Motivated Optimism and Workplace Risk*

A. Yeşim Orhun[†]
University of Michigan

Alain Cohn[‡]
University of Michigan

Collin B. Raymond[§]
Purdue University

November 18, 2021

Abstract

We provide field evidence that individuals engage in motivated optimism in the face of impending risk, and that their belief distortions are time- and stake-dependent. Our study leverages exogenous variation in when people are required to return to their workplaces during the COVID-19 pandemic. Among workers currently staying at home, individuals who are temporally closer to returning to their workplace are relatively more optimistic about the increase in infection risk associated with going back. Temporal belief differences are larger for people who are more likely to get severely ill if infected.

JEL Codes: D84, D91, I12, J28

*We thank Lillian Chen and Basil Isaac for providing excellent research support. We are also thankful for the comments of seminar participants at Carnegie Mellon University, Monash University, Purdue University, Shandong University, University of Chicago, University of Helsinki, University of Michigan, University of Pittsburgh, Virtual Quantitative Marketing Seminar, Yale University, and Zurich University, as well as feedback from David Huffman, Botond Kőszegi, George Loewenstein, Michel Maréchal, Devin Pope, and Peter Schwardmann. The University of Michigan IRB exempt approved this study on May 5, 2020 (HUM00181725).

[†] Ross School of Business and School of Information, University of Michigan (email: aorhun@umich.edu).

[‡] School of Information, University of Michigan (email: adcohn@umich.edu).

[§] Krannert School of Management, Purdue University (email: collinbraymond@gmail.com).

1 Introduction

Beliefs about our future prospects and the world around us are typically more than just inputs for decision-making. They can directly impact well-being by evoking strong feelings such as anxiety and hope. A great deal of anecdotal evidence as well as intuition suggests that individuals often try to manage these emotions by downplaying potential risks in their lives. For example, workers in a benzene plant often denied that the chemicals they worked with were dangerous (Ben-Horin, 1979). In a seminal paper, Akerlof and Dickens (1982) develop the first economic model which explicitly captures the downplaying of risk, focusing on workers who work in dangerous industries and may therefore change their beliefs in order to reduce the psychic costs of "unpleasant feelings of constant fear or unsettling doubts."

Building on these intuitions, an expanding literature proposes that individuals manage their emotional response to the anticipation of future events by engaging in motivated optimism: they overestimate the probability of favorable prospects and underestimate the probability of unfavorable prospects.[1] Motivated optimism is intuitive and compelling in a variety of settings, a point highlighted by Loewenstein (1985), who noted "because higher expectations result in greater pleasure from savoring — people have an incentive to upwardly manipulate their expectations." Indeed, research provides correlational evidence that individuals who face a high chance of natural disasters are also those who refuse to take out insurance (Kunreuther et al., 1978), smokers downplay the risks of smoking (Weinstein et al., 2005), and people deaden worries about their financial well-being (e.g., Brunnermeier and Parker, 2005) and health (e.g., Schwardmann, 2019), even if doing so is costly.

Despite the interest in understanding these types of motivated reasoning, evidence from field settings is scant (see Bénabou and Tirole, 2016, Benjamin, 2019, and Molnar and Loewenstein, 2021 for recent reviews). Identifying motivated optimism in the field is challenging because the ideal setting requires exogenous variation in anxiety-provoking future risk, for example, in impaired health outcomes due to workplace hazards. People typically choose their risk exposure, such as whether to work in a dangerous industry, making it difficult to find a field setting where selection is not a concern. Field experiments that randomly assign individuals to different levels of risk exposure are often not feasible due to ethical concerns.

[1] The class of models studying motivated beliefs based on anticipatory belief-based utility includes Akerlof and Dickens (1982), Loewenstein (1987), Brunnermeier and Parker (2005), Bracha and Brown (2012), Mayraz (2018), and Caplin and Leahy (2019). Other motivations to distort beliefs may include preserving self-esteem, enhancing self-efficacy, and social-signalling (e.g., Bénabou and Tirole, 2002; Köszegi, 2006; Bénabou and Tirole, 2011). Bénabou and Tirole (2016) provide a comprehensive review.

Almost 40 years after the seminal work of Akerlof and Dickens, we are able to explore a setting that parallels their context, offers exogenous imposition of risk, and is experienced by many individuals in the workforce: returning to the workplace during the COVID-19 pandemic. As industries in the U.S. reopened after the initial wave of lockdowns in late April 2020, frontline workers who did not want to lose their jobs had no choice but to return to their workplaces when their employers called them back.[2] Feelings of anxiety about the looming health risks were common in the workforce[3], providing a fruitful setting to probe the importance of motivated optimism in a natural context. Leveraging this unique context, we present field evidence on motivated optimism about health risks associated with returning to the workplace during the COVID-19 pandemic. Intuitively, we would expect the motivation to alter beliefs to be stronger when anticipatory emotions are more gripping, such as when stakes are high or the potential danger is temporally close. We first propose a dynamic model of motivated optimism to analyze the nature and timing of belief adjustments. The model predicts that compared to workers who have a long time until they have to return to work, those who have to return soon are more optimistic about the health risks associated with returning to work relative to staying at home (i.e., *time-dependency*). Our model also predicts that the motivation to distort beliefs is more pronounced when stakes are higher, such as when individuals have a relatively high risk of getting severely ill (i.e., *stake-dependency*).

We then test these predictions using (i) exogenous variation in the time left before workers return to their workplaces and (ii) natural variation in health stakes arising from health conditions that the Centers for Disease Control (CDC) identified as high-risk factors. Through a quick succession of surveys spanning four weeks in May 2020, we elicit subjective risk assessments among workers who are currently staying at home and have no choice but to return to the workplace when called back by their employers. Because we focus on workers whose return dates are not in their control, the staggered nature of the survey generates exogenous variation in return horizons that is orthogonal to baseline risk perceptions, especially after controlling for a rich set of demographics, political affiliation, geographic risk factors, and county-level restrictions. We document that the workers who are about to return to the workplace are relatively more optimistic about the infection risk they will face upon returning to the workplace compared to those who have a longer horizon before they return. Conversely, workers who have a longer horizon are more optimistic about the infection risk associated with staying at home. We also find directional evidence for

[2]Workers had very few legal options to refuse to go back to work without getting fired: https://time.com/5832140/going-back-to-work-coronavirus-rights Our survey focuses on individuals who indicated that they have no choice but to return to their workplaces when called back.

[3]For example, see https://www.bbc.com/worklife/article/20200521-why-do-we-feel-uneasy-about-a-new-normal and https://www.limeade.com/wp-content/uploads/2021/05/2021-EmployeeCareReport-3.0.pdf

increased optimism regarding beliefs about the severity of health outcomes among workers who are temporally closer to returning compared to those with longer horizons. Then, we take a difference-in-differences approach and compare temporal differences in beliefs between high-stakes individuals who are more likely to experience severe health outcomes conditional on getting infected (60 years or older or have medical conditions that put them in the high-risk group) and individuals who are young and healthy. We find that the increase in optimism regarding infection risk and outcome severity is more pronounced among at-risk workers as the return date approaches. Overall, our empirical results confirm that motivated optimism exhibits both time- and stake-dependency.

