based habit model, this concentration of experience and corresponding increase in habit will be paired with a reduction of costly attention paid to hand hygiene. This should result in further increases in compliance as the impacts of fatigue and distractions are diminished. Beyond hand hygiene, if total attention capacity is limited, this effect could spillover into other domains. As automatic habits take over, less attention is paid to hand hygiene and more cognitive resources are available for the many other important tasks that require a healthcare worker's attention.

Without automatic habits, incentives to increase compliance require healthcare workers to pay more attention to hand hygiene. This reallocation of attention may have perverse effects in other domains. Testing this implication in the field would be a valuable area for future research.

Though one of the simpler actions required by healthcare workers, habits in hand hygiene are complex. Cue-based habit formation implies a constellation of many small habits, which

[25]Based on estimates from (Pittet et al., 2000), I assume a 1% increase in compliance results in a 1.2% reduction in hospital-acquired infections. From an average compliance of 60.5%, the regression estimates imply standardization would increase compliance by 2.7pp (or 4.5%) in the first year. Applying this estimate to only the 42% of opportunities by people in their first year, would result in a 2.2% reduction in HAI. Changes are calculated from a baseline of 1.7M infections and 99k deaths as estimated by the CDC (2009) and $35 billion in costs estimated by Scott (2009).

can explain the inability of short-term incentives to successfully build long-term habits. It also implies a trade-off between attention and habits that distinguishes automaticity from other models and has consequences for spillovers to performance on other tasks. For hand hygiene, these findings have direct implications for interventions, where even small changes in behavior can have meaningful impacts for patients and hospitals.

These same lessons about habits, cues, and attention apply more broadly to any organization trying to improve employee behavior on routine tasks and to individuals trying to make or break routines in their own lives.

References

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

Allcott, H., Gentzkow, M., and Song, L. (2020). Digital addiction.

Baker, M. A., Sands, K. E., Huang, S. S., Kleinman, K., Septimus, E. J., Varma, N., Blanchard, J., Poland, R. E., Coady, M. H., Yokoe, D. S., et al. (2021). The impact of covid-19 on healthcare-associated infections. *Clinical infectious diseases: an official publication of the Infectious Diseases Society of America*.

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.

Benabou, R. and Tirole, J. (2003). Intrinsic and extrinsic motivation. *The review of economic studies*, 70(3):489–520.

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*.

Bordalo, P., Conlon, J., Gennaioli, N., Kwon, S. Y., and Shleifer, A. (2021). Memory and probability. Working Paper.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2020). Memory, attention, and choice. *The Quarterly journal of economics*.

Brown, T. L. and Carr, T. H. (1989). Automaticity in skill acquisition: Mechanisms for reducing interference in concurrent performance. *Journal of Experimental Psychology: Human Perception and Performance*, 15(4):686.

Buyalskaya, A., Ho, H., Milkman, K., Duckworth, A., and Camerer, C. (2021). Predicting context-sensitivity of behavior in field data using machine learning. Working Paper.

Camerer, C., Landry, P., and Webb, R. (2018). The neuroeconomics of habit. *Working Paper. Available at SSRN 3752193*.

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.

Centers for Disease Control (CDC) (2009). Cdc to distribute $40 million in recovery act funding to help states fight healthcare-associated infections. `https://www.cdc.gov/media/pressrel/2009/r090901.htm`. [Online; accessed 10-24-2021].

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

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

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.

Danner, U. N., Aarts, H., and de Vries, N. K. (2008). Habit vs. intention in the prediction of future behaviour: The role of frequency, context stability and mental accessibility of past behaviour. *British Journal of Social Psychology*, 47(2):245–265.

Enke, B., Schwerter, F., and Zimmermann, F. (2020). Associative memory and belief formation. National Bureau of Economic Research Working Paper 26664.

Gardner, B., Abraham, C., Lally, P., and de Bruijn, G.-J. (2012). Towards parsimony in habit measurement: Testing the convergent and predictive validity of an automaticity subscale of the self-report habit index. *International Journal of Behavioral Nutrition and Physical Activity*, 9(1):1–12.

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of political economy*, 108(5):874–904.

Gibbons, R. and Roberts, J. (2012). 2. economic theories of incentives in organizations. In *The handbook of organizational economics*, pages 56–99. Princeton University Press.

Gneezy, U. and Rustichini, A. (2000). Pay enough or don't pay at all. *The Quarterly journal of economics*, 115(3):791–810.

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).

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*.

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.

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.

Holmstrom, B. and Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization*, 7:24.

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*.

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

Ji, M. F. and Wood, W. (2007). Purchase and consumption habits: Not necessarily what you intend. *Journal of Consumer Psychology*, 17(4):261–276.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan, London, United Kingdom.

Kelly, A. C. and Garavan, H. (2005). Human functional neuroimaging of brain changes associated with practice. *Cerebral cortex*, 15(8):1089–1102.

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.

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

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.

Malmendier, U. and Wachter, J. A. (2021). Memory of past experiences and economic decisions. Manuscript prepared for Oxford Handbook of Human Memory.

Mullainathan, S. (2002). A memory-based model of bounded rationality. *The Quarterly Journal of Economics*, 117(3):735–774.

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., Wu, M., and Kurlander, D. (2011). The pull of the past: When do habits persist despite conflict with motives? *Personality and Social Psychology Bulletin*, 37(11):1428–1437.

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].

Orbell, S. and Verplanken, B. (2010). The automatic component of habit in health behavior: Habit as cue-contingent automaticity. *Health psychology*, 29(4):374.

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.

Potthoff, S., McCleary, N., Sniehotta, F. F., and Presseau, J. (2018). Creating and breaking habit in healthcare professional behaviours to improve healthcare and health. In *The Psychology of Habit*, pages 247–265. Springer.

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.

Quinn, J. M., Pascoe, A., Wood, W., and Neal, D. T. (2010). Can't control yourself? monitor those bad habits. *Personality and Social Psychology Bulletin*, 36(4):499–511.

Quinn, J. M. and Wood, W. (2005). Habits across the lifespan. *Unpublished manuscript, Duke University, Durham, NC.*

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.

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.

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.*

Semmelweis, I. P. (1861). *Die aetiologie, der begriff und die prophylaxis des kindbettfiebers.* Hartleben.

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.