These results contribute to a broader literature on motivated reasoning in three ways. First, we provide field evidence of motivated beliefs regarding optimism about risks that are relevant to a large part of the population in a major and widely experienced health crisis. There is little field evidence on changes in beliefs arising from motivated reasoning. Most papers focus on documenting overconfidence in a particular sub-population, but do not feature experimental variation (e.g., Park and Santos-Pinto (2010) focus on chess players, while Hoffman and Burks (2020) as well as Huffman et al. (2019) study overconfidence among employees).[4] Thus, these papers try to compare beliefs to some rational benchmark.[5] A smaller set of papers feature quasi-random assignment, but consider very different settings than ours: Di Tella et al. (2007) discuss how access to property rights changes beliefs about political processes (such as rule of law), while Schwardmann et al. (2019) consider self-persuasion by debaters.

To our knowledge, only a single paper, Oster et al. (2013), considers motivated optimism in a field setting.[6] This paper focuses on a very particular subset of individuals. Specifically, it provides evidence of strategic ignorance consistent with motivated optimism from a survey of individuals who had a parent with Huntington's disease and therefore lived with a 50% risk of developing this genetic disease later in life. They document that individuals who chose not to take the genetic test behave like people who know for sure that they do not have the disease, and they believe their chances of developing the disease to still be around 50% even after starting to show symptoms.

[4]There is also a related psychology literature documenting that individuals predict better outcomes for the sports team or political party of their choice (see Krizan and Windschitl 2007 for a discussion of this literature and associated inference issues).

[5]Overconfidence focuses on when individuals alter beliefs about self-relevant traits, such as one's ability or intelligence. Lab studies show that people have a preference for holding inflated beliefs about themselves (e.g., Eil and Rao, 2011; Burks et al., 2013; Charness et al., 2018; Möbius et al., 2014). In the field, people exhibit behaviors consistent with holding overconfident beliefs about their skills (Malmendier and Taylor, 2015). Zimmermann (2020), Huffman et al. (2019), Gödker et al. (2020), and Chew et al. (2020) suggest that asymmetric memory editing may help sustain overconfidence even in the presence of feedback.

[6]Existing evidence on downplaying (inflating) beliefs about the chances of undesirable (desirable) future outcomes mainly comes from the lab (Mayraz, 2011; Coutts, 2019; Engelmann et al., 2019).

Oster et al. (2013) does not feature quasi-experimental variation, and focuses on individuals who have deliberately decided not to get tested. Therefore, it is unclear whether the reported over-optimism is evidence of motivated beliefs or a result of a desire for consistency between one's beliefs about future health outcomes with one's past choices. Furthermore, some researchers (e.g., Klinowski and Paulsen, 2013) have suggested that individuals who will develop Huntington's may have impaired reasoning generated by brain lesions even before they exhibit noticeable symptoms, raising questions about generalizability. Our paper indicates that the desire to downplay risks, as reflected in the findings of Oster et al. (2013), do apply widely, both in terms of context and population.

Second, we examine time dependency in motivated beliefs — an important but understudied phenomenon. Standard approaches in economics suppose that individuals' beliefs evolve over time due to new information. In contrast, we highlight a distinct reason why beliefs may change over time: the passage of time alters the anticipatory utility associated with a future event. Although our empirical evidence is novel, it is presaged in early theoretical work by Loewenstein (1985), who considers the adaptation of motivated beliefs over time; yet this vein of research was not pursued.[7] Studying temporal differences in subjective beliefs is not only of interest in itself, but it also serves as a useful identification strategy for three reasons. First, our approach to identify belief distortions does not necessitate having to compare beliefs to a rational benchmark (i.e., objective risks) which can be hard to ascertain, especially in this context.[8] Second, comparing the beliefs of people who are exposed to risk against those who are not exposed is often confounded with individuals self-selecting into environments with different risks. In our field setting, all workers eventually have to return to their workplaces, return and survey timing is externally imposed, and thus the *time left* to return varies exogenously across workers. Lastly, the temporal differences in optimism we document are unlikely to be the result of an internal desire to align past choices and beliefs about future outcomes, as the work-return date is imposed rather than chosen by workers.

Third, the evidence we provide on stake-dependency speaks to both positive and normative implications of motivated beliefs. Stake-dependency is an important comparative static emerging from many models of motivated beliefs (see Bénabou, 2015, for a discussion). On the positive side,

[7] Gottlieb (2014) is an exception, although the mechanism (imperfect memory) he explores is very different than what we consider in this paper. Moreover, there is a psychology literature on "bracing" that shows that although individuals generally hold optimistic beliefs as they wait for self-relevant news, they conjure up pessimistic expectations right before receiving news about the outcome to protect themselves against disappointment (e.g., Shepperd et al. 1996, Sweeny and Krizan 2013). Macera (2014) formalizes this intuition. However, it is unclear how these results would extend to situations with ongoing risk, i.e., without a certain date for uncertainty resolution.

[8] In particular, rather than asking how beliefs are distorted relative to some objective baseline, we ask about the change in relative distortions over time. Such an approach may be very useful in applied settings where it can be difficult to know what objective beliefs are, or whether subjects have access to them.

providing empirical support lends credence to these models.[9] On the normative side, one narrative about motivated optimism is that it is a "mistake" (e.g., Krizan and Windschitl, 2009). If so, models where mistakes can be minimized via costly effort would predict that individuals with the biggest concerns about the health risks have the weakest incentives to distort their beliefs away from the truth. Instead, we find that these individuals distort their beliefs the most. This is in line with models where anticipatory utility features prominently in the utility function so that higher stakes actually lead to increased belief adjustments.

The next section presents a theoretical model of the dynamics of motivated belief adjustments that underpins our empirical approach. Section 3 details our study design, identification strategy and presents our empirical results. Section 4 concludes with a discussion of implications and future research avenues.

2 Theoretical Framework

Here we provide formal details of the theoretical model that underpins our empirical approach. This model allows us to derive predictions about when and how individuals adjust their beliefs about uncertain and anxiety provoking future outcomes. Our model extends Bracha and Brown (2012), one of the workhouse approaches to motivated beliefs, in two ways.[10] First, individuals face explicit intertemporal concerns. Second, we allow for individuals to hold multiple beliefs and distort each independently.

We consider a representative individual who is initially staying at home (h) until period T before returning to her workplace (w). In any given period t while the individual is in h, the chance of getting sick is p_h. That chance increases to p_w when the individual transitions to w. Each period, the individual gets a flow utility from being healthy (G) or sick (B), which she discounts at a rate δ. To clarify our exposition, and without loss of generality, we normalize the flow utility of being healthy to 1 and the present discounted value (PDV) of being sick to 0.[11]

[9] In the domain of monetary stakes, Kunda (1987), Babcock et al. (1995), Mijović-Prelec and Prelec (2010), and Mayraz (2011) present evidence from lab experiments consistent with stake-dependency of motivated beliefs, although not always in the context of motivated optimism.

[10] Although we focus on a dynamic model of motivated reasoning, one could imagine alternative models that could also generate our pattern of belief distortions. For example, individuals could be trying to rationalize their decision to stay employed (and so have to return to work), and distort their beliefs in service of this. Or it could be that individuals want to avoid anxiety at the point they return to work, and face a non-trivial time to adjust beliefs before the point of their return. To the extent these other models could generate similar comparative statics, we remain agnostic about the exact functional form, and leave it to future research to distinguish more clearly between them.

[11] Given our cardinal utility representation, we can normalize two payoffs. We will later allow the flow utility from being healthy to vary when we examine the role of stakes. Note that the PDV of being sick today includes the stream of values which takes into account the flow values of being sick and being healthy (post recovery), given the discount

We assume that the actual probabilities of getting sick are time invariant. Like Bracha and Brown (2012), we model individuals as experiencing anticipatory utility from thinking about future outcomes. In particular, in any given period the individual gains flow utility from three sources. The first component is standard utility, which depends on actual outcomes — the PDV of flow utility based on true probabilities (p_h, p_w). The second component is anticipatory utility: each period individuals receive utility from thinking about their future outcomes (i.e., the PDV of their flow utility over outcomes) based on their *perceived* probabilities of future events.[12] We denote the perceived risk of getting sick after returning to the workplace which is held in Period t (and used when calculating the perceived PDV of payoffs from the perspective of Period t) as $\hat{p}_{w,t}$, and similarly $\hat{p}_{h,t}$ for the perceived risk of infection when staying at home. The third component is a convex cost of adjusting beliefs to $(\hat{p}_h, \hat{p}_w)$ given true probabilities (p_h, p_w) — this is a cost that prevents beliefs from straying too far from reality (for simplicity we also assume costs are additively separable in each belief).[13] A weight α is placed on anticipatory utility. As it goes to infinity, the individual cares only about anticipatory utility, whereas if the weight goes to 0 we recover the standard model.

Such a framework can be thought of as a "dual-self" model, where the subconscious self knows the true probabilities and cares about realized outcomes, and the conscious self only knows the perceived probabilities and consumes anticipatory utility. The "subconscious" self chooses a sequence of perceived beliefs in order to maximize the discounted sum of their flow utility. The utility function that individuals maximize is

$$\sum_{t=0}^{T-1} \delta^t \underbrace{(1-p_h)^t}_{\text{Probability of staying healthy until period } t<T} \left(\overbrace{\alpha\left[\frac{1}{1-\delta(1-\hat{p}_{h,t})} - \frac{\delta^{T-t}(1-\hat{p}_{h,t})^{T-t}}{1-\delta(1-\hat{p}_{h,t})} + \frac{\delta^{T-t}(1-\hat{p}_{h,t})^{T-t}}{1-\delta(1-\hat{p}_{w,t})}\right]}^{\text{Anticipatory utility in period } t<T} \underbrace{-c_h(p_h,\hat{p}_{h,t}) - c_w(p_w,\hat{p}_{w,t})}_{\text{Cost of adjusting beliefs in period } t} \right)$$
$$+ \sum_{t=T}^{\infty} \delta^t \underbrace{(1-p_h)^{T-1}(1-p_w)^{t-T}}_{\text{Probability of staying healthy until period } t\geq T} \left(\overbrace{\alpha \frac{1}{1-\delta(1-\hat{p}_{w,t})}}^{\text{Anticipatory utility in period } t\geq T} \underbrace{-c_h(p_h,\hat{p}_{h,t}) - c_w(p_w,\hat{p}_{w,t})}_{\text{Cost of adjusting beliefs in period } t} \right)$$
$$+ \overbrace{\left[\frac{1}{1-\delta(1-p_h)} - \frac{\delta^T(1-p_h)^T}{1-\delta(1-p_h)} + \frac{\delta^T(1-p_h)^T}{1-\delta(1-p_w)}\right]}^{\text{Realized utility}}$$

factor.

[12]The model focuses on anticipatory utility from beliefs about infection risk. In the empirical section, we also show that we observe similar changes in beliefs about the severity of infection, conditional on being infected. The model could be extended to show similar results for these kinds of beliefs, but such an extension would come at the expense of clarity and simplicity, as we then need to consider multiple "levels" of beliefs.

[13]One can think of this formulation as a reduced form way of potentially capturing psychological distortion costs, as well as any costs imposed by taking incorrect actions, given a fixed set of actions (e.g., fixing a set of actions, we obtain a convex loss function for taking actions that are optimal for adjusted, rather than true, probabilities).

We now turn to explaining each component of the utility function, starting with anticipatory utility. The first line considers the situation where an individual has not yet returned to the workplace in Period $t < T$ and is still healthy, which happens with chance $(1 - p_h)^t$. In such situations, the PDV of current and future anticipatory utility flows depends on both $\hat{p}_{h,t}$ (for the periods the individual has not yet returned to work) and $\hat{p}_{w,t}$ (for the periods after the individual has returned to work). In particular, the PDV can be written in a simple way: begin with the PDV of staying at home forever , $\frac{1}{1-\delta(1-\hat{p}_{h,t})}$; subtract off the part of the PDV that captures the value of being at home after period T when the individual returns to work, $\frac{\delta^{T-t}(1-\hat{p}_{h,t})^{T-t}}{1-\delta(1-\hat{p}_{h,t})}$;[14] and add in the PDV of being at work starting in Period T, conditional on having stayed healthy at home until T, $\frac{\delta^{T-t}(1-\hat{p}_{h,t})^{T-t}}{1-\delta(1-\hat{p}_{w,t})}$.[15] Note that once an individual gets sick they can no longer gain anticipatory utility from believing they are healthy. We therefore normalize the flow anticipatory utility to 0 for the rest of time. Of course, after returning to work (i.e., $t \geq T$), the probability of not having gotten sick while at home and still being healthy after returning to work is $(1 - p_h)^{T-1}(1 - p_w)^{t-T}$ and the calculation of anticipatory utility becomes simpler, as it is just the PDV of being at work ($\frac{1}{1-\delta(1-\hat{p}_{w,t})}$), given the perceived probability of being sick at work ($\hat{p}_{w,t}$). We assume that costs are additively separable across the two kinds of beliefs. The third component of utility, realized utility, is computed as per standard, and is simply the PDV of payoffs using the true, time invariant, probabilities.

This model embeds the assumptions that the choice of perceived beliefs is made independently across time periods, beliefs are adjusted every period, that in any given period t, the cost of holding subjective beliefs (e.g., $\hat{p}_{w,t}$) is always calculated by comparing the perceived probability in period t to the objective probability (i.e., p_w), and that payoffs are additively separable across time periods.[16] Of course, if we are interested in eliciting the current period beliefs before returning to work (i.e., $t = 0$), under the assumption that the individual is currently healthy, the optimization problem becomes simpler. The optimal beliefs for $t = 0$ (i.e., the current period) can be shown to be the solution to maximizing the following formula:

[14]By assumption, the value, from Period T's perspective, of being healthy in Period T and staying at home for the rest of time is $\frac{1}{1-\delta(1-\hat{p}_{h,t})}$. The value of this in Period t must be discounted by $\delta^{T-t}(1 - \hat{p}_{h,t})^{T-t}$.