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

Tricomi, E., Balleine, B. W., and O'Doherty, J. P. (2009). A specific role for posterior dorsolateral striatum in human habit learning. *European Journal of Neuroscience*, 29(11):2225–2232.

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.

Verplanken, B. (2018). *Psychology of habit*. Springer, Cham, Switzerland.

Verplanken, B. and Orbell, S. (2003). Reflections on past behavior: a self-report index of habit strength. *Journal of applied social psychology*, 33(6):1313–1330.

Verplanken, B. and Roy, D. (2016). Empowering interventions to promote sustainable lifestyles: Testing the habit discontinuity hypothesis in a field experiment. *Journal of Environmental Psychology*, 45:127–134.

Volpp, K. G. and Loewenstein, G. (2020). What is a habit? diverse mechanisms that can produce sustained behavior change. *Organizational Behavior and Human Decision Processes*, 161:36–38.

Von Thadden, E.-L. and Zhao, X. (2012). Incentives for unaware agents. *The Review of Economic Studies*, 79(3):1151–1174.

Wachter, J. A. and Kahana, M. J. (2019). A retrieved-context theory of financial decisions. National Bureau of Economic Research Working Paper 26200.

Weiner-Lastinger, L. M., Pattabiraman, V., Konnor, R. Y., Patel, P. R., Wong, E., Xu, S. Y., Smith, B., Edwards, J. R., and Dudeck, M. A. (2021). The impact of coronavirus disease 2019 (covid-19) on healthcare-associated infections in 2020: A summary of data reported to the national healthcare safety network. *Infection Control & Hospital Epidemiology*, pages 1–14.

Weiss, H. M. and Ilgen, D. R. (1985). Routinized behavior in organizations. *Journal of Behavioral Economics*.

Wood, W. and Neal, D. T. (2007). A new look at habits and the habit-goal interface. *Psychological review*, 114(4):843.

Wood, W., Quinn, J. M., and Kashy, D. A. (2002). Habits in everyday life: Thought, emotion, and action. *Journal of personality and social psychology*, 83(6):1281.

Wood, W., Tam, L., and Witt, M. G. (2005). Changing circumstances, disrupting habits. *Journal of personality and social psychology*, 88(6):918.

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

Table 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)	45%		
Long break (2-4 weeks off)	0.9%		
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 2. Impact of Fatigue and Habit Disruptions on Compliance

Dependent Variable: Compliance	(1)	(2)
Fatigue		
Shift end (12 hr)	-5.88***	-5.89***
	(0.08)	(0.08)
Night shift	-1.41***	-1.41***
	(0.19)	(0.19)
Consecutive shift (no day off)	-0.71***	-0.77***
	(0.02)	(0.02)
Habit disruptions		
Long break (2-4 weeks off)		-1.40***
		(0.14)
Float day (department switch)		-3.04***
		(0.19)
Individual FE	X	X
Department FE	X	X
Network-Month-Year FE	X	X
Benefit Proxies	X	X
Very long break (4+ weeks off)		X
Opportunities	122,804,335	122,804,335
Workers	13,606	13,606
Mean Compliance	60.5	60.5
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. Consecutive shifts are worked with no days off in between. Long break indicates shift that follows a 2-4 week break. Very long breaks last more than 4 weeks. Reference group is 1 to 13 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 weekend (vs. weekday), opportunities per hour (normalized within individual), room entry (vs. exit), a quadratic trend over months in the data, and an indicator for the first shift in the data. Standard errors are clustered by individual.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3. Heterogeneity in Response to Fatigue and Habit Disruptions

Dependent Variable: Compliance	(1) Low Compliance	(2) Medium Compliance	(3) High Compliance
Fatigue			
Shift end (12 hr)	-6.71***	-6.59***	-3.82***
	(0.15)	(0.15)	(0.11)
Night shift	-2.28***	-1.17***	-0.54**
	(0.43)	(0.36)	(0.24)
Consecutive shift (no day off)	-0.81***	-0.88***	-0.60***
	(0.05)	(0.04)	(0.04)
Habit disruptions			
Long break (2-4 weeks off)	-1.21***	-1.48***	-1.44***
	(0.26)	(0.23)	(0.22)
Float day (department switch)	-2.36***	-3.12***	-3.29***
	(0.27)	(0.25)	(0.21)
Individual FE	X	X	X
Department FE	X	X	X
Network-Month-Year FE	X	X	X
Benefit Proxies	X	X	X
Very long break (4+ weeks off)	X	X	X
Opportunities	33,382,442	32,104,567	31,616,314
Workers	4,495	4,495	4,495
Mean Compliance	45.3	64.0	78.4
R^2	0.171	0.115	0.085

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 estimated individual fixed effects from a regression of compliance on distractions, disruptions, and benefits in the first three months in the data (as described in Equation (7)). Coefficients estimated on months 4+. 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. Consecutive shifts are worked with no days off in between. Long break indicates shift that follows a 2-4 week break. Very long breaks last more than 4 weeks. Reference group is 1 to 13 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 weekend (vs. weekday), opportunities per hour (normalized within individual), room entry (vs. exit), a quadratic trend over months in the data, and an indicator for the first shift in the data. Standard errors are clustered by individual.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4. Consistency of Empirical Findings with Different Models

Empirical Finding	Standard Model Heterogeneity in b	Cue-based Habit Model Heterogeneity in h
(1) Compliance lower late in the shift, on weekends, and on nights	✓	✓
(2) Compliance lower on float days and after long breaks	✓ Assuming $\uparrow d$	✓
(3) Effect (1) larger for low types	X Opposite prediction	✓
(4) Effect (2) affects high and low types	✓ Assuming $\uparrow d$	✓
(5) Compliance is increasing in location-specific experience	X	✓

Notes: Table presents the empirical findings from Section 4 and whether they can be explained by the models from Section 2.

Appendix A Generalized Theoretical Framework

In this section, I derive results for a more general version of the theoretical framework described in Section 2.

Specifically, I assume that the utility from paying attention level a can be expressed as

$$U(a) = b_i p(a, h_{ilt}) - d_t c(a). \tag{10}$$

Benefits b_i and distractions d_t are as described in Section 2. Attention costs $c(a)$ are convex with $c_a(a) > 0$and $c_{aa}(a) > 0$. $p(a, h_{ilt})$ is compliance as a function of attention and habits. This set-up nests both the standard model, with $p(a, h_{ilt}) = a$, and the automatic cue-based habit model, with $p(a, h_{ilt}) = a + (1 - a)h_{ilt}$.