[15]When calculating the PDV individuals assume that the risk of infection depends only on period t perceived beliefs.

[16]This gives rise to a form of dynamic inconsistency in beliefs. This implies that although in Period t the individual calculates their Period t anticipatory utility by assuming that the risk in Period $t' > t$ will be $\hat{p}_{w,t}$ and $\hat{p}_{h,t}$. However, when Period t' arrives, the perceived beliefs they use to calculate their anticipatory utility in Period t' will be $\hat{p}_{w,t'}$ and $\hat{p}_{h,t'}$. Because we only elicit current beliefs about risk, and not current beliefs about future beliefs about risk, these assumptions are not testable.

$$\left(\overbrace{\alpha[\frac{1}{1-\delta(1-\hat{p}_{h,0})} - \frac{\delta^T(1-\hat{p}_{h,0})^T}{1-\delta(1-\hat{p}_{h,0})} + \frac{\delta^T(1-\hat{p}_{h,0})^T}{1-\delta(1-\hat{p}_{w,0})}]}^{\text{Anticipatory utility}} \underbrace{-c_h(p_h, \hat{p}_{h,0}) - c_w(p_w, \hat{p}_{w,0})}_{\text{Cost of adjusting beliefs}} \right)$$

$$+ \overbrace{[\frac{1}{1-\delta(1-p_h)} - \frac{\delta^T(1-p_h)^T}{1-\delta(1-p_h)} + \frac{\delta^T(1-p_h)^T}{1-\delta(1-p_w)}]}^{\text{PDV of realized utility}}$$

In particular, individuals maximize this sum by choice of $(\hat{p}_{h,t}, \hat{p}_{w,t})$ for each period t.

In line with our empirical analyses, we compare three types of individuals. One type, NOW, returns to work immediately ($T = 0$). The second type, LATE, returns to work in an intermediate period ($T \geq 1$). The third type, LATER, returns to work in the distant future ($T >> 1$).[17] We define $\Delta_i = \hat{p}_{w,i} - \hat{p}_{h,i}$ as the perceived infection risk differential between returning to work and staying at home for type $i \in \{NOW, LATE, LATER\}$. We assume that belief adjustments preserve the natural ordering of infection risk $\hat{p}_w > \hat{p}_h$ (as is true in the data), such that $\Delta_i > 0$. We show five results.

In stating our results, we will maintain one assumption, namely that belief adjustments preserve the natural ordering of infection risk:[18]

A0: At the optimum $\hat{p}_w > \hat{p}_h$

Our first two results compare individuals returning to work NOW, compared to those who are returning LATER. They are collected in the first proposition. The first part of the proposition says that we should observe that the individuals who have to return to work NOW have a smaller difference between $\hat{p}_w$ and $\hat{p}_h$, compared to those who have to return to work LATER: individuals who are temporally closer to returning decrease the gap between their perceived risk of infection in w and in h: $\Delta_{NOW} < \Delta_{LATER}$. Parts two and three say this change in the belief differential is driven by both a change in the beliefs about the chance of getting sick after returning to the workplace, as well as the chance of getting sick while staying at home: individuals who must return to work sooner become more optimistic about the risk of infection after returning to work

[17]Similar results, albeit with slightly more restrictive conditions, can be obtained so long as we have three dates, one sufficiently early and one sufficiently late.

[18]This ordering is preserved in the beliefs we observe in the data, so we view this assumption as essentially innocuous.

($\hat{p}_w$ decreases) and become more pessimistic about the risk of infection while still at home ($\hat{p}_h$ increases).

The intuition driving the results is quite simple: as the time to return to work draws nearer, the risk in more periods is being driven by $\hat{p}_w$, and the risk in fewer periods is being driven by $\hat{p}_h$. Because individuals discount the future by $\delta < 1$, which leads to a balancing act between belief distortions regarding risks that govern temporally close and distant periods. In particular, as the return date draws nearer, the anticipated utility from future risk is being driven more by $\hat{p}_w$ than $\hat{p}_h$, increasing the marginal return to distorting $\hat{p}_w$ relative to $\hat{p}_h$.[19] Thus, the marginal benefit to distortion increases for $\hat{p}_w$ and falls for $\hat{p}_h$. As a result, individuals become more optimistic about the infection risk at work and reduce their optimism regarding the infection risk at home.

Proposition 1

1. *The perceived chance of getting sick at work less the chance of getting sick at home is smaller for NOW compared to LATER:* $\Delta_{NOW} < \Delta_{LATER}$

2. *The perceived chance of getting sick at work is smaller for NOW compared to LATER:* $\hat{p}_{w,NOW} < \hat{p}_{w,LATER}$

3. *The perceived chance of getting sick at home is larger for NOW compared to LATER:* $\hat{p}_{h,NOW} > \hat{p}_{h,LATER}$

Proof of Proposition 1: We focus on the case of $t = 0$, i.e., how beliefs are distorted in the current period (since this is what we measure). The payoff for determining current period beliefs can be written as

$$\overbrace{\alpha\left[\frac{1}{1-\delta(1-\hat{p}_{h,0})} - \frac{\delta^T(1-\hat{p}_{h,0})^T}{1-\delta(1-\hat{p}_{h,0})} + \frac{\delta^T(1-\hat{p}_{h,0})^T}{1-\delta(1-\hat{p}_{w,0})}\right]}^{\text{Anticipatory Utility}} \underbrace{-c_h(p_h,\hat{p}_{h,0}) - c_w(p_w,\hat{p}_{w,0})}_{\text{Cost of Adjusting Beliefs}}$$
$$+\overbrace{\left[\frac{1}{1-\delta(1-p_h)} - \frac{\delta^T(1-p_h)^T}{1-\delta(1-p_h)} + \frac{\delta^T(1-p_h)^T}{1-\delta(1-p_w)}\right]}^{\text{Realized Utility}}$$

[19] We suspect that in reality additional forces could be at play which would amplify the effects described here, such as selective attention or salience of risks that increase as time to return gets closer. Such considerations would reinforce the effects generated by discounting alone.

Because we focus on $t = 0$ beliefs, we drop the subscript in beliefs indicating time period. The payoff function for $T = 0$ is

$$\alpha[\frac{1}{1-\delta(1-\hat{p}_w)}] + [\frac{1}{1-\delta(1-p_w)}] - c(p_h, \hat{p}_h) - c(p_w, \hat{p}_w)$$

The second order conditions are not necessarily globally negative, so we cannot simply analyze the first order conditions. Instead, our analysis will focus on deriving comparative statics in a way that does not rely on first-order conditions.

Without any additional analysis, because there is no benefit to distorting beliefs about the risk at home, $\hat{p}_{h,NOW} = p_h$. In contrast, $\hat{p}_{w,NOW} \leq p_w$.

As t gets large (for LATER) utility converges to

$$\alpha[\frac{1}{1-\delta(1-\hat{p}_h)}] + [\frac{1}{1-\delta(1-p_h)}] - c(p_h, \hat{p}_h) - c(p_w, \hat{p}_w)$$

We have the opposite situation from before $\hat{p}_{w,LATER} = p_w$ and $\hat{p}_{h,NOW} \leq p_h$. □

Our third result, captured in the second Proposition, is that the perceived probabilities of those who have to return in an intermediate period (LATE) lie between the beliefs of NOW and LATER individuals: $\Delta_{NOW} - \Delta_{LATER} < \Delta_{LATE} - \Delta_{LATER}$. Again, this result is driven by the fact that temporal proximity to return increases the marginal benefit of distorting the infection risk at work, and reduces the marginal benefit of distorting the infection risk at home.

Proposition 2 $\Delta_{NOW} - \Delta_{LATER} < \Delta_{LATE} - \Delta_{LATER}$

Proof of Proposition 2: We will show that $\hat{p}_{w,NOW} \leq \hat{p}_{w,LATE} \leq \hat{p}_{w,LATER}$. The case that $\hat{p}_{h,NOW} \geq \hat{p}_{h,LATE} \geq \hat{p}_{h,LATER}$ follows analogously.

By construction $\hat{p}_{w,LATE} \leq \hat{p}_{w,LATER}$. Thus, we just need to ensure that $\hat{p}_{w,NOW} \leq \hat{p}_{w,LATE}$

Consider the portion of the utility payoff that depends only on the selection of $\hat{p}_w$:

$$\alpha\frac{\delta^T(1-\hat{p}_h)^T}{1-\delta(1-\hat{p}_w)} - c_w(p_w, \hat{p}_w)$$

Observe that the first term represents the benefit of distortion, and the second the cost. Suppose $\hat{p}^*_{w,NOW}$ is optimal at $T = 0$. Then $\alpha\frac{1}{1-\delta(1-\hat{p}^*_{w,NOW})} - c_w(p_w, \hat{p}_w) \geq \alpha\frac{1}{1-\delta(1-\hat{p}_w)} - c_w(p_w, \hat{p}_{w,0})$ for all $\hat{p}_w \in [0, 1]$.

Now consider what happens to payoffs when we move to $T > 0$. Observe that the benefits at all $\hat{p}_w \in [0,1]$ fall, since $\delta^T(1-\hat{p}_h)^T < 1$, but that they fall greatest for lower levels of p_w (one can easily confirm this by looking at the cross partial derivative of $\frac{z}{1-\delta(1-\hat{p}_w)}$ with respect to z and $\hat{p}_w$). Moreover, the costs of distortion do not change. Thus, if $\hat{p}_w < \hat{p}^*_{w,NOW}$, then $\hat{p}_w$ cannot be optimal for any $T > 0$. Therefore, $\hat{p}_{w,NOW} \leq \hat{p}_{w,LATE}$. □

The third proposition highlights stake-dependence of belief adjustments. Individuals naturally differ in how badly they may be affected by the disease. The proposition says that those who experience a larger difference between being healthy and sick will adjust more, because the marginal benefit of adjustment has risen. We call people with a higher flow payoff difference between being healthy and sick "high stakes" (HS), and others "low stakes" (LS) individuals. Formally, we assume (without loss) that the flow payoff to being healthy is $H > 1$ for "high stakes" (HS) individuals. We refer to our original individuals as having "low stakes" (LS).

The result says that we should observe a larger effect when comparing the beliefs of NOW individuals to LATER individuals when we look at HS than LS. Formally, we show that temporal adjustments are larger among people who face higher stakes: $\Delta_{NOW,HS} - \Delta_{LATER,HS} > \Delta_{NOW,LS} - \Delta_{LATER,LS}$. This occurs because a higher utility difference between the desired and undesired state of the world increases the marginal return to adjusting beliefs, whereas the marginal costs remain the same.

Proposition 3 $\Delta_{NOW,HS} - \Delta_{LATER,HS} < \Delta_{NOW,LS} - \Delta_{LATER,LS}$

Proof of Proposition 3: Normalize flow payoff of not being sick for HS people as $H > 1$ so that they face a bigger gap between payoffs. We will prove that $\hat{p}_{w,LATER,HS} = \hat{p}_{w,LATER,LS} > \hat{p}_{w,NOW,LS} > \hat{p}_{w,NOW,HS}$. The analogous results follow for $\hat{p}_h$. The proposition follows since we know that $\hat{p}_{h,LATER,HS} + \hat{p}_{w,NOW,HS} < \hat{p}_{h,LATER,LS} + \hat{p}_{w,NOW,LS}$, which is equivalent to $\Delta_{NOW,HS} - \Delta_{LATER,HS} < \Delta_{NOW,LS} - \Delta_{LATER,LS}$

The utility for LATER does not depend on $\hat{p}_w$, and the maximization problem is equivalent to

$$\alpha \frac{H}{1-\delta(1-\hat{p}_h)} - c(p_h, \hat{p}_h)$$

Thus $\hat{p}_{w,LATER,HS} = \hat{p}_{w,LATER,LS}$.

The payoff for NOW is

$$\alpha \frac{H}{1-\delta(1-\hat{p}_w)} - c_w(p_w, \hat{p}_w)$$

Suppose $\hat{p}^*_{w,NOW,LS}$ is optimal for LS individuals (where the payoff to being health is H). Then

$\alpha\frac{1}{1-\delta(1-\hat{p}^*_{w,NOW,LS})} - c_w(p_w, \hat{p}^*_{w,NOW,LS}) \geq \alpha\frac{1}{1-\delta(1-\hat{p}_w)} - c_w(p_w, \hat{p}_w)$ for all $\hat{p}_w \in [0,1]$.

Now consider what happens to payoffs when we move to payoffs to being healthy to $H' > H$ and so consider HS individuals. Observe that the benefits at all $\hat{p}_w \in [0,1]$ increase, but that they increase greatest for lower levels of p_w (one can easily confirm this by looking at the cross partial derivative of $\frac{z}{1-\delta(1-\hat{p}_w)}$ with respect to z and $\hat{p}_w$). Moreover, the costs of distortion do not change. Thus, if $\hat{p}_w > \hat{p}^*_{w,NOW,LS}$, then $\hat{p}_w$ cannot be optimal at $H' > H$. Therefore, $\hat{p}_{w,NOW,HS} \leq \hat{p}_{w,NOW,LS}$. □

With an additional strengthening of the assumptions on the cost function we can show a final result, that temporal adjustments in $\hat{p}_w$ are more pronounced than adjustments in $\hat{p}_h$: $\hat{p}_{w,LATER} - \hat{p}_{w,NOW} > \hat{p}_{h,NOW} - \hat{p}_{h,LATER}$. Our first additional assumption is that we put more structure on the cost function:

A1: $c_h(\hat{p}_h, p_h) + c_w(\hat{p}_w, p_w) = \gamma(|.5 - p_h|)\frac{k}{2}(p_h - \hat{p}_h)^2 + \gamma(|.5 - p_w|)\frac{k}{2}(p_w - \hat{p}_w)^2$

Second, we suppose that the true probability of getting sick at work is larger than than the true probability of getting sick at home, but is still less than $\frac{1}{2}$.