To define some notation, I use a^* to refer to the optimal attention level, which will be a function of habit h, benefits b, and distractions d and p^* to refer to the optimal compliance under the same conditions. I use $p_h(a^*, h)$ as shorthand for $\left.\frac{\partial p(a,h)}{\partial h}\right|_{a=a^*}$ and p_h^* as shorthand for $\frac{\partial p^*(b,d,h)}{\partial h}$.

In the results that follow, I make several assumptions about the structure of $p(a, h)$..

First, I assume that compliance is increasing in attention, or $p_a(a, h) \geq 0$. This is a trivial restriction given the set-up of the model.

Second, I assume that the p functions are linear or at most multiplicative combinations of attention and habits (i.e., $p(a, h) = x_0 + x_1 a + x_2 h + x_3 ah$). This delivers zero higher-order derivatives $p_{aa} = 0$, $p_{hh} = 0$. This assumption gives structure to the derivatives of optimal compliance with respect to b, d, and h which allow us to make progress in identifying patterns of heterogeneity in responsiveness.

Finally, I assume that optimal compliance is increasing in habit, or that the parameters are such that $p_h^*(b, d, h) \geq 0$. As discussed in 2, this assumes that the direct effect of habits on compliance outweighs the indirect effect (through a potential change in attention). This assumption is non-trivial, but seems plausible in this context. Note that under this assumption, habits that result in a lower level of optimal attention ($a_h^*(b, d, h) < 0$), like cue-based habit, must also directly increase compliance through $p(a, h)$ in order to offset the effect of lower attention.

The first order condition for this generalized model yields

$$c_a(a^*) = \frac{b_i}{d_t} p_a(a^*, h). \tag{11}$$

Result 1. *Attention and compliance are decreasing with other drains on attention.*

Proof. Taking the derivative of the first order condition with respect to the d,

$$\begin{aligned}
a_d^*(b_i, d_t, h_{ilt}) &= -\frac{1}{d_t^2 c_{aa}(a^*)} \left(b_i p_a(a^*, h)\right) \\
&= -\frac{1}{d_t} \frac{c_a(a^*)}{c_{aa}(a^*)} \\
&\leq 0.
\end{aligned}$$

Substituting in $c_a(a^*)$, with convex costs, the derivative is negative.
Taking the derivative of optimal compliance with respect to distractions,

$$p_d^*(b_i, d_t, h_{ilt}) = p_a(a^*, h) a_d^*(b_i, d_t, h_{ilt}) \leq 0$$

as $p_a(a, h) \geq 0$ and $a_d^*(b_i, d_t, h_{ilt}) \leq 0$. □

Result 2. *Habits are automatic and reduce costly attention when $p_{ah}(a^*, h) \leq 0$.*

Proof. Taking the derivative of the first order condition with respect to habit h,

$$a_h^*(b_i, d_t, h_{ilt}) = \frac{b_i}{d_t c_{aa}(a^*)} p_{ah}(a^*, h). \tag{12}$$

Attention is declining in habits when $a_h^*(b_i, d_t, h_{ilt}) \leq 0$. □

Result 2 identifies the key feature that delivers automaticity in this model of habit; with higher habits, the worker washes more while also paying less attention.

A.1 Heterogeneity in the standard model

In this section, I discuss sources of heterogeneity across individuals in the standard model with $p(a, h) = a$.

Section 2.1 showed that heterogeneity across individuals in b leads to variation in compliance levels and responses to drains on attention. In this section, I show that the results apply for any cost function that is not significantly increasing in convexity (i.e., $c_{aaa}(a^*)$ sufficiently small).

Result 3. *In the standard model, attention and compliance are increasing in benefits and decreasing in drains on attention.*

Proof. Taking the derivative of the first order condition with respect to b and d, yield

$$a_b^*(b_i,d_t) = p_b^*(b_i,d_t) = \frac{1}{d_t c_{aa}(a^*)} \geq 0 \tag{13}$$

$$a_d^*(b_i,d_t) = p_d^*(b_i,d_t) = -\frac{b_i}{d_t^2 c_{aa}(a^*)} \leq 0 \tag{14}$$

as $c_{aa}(a^*) > 0$. □

Result 4. *In the standard model, the impact of drains on attention are larger for individuals with high b_i whenever $c_{aaa}(a^*)$ is sufficiently small.*

Proof. The impacts of drains on attention will be larger (more negative) for individuals with high b_i if $p_{bd}^*(b_i,d_t) \leq 0$. Taking the derivative of Equation 13 with respect to d, yields

$$a_{bd}^*(b_i,d_t) = p_{bd}^*(b_i,d_t) = -\frac{1}{d_t^2 c_{aa}(a^*)} - \frac{a_d^*(b_i,d_t)}{d_t(c_{aa}(a^*))^2} c_{aaa}(a^*).$$

The first term is negative and the second term matches the sign of $c_{aaa}(a^*)$. The derivative will be negative whenever

$$c_{aaa}(a^*) \leq -\frac{1}{d_t}\frac{c_{aa}(a^*)}{a_d^*(b_i,d_t)}.$$

□

A.2 Heterogeneity in the habits model

Section 2.2 showed that heterogeneity in habits, holding constant b and d, drive heterogeneity in hand hygiene compliance and differential heterogeneity across types of shocks. Individuals with higher habit respond less to drains on attention and more to shocks to habit. As in the previous section, the predictions described in Section 2.2 hold for any cost function that is not significantly increasing in convexity (i.e., $c_{aaa}(a^*)$ sufficiently small).