A2: $\frac{1}{2} > p_w > p_h$

Third, we suppose that γ is sufficiently steep. We focus, as is true in the data, on situations where the distorted beliefs are interior. The intuition here is slightly different than that for the previous propositions. Here, the intuition revolves around shifts in the marginal cost. Our assumptions imply that as p_h or p_w get more extreme (i.e., closer to 0) the marginal cost of distorting them increases. This implies that extreme values of p_h and p_w (i.e., beliefs that are more certain) dampen the effects described in Proposition 1.

Proposition 4 *Suppose A1 and A2 hold, $\gamma'(x)$ is large enough for all $x \in [0,1]$ and both $\hat{p}_h$ and $\hat{p}_w$ are strictly greater than 0 for both NOW and $LATER$. Then the change in beliefs about getting sick at work is larger than the change about getting sick at home between NOW and $LATER$: $\hat{p}_{w,LATER} - \hat{p}_{w,NOW} > \hat{p}_{h,NOW} - \hat{p}_{h,LATER}$*

Proof of Proposition 4: Notice that answering which of the two beliefs moves more boils down to asking the optimal solution for $\hat{p}_w$ at NOW, i.e., the argmax of

$$\alpha\frac{1}{1-\delta(1-\hat{p}_w)} - \gamma(|.5 - p_w|)k(p_w - \hat{p}_w)^2$$

is larger or smaller than the solution for $\hat{p}_h$ at LATER, i.e., the argmax of

$$\alpha \frac{1}{1-\delta(1-\hat{p}_h)} - \gamma(|.5-p_h|)k(p_h-\hat{p}_h)^2.$$

Since we assume $p_h < p_w$, another way of answering our question is to see whether the optimum (with respect to $\hat{p}$) of

$$\alpha \frac{1}{1-\delta(1-\hat{p})} - \gamma(|.5-p|k(p-\hat{p})^2$$

is increasing more than one-for-one or less than one-for-one with p.

First, we note that the first order condition is

$$\gamma(|.5-p|)k(p-\hat{p}) - \alpha \frac{\delta}{(1-\delta(1-\hat{p}))^2} = 0$$

which is a third-order polynomial in $\hat{p}$. Moreover, at $\hat{p} = p$ the first order condition is negative. This immediately implies that there are at most two local maxima for $\hat{p}$ between 1 and p; up to one interior one, and a potential edge case at $\hat{p} = 1$.

By assumption we rule out the edge case, so we just need to consider the interior solution which satisfies the first order condition. Thus, we just need to ask how any solution to the equation[20]

$$\gamma(|.5-p|)k(p-\hat{p}) - \alpha \frac{\delta}{(1-\delta(1-\hat{p}))^2} = 0$$

changes with p (with $p < \frac{1}{2}$); and in particular if it increases more or less than one for one. Taking the derivative of the FOC with respect to p gives

$$\gamma(|.5-p|)k - \gamma'(|.5-p|)k(p-\hat{p}).$$

Taking the derivative of the FOC with respect to $\hat{p}$ gives

$$-\gamma(|.5-p|)k + \alpha \frac{2(1-\delta(1-\hat{p}))\delta^2}{(1-\delta(1-\hat{p}))^3}.$$

Thus, the change in $\hat{p}$ with respect to p is

[20]Note that although there is only one interior maximum, there may also be an interior minimum (i.e., there may be two local extrema, one a maximum and one a minimum). Our result shows that in both cases the comparative static result holds.

$$\frac{\gamma(|.5-p|)k-\gamma'(|.5-p|)k(p-\hat{p})}{\gamma(|.5-p|)k-\alpha\frac{2(1-\delta(1-\hat{p}))\delta^2}{(1-\delta(1-\hat{p}))^3}}.$$

The first term in the numerator and denominator are the same. So the question becomes whether the second term in the numerator is larger or smaller (in absolute value) than the second term in the denominator. So long as γ' is large enough, the former is larger (in absolute value) than the latter. Therefore $\hat{p}$ increases less than one for one with p. This means that $p_h - \hat{p}_h$ is smaller than $p_w - \hat{p}_w$ since $p_h < p_w$.$\square$

Thus, under the conditions stated in the proposition we observe work beliefs shifting more than home beliefs. If we assumed that γ was a constant (so that γ' was 0) the result would not be true, we would in fact observe the opposite: home beliefs shifting more than work beliefs. This is because the benefit function embeds some convexities with respect to decreases in perceived probabilities (i.e., consider an individual who has to return to work NOW; the marginal benefit of a .1 reduction in perceived probabilities is larger if p_w is .2 compared to if it is 1). There is some pre-existing evidence to support our assumptions: Loewenstein (1985), Engelmann et al. (2019) and Sloman et al. (2010) find evidence that more extreme beliefs impede motivated reasoning.

3 Empirical Investigation

3.1 Field Survey

We conducted a series of surveys during the second, third, and fourth week of May 2020 via Lucid, a marketplace that recruits survey participants from several different online panels. The target group for our survey are working individuals in the U.S. who were furloughed or working from home at the time of the survey. A pre-screening survey administered by Lucid made sure that only qualified individuals who specified that they (1) "worked outside the home before and working from home now (or furloughed)", and (2) "have no choice but to return to the workplace" were admitted to our survey.[21] These requirements guarantee that our respondents are not yet back to work at the time we survey them, and that they will have to return to their workplaces at some point in the future. The survey itself included employment status verification and attention checks

[21] Individuals who indicated that they worked from home before, are unemployed, or who kept working outside the home during the pandemic did not qualify. Individuals who thought they would have a choice to decide on when to return to workplace, who thought they may not return to work, or who had the option to continue working from home also did not qualify. Of the 3,904 respondents who qualified, 6 reported invalid zipcodes and 21 had already tested positive for COVID-19, and are therefore excluded from the analysis.

at the beginning to disqualify any remaining respondents who were not eligible to participate or were not paying sufficient attention to the survey questions. A total of 3,877 individuals completed the survey.

We obtained demographic information about each respondent from Lucid, including their age, gender, ethnicity, education, household income, and political leaning. The survey also asked respondents to indicate their zipcode, which we used to match geographic data such as population density, COVID-19 cases and deaths by county, state identifiers, and county-level restrictions on mobility and economic activity. We report summary statistics of these descriptors in Tables A.1-A.4 of the Supplemental Appendix. As reflected by the first two of these tables, our sample draws from a diverse population in terms of race, education, age, income, and political attitudes, and has an over-representation of women.

After presenting the consent form, and a few general questions about the health and economic impact the pandemic already had on the respondent, the survey asked respondents to describe their jobs and industry they work in, as well as the date of when the industry is being allowed to reopen (next week, in two weeks, etc.). The survey also elicited when respondents expected to return to work.[22] We denote state mandated reopening dates with κ and an individual's expected return date with T. Among the 3,877 respondents, 3,172 expected to return within a week of the reopening date (T and κ are separated by less than a week), and 705 expected to go back with a delay of more than a week between κ and T (70% of which expected to return within a month of the reopening date). A temporal gap between the reopening of the industry and workers' expectations may be due to exogenous reasons, such as less need for employees (e.g., curbside pickup only) or capacity restrictions (e.g., only 3 employees at a time), or for endogenous reasons, such as when a workplace cannot be made safe for high-risk individuals. If a participant indicated that they expect to return more than a week after the industry reopening date, the survey also asked an open ended question about the reason for the delay. Answers were coded by independent research assistants so that we could determine whether the reason is exogenous to risk or risk perceptions. Among the 705 individuals who expected to return with a delay, 140 of them mentioned an exogenous reason, which gives us a total sample of 3,312 respondents.