Result 5. *When $p_{ah}(a^*,h_{ilt}) \leq 0$, low habit types are more responsive to changes in drains on attention whenever $c_{aaa}(a^*)$ is sufficiently small.*

Proof. Low habit types will be more responsive to changes in drains on attention when $p_{dh}^*(b_i,d_t,h_{ilt}) \geq 0$. Taking the derivative of $a_d^*(b_i,d_t,h_{ilt})$ in Equation 12 with respect to h,

$$a_{dh}^*(b_i,d_t,h_{ilt}) = \frac{1}{d_t} a_h^*(b_i,d_t,h_{ilt}) \left(-1 + \frac{c_a(a^*)c_{aaa}(a^*)}{c_{aa}(a^*)^2}\right)$$

and a cross partial of compliance with respect to distractions and habits of

$$p^*_{dh}(b_i, d_t, h_{ilt}) = p_a(a^*, h_{ilt})a^*_{dh}(b_i, d_t, h_{ilt}) + p_{ah}(a^*, h_{ilt})a^*_d(b_i, d_t, h_{ilt}). \tag{15}$$

When $p_{ah}(a^*, h_{ilt}) \leq 0$, as in the cue-based habit model, the second term is always positive. The first term will also be positive when $a^*_{dh}(b_i, d_t, h_{ilt}) \geq 0$ which is true whenever $c_{aaa}(a^*) \leq c_{aa}(a^*)^2/c_a(a^*))$. (15) will be positive whenever $c_{aaa}(a^*)$ is sufficiently small so the first term is either positive or negative but smaller than the second term. □

The next result generalizes the prediction that individuals with a high level of habit will respond more to any shocks to habit.

Result 6. *When $p_{ah}(a^*, h_{ilt}) \leq 0$, high habit types are more responsive to changes in habits whenever $c_{aaa}(a^*)$ is sufficiently small.*

Proof. Taking the derivative of $a^*_h(b_i, d_t, h_{ilt})$ with respect to habit h,

$$\begin{aligned} a^*_{hh}(b_i, d_t, h_{ilt}) &= -\frac{b_i}{d_t}\frac{c_{aaa}(a^*)}{c_{aa}(a^*)^2}p_{ah}(a^*, h)a^*_h(b_i, d_t, h_{ilt}) \\ &= -\frac{c_{aaa}(a^*)}{c_{aa}(a^*)}\left(a^*_h(b_i, d_t, h_{ilt})\right)^2 \end{aligned} \tag{16}$$

which is positive when $c_{aaa}(a^*) \leq 0$.
Under the assumptions of the functional form of $p(a, h)$,

$$p^*_{hh}(b_i, d_t, h_{ilt}) = p_a(a^*, h_{ilt})a^*_{hh}(b_i, d_t, h_{ilt}) + p_{ah}(a^*, h_{ilt})a^*_h(b_i, d_t, h_{ilt}). \tag{17}$$

As $p^*_h(b_i, d_t, h_{ilt}) \geq 0$ by assumption, high habit types are more responsive to changes in habits when (17) is positive. When $p_{ah}(a^*, h_{ilt}) \leq 0$, the second term is positive as $a^*_h(b_i, d_t, h_{ilt}) \leq 0$. When $c_{aaa}(a^*)$ is sufficiently small, or when

$$c_{aaa}(a^*) \leq \frac{c_{aa}(a^*)a^*_h(b_i, d_t, h_{ilt})p_{ah}(a^*, h_{ilt})}{p_a(a^*, h_{ilt})}$$

(17) will be positive. □

Appendix B Electronic Monitoring System

Figure A1. SwipeSense Hand Hygiene Monitoring System

Note: SwipeSense system of caregiver badges, hygiene sensors, location hub, and communication hubs. Graphic from `https://www.swipesense.com/electronic-hand-hygiene-monitoring`.

Appendix C Sample Selection

Table A1. Sample Selection

	All Data	Main Sample	Compliance Subsamples		
			Low	Medium	High
Total					
Opportunities	166,557,886	123,465,032	33,508,404	32,290,174	31,837,578
Shifts	3,882,105	2,123,320	630,682	552,813	506,973
Workers	40,356	13,606	4,495	4,495	4,495
Networks	25	21	20	21	21
Average					
Compliance	59%	61%	46%	64%	79%
Shift length (hours)	7.4	9.7	9.5	9.7	10.0
Opportunities per hour	11.1	5.7	5.4	5.8	6.1
Tenure (months)	11.1	17.5	19.3	17.3	15.8
Shifts					
Weekday		77%	78%	77%	76%
Day shift		64%	65%	67%	62%
Float days		1.8%	1.8%	1.9%	1.8%
Consecutive shift		45.2%	46.0%	45.1%	45.0%
Long break (2-4 weeks)		0.9%	1.0%	1.0%	0.9%

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 individual fixed effects from a regression of compliance on proxies for fatigue, disruptions, and benefits in the first three months in the data (as described in Equation (7)). Data summarized for the compliance subsamples excludes the first three months used to categorize.

Appendix D Hand Hygiene During the COVID-19 Pandemic

In Appendix Figure A2, 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 A3 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.

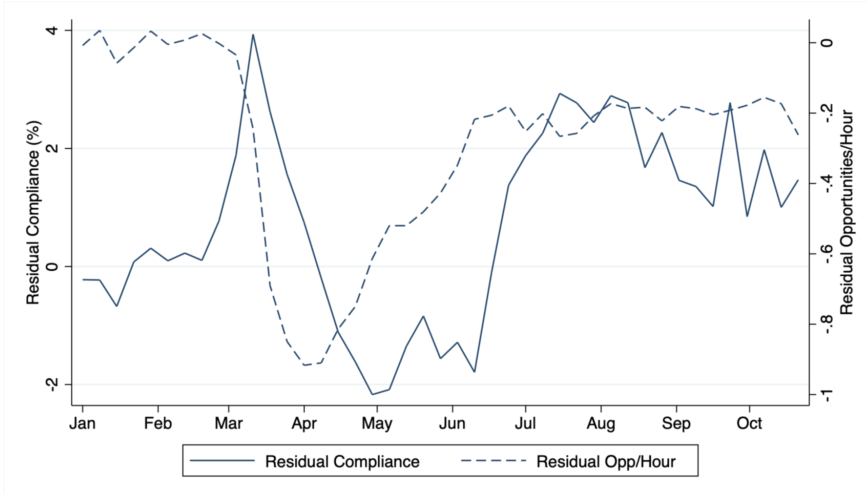

Figure A2. Compliance and number of opportunities in 2020

Note: Residualized compliance and number of opportunities per hour (residualizing individual levels in Jan and Feb 2020).

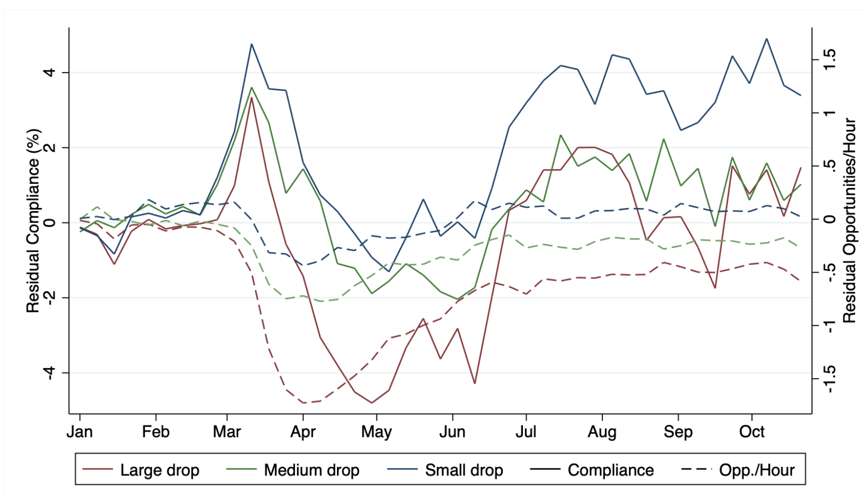

Figure A3. Compliance and number of opportunities in 2020, by change in opportunities

Note: 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.

Appendix E Additional Event Study Plots for Disruptions to Habit

In Appendix Figure A4, I plot event time dummies for department switches that last 3+ days in addition to those shown in the main text.

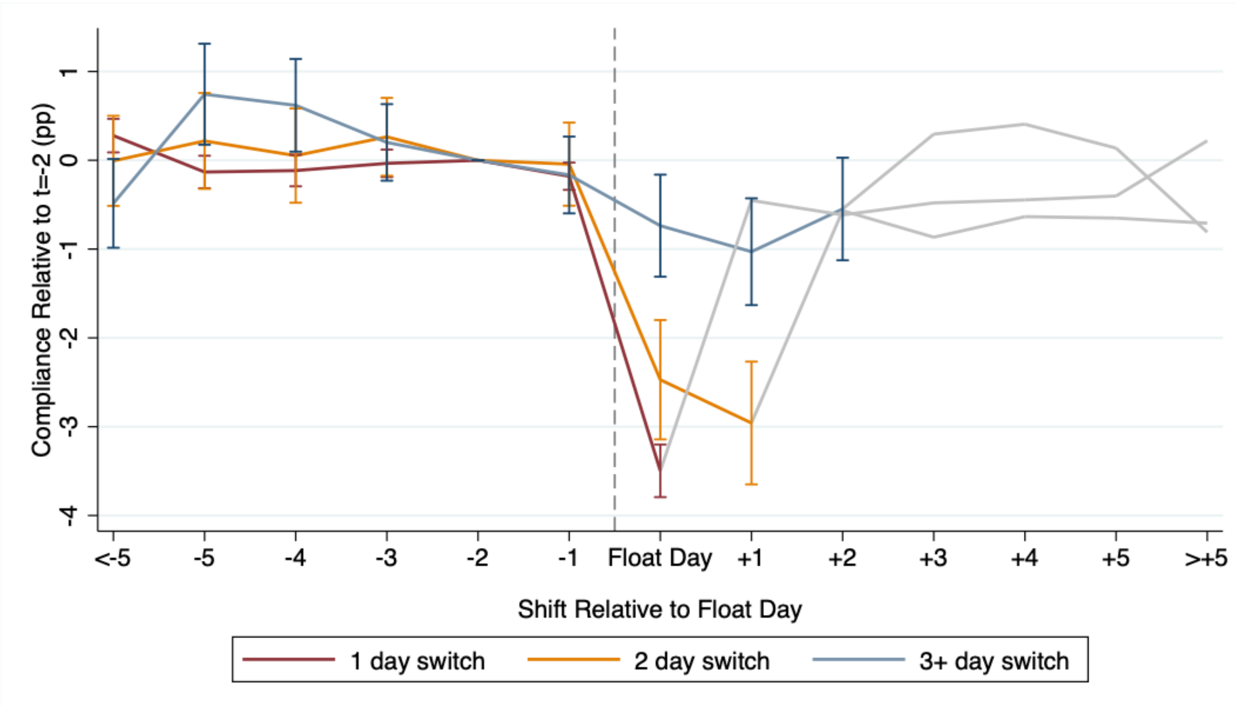

Figure A4. Event study graphs: float days (department switches)

In Appendix Figure A5, I plot event time dummies for breaks of all length including very long breaks of more than 4 weeks. A these long breaks exhibit a significant pre-trend, I treat them separately from long, but more reasonable length breaks of 2-4 weeks.