The survey proceeded to elicit participants' counterfactual beliefs about getting infected with the virus under two scenarios:

1. **Probability of getting sick while staying at home**: What are the chances that you will

[22] Almost half of the respondents in our sample expected to return in 4 weeks or later, ranging from 40%-52% across survey waves. The percentage of respondents who expected to return within a week ranges from 18-39%.

get infected with the coronavirus in the next three months if your current living/working conditions did not change? [0%, 1-10%, . . . , 91-100%]

2. **Probability of getting sick after returning to work**: What would your infection chance be if you went back to working outside the home next week? [0%, 1-10%, . . . , 91-100%]

The answers to these questions correspond to respondents' $\hat{p}_h$ and $\hat{p}_w$, respectively. Crucially, the time frame for the questions was kept the same for all participants. Thus, even though there is variation in the time before respondents return to their workplaces, they all have to consider the same time horizon when answering those questions (e.g., risk of infection if returning *next week*). This feature allows us to infer the impact of time left before returning to work on individual infection beliefs without confounds about the time horizon for which beliefs are elicited. The top panel in Figure 1 plots the distribution of $\hat{p}_h$ and $\hat{p}_w$. Whereas a majority of participants thought the infection risk when staying at home is 1-10%, the most common response for the infection risk after returning to the workplace is 41-50%. We take the middle points of the elicited ranges to construct $\hat{p}_w$ and $\hat{p}_h$. The resulting average infection risk belief associated with staying at home is 17.7% and the same associated with returning to work is 38.8%, which are largely consistent with other papers studying the same time period (e.g., Fan et al., 2020 and Heffetz and Ishai, 2021). Figure A.1 in Supplemental Appendix A shows the distribution of Δ, the difference between $\hat{p}_w$ and $\hat{p}_h$. As evident from this figure, most respondents associated a higher risk of infection with returning to the workplace. The median participant believes that the infection risk will be 20 percentage points larger if they return to the workplace compared to staying at home — reflecting an odds ratio of approximately two between the risk associated with staying at home and the risk after returning to the workplace. This ratio is in line with much of the evidence regarding relative risks at that time period.[23]

The base level beliefs (i.e., $\hat{p}_h$ and $\hat{p}_w$) seem high at first glance, especially compared to the 1% infection rate suggested by the 3,577,132 cases recorded between May 1 and July 31st, 2020. However, it would be misleading to rely on data regarding case counts to predict p_h or p_w. Not only do overall cases not inform us about the risk of being infected in h versus in w, but also, the number of actual infections are known to be at least 10 times under-counted due to both lack of testing and asymptomatic infections (e.g., see Anand et al., 2020; Rosenberg et al., 2020; Sood

[23] Fisher et al. (2020) finds that approximately twice the number of positive COVID tests are acquired at work compared to at home, news reports said that individuals have twice the chance of catching COVID at work (e.g., `https://www.yahoo.com/lifestyle/going-every-day-doubles-chances-113444265.html`) and Herstein et al. (2021) finds that meat-packers had between 2.8 and 15 times the chance of catching COVID relative to surrounding community members. Mulligan (2021) provides a recent survey of some results.

et al., 2020 for seroprevalence studies in that time period). Also, at the time of our study, several salient news reports cited studies that found higher infection rates than what was suggested by case counts, (e.g., about 20% among meat-packers, Herstein et al., 2021; 14% among pregnant women in New York City, Sutton et al., 2020; about 8.5% of a national sample of dialysis patients, Anand et al., 2020). Therefore, these estimates may not be gross overestimates. Additionally, note that individuals generally tend to overestimate (by orders of magnitude) low probability events (see, e.g., Gerber et al., 2020 in the context of gubernatorial elections and Barseghyan et al., 2013 on insurance choice), as consistent with predictions of well-known formulations of probability weighting.

Figure 1: Distribution of beliefs about infection risk and severe health outcomes

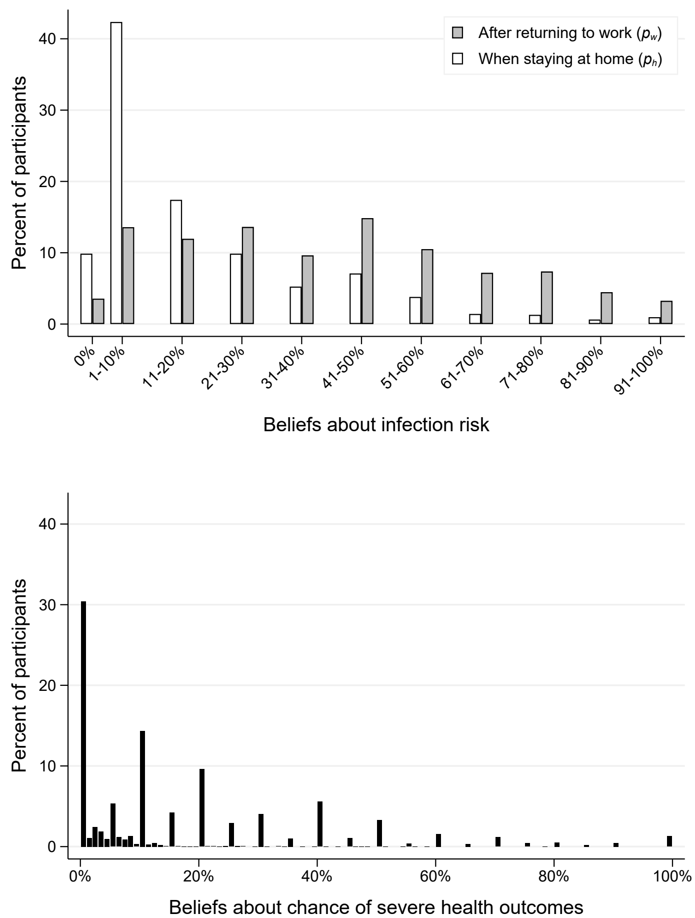

Notes: The top panel of this figure provides a histogram of the beliefs about the infection risk after returning to the workplace ($\hat{p}_w$) and the infection risk associated with staying at home ($\hat{p}_h$). The bottom panel shows a histogram of the beliefs about the chance to be hospitalized or needing treatment in the ICU conditional on getting infected.