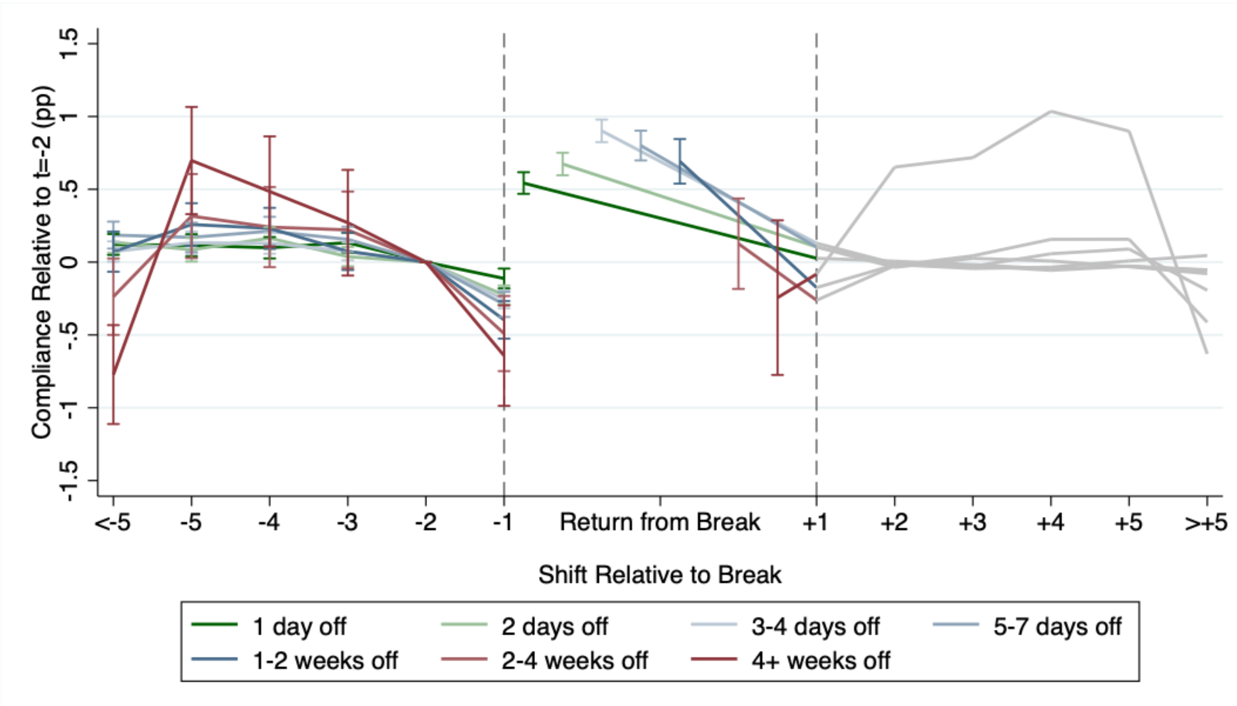

Figure A5. Event study graphs: days off between shifts

Note: See text for description of plotted coefficients.

Appendix F Alternative Categorizations

(a) Compliance in First 3 Months

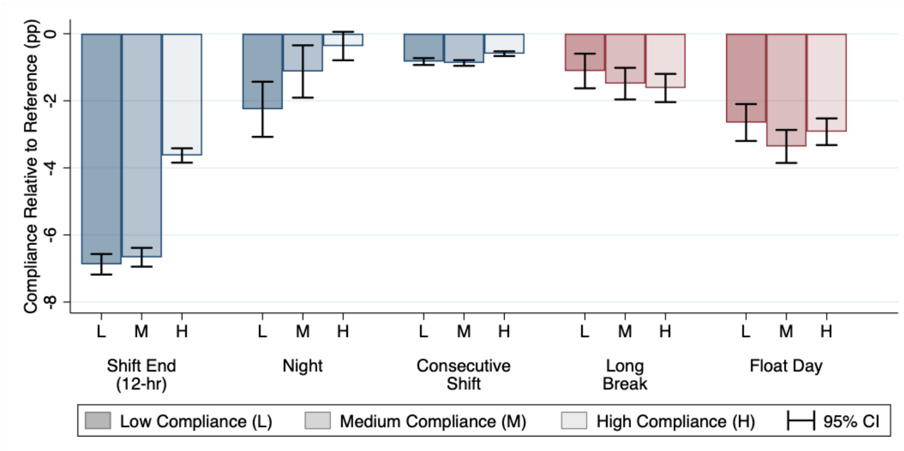

(b) Compliance in First 3 Months, Relative to Network-Role

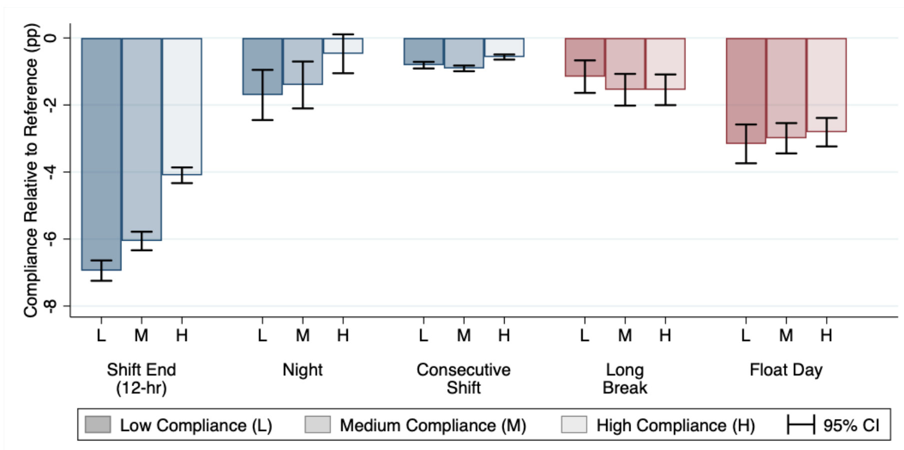

(c) Compliance in Last 3 Months (Dec 2019-Feb 2020)

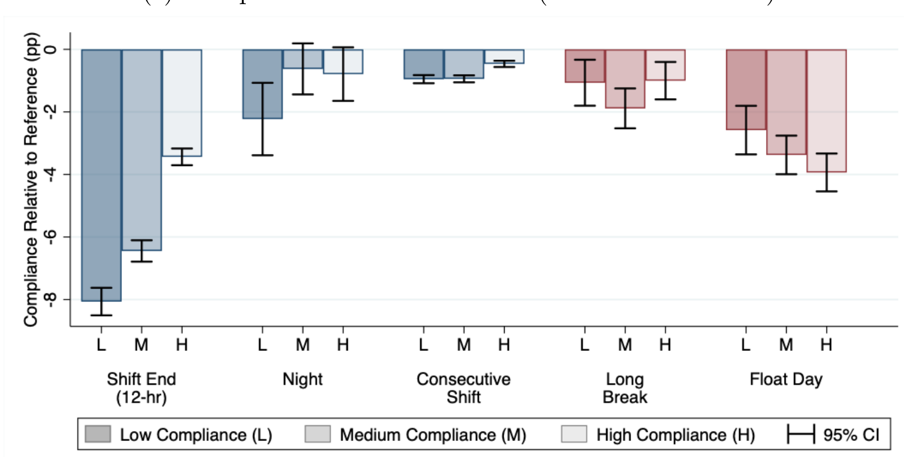

Figure A6. Effect of fatigue and habit disruptions on compliance: alternative categorizations

Appendix G Within-Shift Variation

Table A2. Heterogeneity: Fatigue Within-Shift

Dependent Variable: Compliance	(1) Low Compliance	(2) Medium Compliance	(3) High Compliance
Shift end (12 hr)	-6.91***	-6.67***	-3.93***
	(0.15)	(0.15)	(0.11)
Individual-Shift FE	X	X	X
Within-Shift Benefit Proxies	X	X	X
Department FE	X	X	X
Opportunities	33,382,442	32,104,567	31,616,314
Workers	4,495	4,495	4,495
Mean Compliance	45.3	64.0	78.4
R^2	0.256	0.188	0.150

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 estimated individual fixed effects from a regression of compliance on proxies for fatigue, disruptions, and benefits in the first three months in the data (as described in Equation (7)). Coefficients estimated on months 4+. 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). Within-shift benefit proxies include opportunities per hour (normalized within individual) and room entry (vs. exit). Standard errors are clustered by individual.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix H Alternative Compliance

Table A3. Heterogeneity by Type: Any Compliance Between Interactions

Dependent Variable: Any Compliance Between	(1) Low Average Compliance	(2) Medium Average Compliance	(3) High Average Compliance
Fatigue			
Shift end (12 hr)	-6.05***	-4.97***	-2.22***
	(0.17)	(0.14)	(0.09)
Night shift	-2.83***	-1.51***	-0.64***
	(0.44)	(0.33)	(0.18)
Consecutive shift (no day off)	-0.86***	-0.77***	-0.38***
	(0.05)	(0.04)	(0.03)
Habit disruptions			
Long break (2-4 weeks off)	-1.08***	-1.23***	-1.02***
	(0.28)	(0.23)	(0.19)
Float day (department switch)	-2.37***	-3.05***	-2.81***
	(0.27)	(0.24)	(0.19)
Individual FE	X	X	X
Department FE	X	X	X
Network-Month-Year FE	X	X	X
Benefit Proxies	X	X	X
Very long break (4+ weeks)	X	X	X
Observations (Entries)	16,446,109	15,857,170	15,644,628
Workers	4,495	4,495	4,495
Mean Any Comp. Btwn.	59.6	78.3	89.5
R^2	0.344	0.249	0.188

Notes: Table shows coefficients from OLS regressions. Each observation is a room entry. The dependent variable is an indicator, rescaled to 100, for compliance at room entry or compliance in the previous room exit (if not the first opportunity of the shift). Data split by estimated individual fixed effects from a regression of compliance on proxies for fatigue, disruptions, and benefits in the first three months in the data (as described in Equation (7)). Coefficients estimated on months 4+. 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. Consecutive shifts are worked with no days off in between. Long break indicates shift that follows a 2-4 week break. Very long breaks last more than 4 weeks. Reference group is 1 to 13 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 weekend (vs. weekday), opportunities per hour (normalized within individual), a quadratic trend over months in the data, and an indicator for the first shift in the data. Standard errors are clustered by individual.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix I Additional Tables and Figures on Room-Specificity

In Column 1 of Appendix Table A4, I estimate a variant of Equation (9) including all opportunities in all rooms visited in the first year by this new hire sample. With the larger pool of rooms, including the 50% with less than 30 opportunities, I estimate room fixed effects, instead of the individual-room pair fixed effects as in Figure 11. The estimates imply slightly larger benefits to experience, with compliance in 100th opportunity 3.7pp higher than the first.

The habit model assumes the relevant room-experience measure is the number of past washes, rather than opportunities. In Column 2 of Appendix Table A4, I estimate the impact of compliant opportunities in addition to total opportunities. Consistent with habit formation, the positive effect of experience is driven by compliant opportunities. The problem with the interpretation of this specification is that past compliance in a room may be correlated with current benefits. For example, a room with an immunocompromised patient has a temporary increase in the benefits of hand hygiene generating a string of compliant opportunities. This will appear as past compliant opportunities predicting future compliance, while in fact it's the result of correlated benefits in a room over time. Using all past opportunities as the room-experience measure, as in Figure 11 or Column 1 of Appendix Table A4, is like an intent-to-treat analysis with past opportunities as a plausibly more exogenous proxy for past hand washing.

While the model assumes habits are a function of cumulative experiences, there are alternative ways of comparing experience across rooms. In Appendix Figure A9, I plot the relationship between compliance and the share of total past opportunities that have occurred in a room. Specifically, I regress compliance on binned values of the share of prior opportunities in a room (plotted), room fixed effects, individual-month fixed effects and controls for distractions, disruptions, and benefits as described in Equation (9). As with the number of past room opportunities, compliance is increasing with the concentration of past opportunities in a room. Compliance in rooms with a 2% share of past opportunities (the 25th percentile) is about 3pp higher than a room with no past opportunities. The relationship flattens at a 3.75pp increase around the median at a 4% share of past opportunities. For the most-highly visited rooms ($\geq$10% share of opportunities, approximately the 90th percentile), compliance is about 1.5pp higher than rooms with no prior opportunities.[26]

Some of this reversal in the relationship between very high experience and compliance may be driven by recency effects. If habits decay over time, as indicated by long breaks,

[26] These outliers are concentrated in the initial months when the total rooms visited are low.

past washes that happened more recently should carry more weight. Rooms that have been visited very frequently may have many opportunities that occurred long ago. In Column 3 of Appendix Table A4, I break down past room opportunities into those occurred in each of the last three months and more than three months ago. In this analysis, I limit the sample to 3 months in the data and later. Breaking down experiences across time, there is a clear recency effect, with opportunities in the last month weighted more than those in preceding two months. This supports the assumption that habits decay over time. Note that there is also a large positive effect of opportunities that occurred more than three months ago, suggesting that there is also a long-run persistence in the impact of room-specific experience.

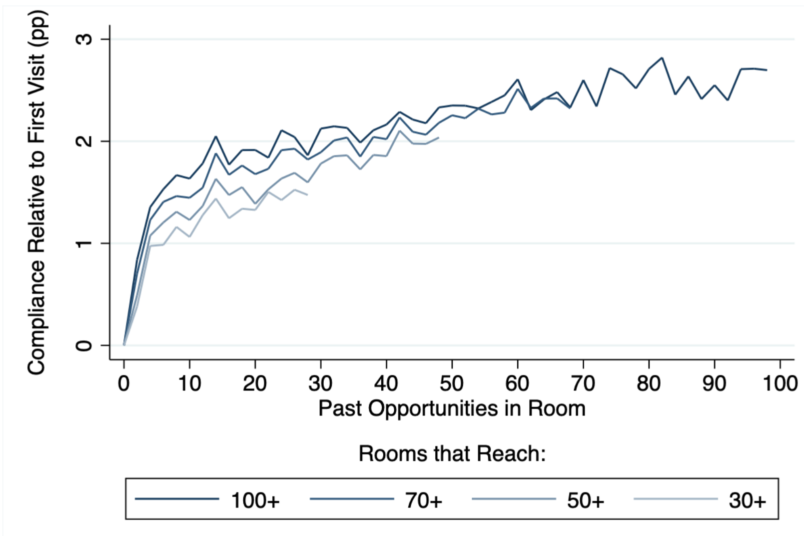

Figure A7. Initial Opportunities in a Room

Note: Figure shows compliance relative to an individual's first visit to a room as a function of the number of past opportunities in that room. Lines are estimated coefficients from a regression of compliance on binned values of the number of prior opportunities in a room, individual-room fixed effects, individual-month fixed effects and controls for distractions, disruptions, and benefits as described in Equation (9). Each line plots coefficients from separate regressions limited to the first 30, 50, 70 or 100 opportunities. Sample includes only the first year in the data and is limited to workers with at least 12 months in the data who first appear in the data at least 8 weeks after the start of the hospital's data.

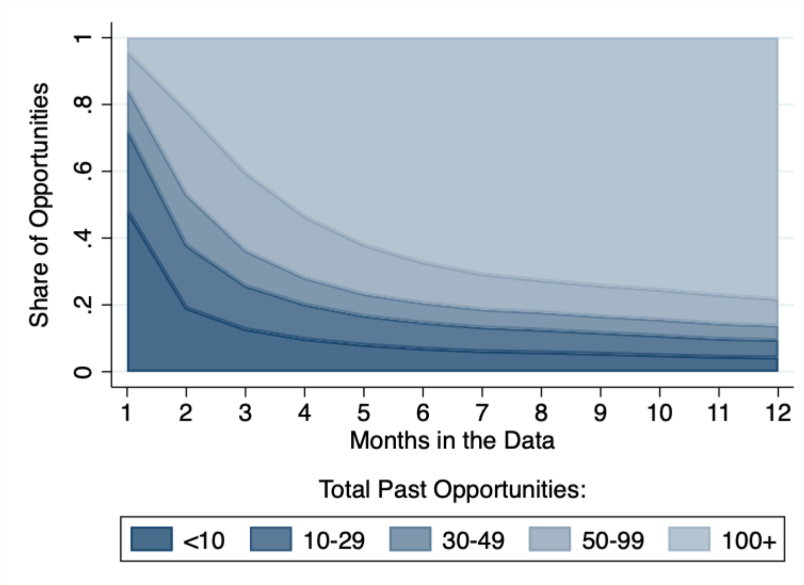

Figure A8. Share of Opportunities by Frequency of Past Room Visits

Note: Figure plots the share of opportunities that occur in rooms with different levels of past experience, as measured at the start of the shift. Y-axis is the share of opportunities, which always sums to 1 across the bins of past room opportunities. X-axis is the worker's month in the data. Sample includes only the first year in the data and is limited to workers with at least 12 months in the data who first appear in the data at least 8 weeks after the hospital's start of the data.

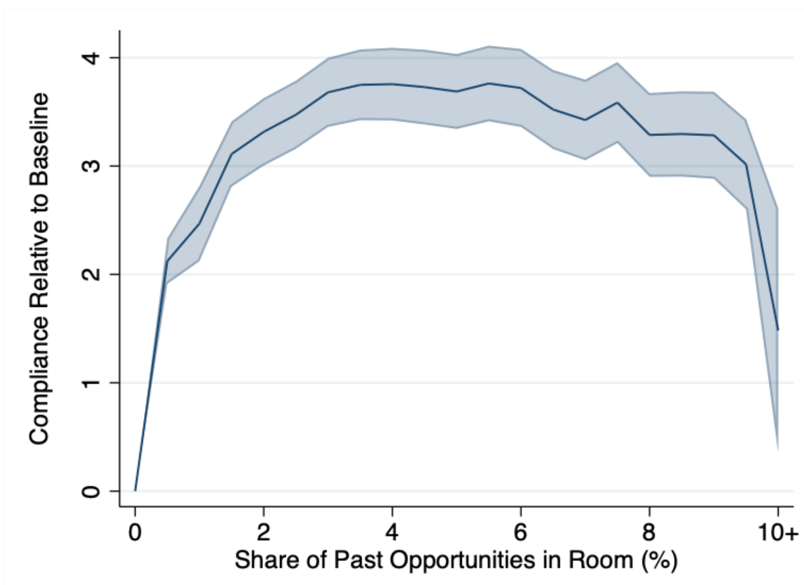

Figure A9. Concentration of Opportunities in a Room

Note: Figure shows compliance relative the first time in a room as a function of the share of past opportunities that have occurred in that room. Lines are estimated from a regression of compliance on binned values of the share of prior opportunities in a room (plotted), room fixed effects, individual-month fixed effects and controls for distractions, disruptions, and benefits as described in Equation (9). Sample includes only the first year in the data and is limited to workers with at least 12 months in the data who first appear in the data at least 8 weeks after the start of the hospital's data. Shaded region is the 95% confidence interval.

Table A4. Room-Specific Habits, First Year in the Data

Dependent Variable: Compliance	(1)	(2)	(3)
Log of Past Room Opportunities	0.80***	-15.66***	
	(0.06)	(0.25)	
Log of Past Compliant Room Opportunities		17.50***	
		(0.24)	
Log of Room Opportunities Last Month			0.23***
			(0.03)
Log of Room Opportunities 2 Months Ago			0.06***
			(0.02)
Log of Room Opportunities 3 Months Ago			0.05**
			(0.02)
Log of Room Opportunities > 3 Months Ago			0.49***
			(0.03)
Room FE	X	X	X
Individual-Month-Year FE	X	X	X
Distractions	X	X	X
Habit Disruptions	X	X	X
Benefit Proxies	X	X	X
Opportunities	26,144,292	26,144,292	19,154,310
Workers	3,995	3,995	3,986
Mean Comp	64.4	64.4	65.9
R^2	0.234	0.242	0.236

Notes: Table shows coefficients from OLS regressions of compliance (rescaled to 100) in each hand hygiene opportunity on measures of room-specific experience. All log measures are log of described opportunities plus one to account for zeros. Sample is limited to individuals who enter the data at least 8 weeks after the start of the hospital's data and are in the data for at least one year. Data limited to the first year of an individual's data. In Column 3, the data is further limited to months 4-12. Distractions include hour at work and indicators for nights, weekends, and consecutive shifts. Habit disruptions are returns from long (14+ day) breaks. Benefit proxies include opportunities per hour (normalized within individual) and room entry (vs. exit). Standard errors are clustered by individual. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$