Respondents reported whether they have certain medical conditions (e.g., asthma or chronic lung disease) that the CDC has identified as conditions that increase the likelihood of severe outcomes. A total of 32% of participants report having a health condition that may put them at a higher risk for severe illness from COVID-19. The top drivers are cardiovascular disease, diabetes, respiratory disease, hypertension, and cancer. When combined with people who are 60 years of age or older, 40% of our respondents are classified as high-risk. Based on this list of conditions, CDC estimated that about 45% of the U.S. population may have heightened risks.[24]

The survey also elicited beliefs regarding the potential severity of disease if they contract the virus. In particular, they indicated the chances they believe they would be symptom free, have a mild version of the disease, have a more severe version without hospitalization, would be hospitalized without further interventions, and the chances they would require intensive care. In our analyses, we operationalize adverse outcome expectations as the probability of hospitalization and/or intensive care. The distribution of beliefs about the probability of severe health outcomes is plotted at the bottom of Figure 1. On average, participants believed there is about a 16.5% chance that they would end up in hospital (or worse) if they get infected. Yet, roughly a third of participants thought that the chance of developing severe outcomes is close to zero. As with infection rates, it is difficult to calculate the rate of hospitalization. Based on models developed by Reese et al. (2021) and Iuliano et al. (2021), the CDC estimates 18% hospitalization rate among those who are 65 and older, 7% for 50-64 year-olds and 2.7% for 18-49 year olds.[25] Also, the CDC announced that individuals with underlying conditions are more likely to be hospitalized than otherwise healthy individuals.[26] Congruently, those who were categorized as high-risk by the CDC guidelines reported higher expectations of hospitalization and/or ICU care (see Table A.6 in the Supplemental Appendix).

We also asked participants about their expectations of the number of COVID-19 deaths in the U.S. as well as the impact of the pandemic on the economy (unemployment rate, GDP growth). These general beliefs allow us to perform a placebo test for whether changes in beliefs are due to differences in people's knowledge of the current situation of the pandemic. Table A.5 summarizes all elicited beliefs about individual risks and expectations of national outcomes, and Tables A.6 and A.7 report how these beliefs vary with various demographics It is reassuring to see that patterns of belief variations by age, gender, income, and political attitudes are similar to those documented

[24]See Adams et al. (2020) for details.

[25]Source: Table 2 in https://www.cdc.gov/coronavirus/2019-ncov/cases-updates/burden.html In other studies, higher hospitalization rates have been reported. For example, Jehi et al. (2020) reports an overall 21% rate among patients who tested positive at a Cleveland clinic between March and June 2020.

[26]The CDC report in early May, 2020 can be found here: https://www.cdc.gov/coronavirus/2019-ncov/covid-data/covidview/past-reports/05082020.html

by other surveys in the same time period. For example, we see that older people are more worried about the severity of the disease, even though they estimate their risk of infection to be lower (as in Bordalo et al., 2020).[27] Also, women and those who consider their political view to be on the liberal side are more pessimistic about infection risks, consistent with these groups being more worried about the pandemic in general (as in Allcott et al., 2020; Bundorf et al., 2021; Galasso et al., 2020; Fan et al., 2020). Supplemental Appendix A provides further details on these analyses.

The end of the survey included a other questions that do not inform our main analyses.[28] The survey concluded by asking about the respondents' news consumption habits and presenting an open-ended feedback opportunity. Supplemental Appendix D reproduces the survey.

3.2 Identification Strategy

Recall that in our field setting, all workers eventually have to return to their workplaces, return and survey timing is externally imposed, and thus the *time left* to return varies exogenously across workers. We refer to the time left before people have to go back to the workplace as $\tilde{t} = T - \tau$, where T denotes expected time of return and τ is the time of the survey. The data feature two sources of independent variation in $\tilde{t}$. The first source of variation comes from the staggered nature of the survey: fixing T, the variation in $\tilde{t}$ arises from individuals participating in different waves τ of the survey. The second source comes from the staggered re-opening dates: fixing the timing of the survey τ, we can leverage the fact that re-opening week κ (and, thus, also the return time T) varies exogenously across workers as (i) within states different industries reopened at different points in time, and (ii) across states the same industries reopened on different dates.

To identify how beliefs about the infection risk regarding each state of the world (staying at home vs. going back to work) responds to $\tilde{t}$, the variation in return timing should be independent of an individual's risk of infection in h and in w. In our main analyses, we assume that re-opening week κ is uncorrelated with p_w. Therefore, we take advantage of both the variation generated by the survey waves and by the staggered relaxation of state-mandated restrictions. In follow-up analyses, we address the potential issue that the reopening week κ may serve as a signal about the infection risk associated with going back to the workplace by including return-week fixed effects in

[27] This may be perhaps because older people have been reported to adopt more preventive behaviors (Kim and Crimmins, 2020), or as proposed in Bordalo et al. (2020), the pandemic may have made mortality salient for the first time to young people.

[28] The set of questions varied across waves. They included questions about the preventative measures participants take and plan to take, a measure of information avoidance presented as a choice between reading a short summary of a research article about the potential long-term health effects of COVID-19 or about sleep deprivation, and a hypothetical choice about health insurance. In Supplemental Appendix C.2 we shows that there are no consistent patterns in these outcome variables as a function of time left to return, and we discuss these null findings in Section 4.

our regressions. For those regressions, we rely only on the variation in $\tilde{t}$ generated by the survey waves.

3.3 Testing the role of temporal distance and health stakes on beliefs

Our main specification tests whether beliefs differ systematically with temporal distance to the return date by estimating the following OLS equation:

$$\Delta_{i(\tau)} = \beta_0 + \beta_1 1(\tilde{t}_i \in \{0,1\}) + \beta_2 1(\tilde{t}_i \in \{2,3\}) + \mathbf{X}_i + \mathbf{\Gamma}_{i(\tau)} + \eta_\tau + \varepsilon_{i(\tau)} \tag{1}$$

where i denotes a respondent, (τ) reminds us of the time variation across survey waves, $\Delta_{i(\tau)}$ is the difference between $\hat{p}_w$ and $\hat{p}_h$, and $\tilde{t}_i$ is the number of weeks a respondent expects to stay at home before returning to the workplace. The coefficients of interest, β_1 and β_2, measure the belief differences between respondents who are returning back to work in four weeks or later (reference group, n=1,546) and those who are returning in the same or next week of the survey ($\tilde{t}_i \in \{0,1\}$, n=891), and also those returning back in two to three weeks after the survey ($\tilde{t}_i \in \{2,3\}$, n=875), respectively. We hypothesize that $\beta_1 < 0$ and $\beta_2 < 0$ and $|\beta_1| > |\beta_2|$. Moreover, our model predicts that people will alter $\hat{p}_w$ more than $\hat{p}_h$. We therefore estimate equation (1) separately for these two beliefs as well.

The model includes wave fixed effects η_τ to control for common changes in beliefs and risks over time. It also features individual specific controls $\mathbf{X}_i$, which includes demographics (indicator variables for gender, ethnicity, education level, household income groups, political leaning of the individual) and health risk factors (indicators for age groups, indicator for whether the individual has any comorbidities). These controls capture any belief differences across respondents arising from differences in the demographic characteristics of workers returning to work at different times. The regression also features county-date level controls $\mathbf{\Gamma}_{i(\tau)}$, which includes population density, the number of COVID-19 cases and deaths per million at the time of the survey, the most recent two-week changes in the number of cases and deaths to control for geographic COVID-19 risks. It also includes indicators for interventions (i.e., stay-at-home orders, closures of entertainment venues, closures of restaurants, restrictions on 50+ or 500+ gatherings) at the county-day level to account for changes in policies and accompanying attitudes and communications that may impact beliefs.

Figure 2 plots β_1 and β_2 (and the associated 95% confidence intervals) both where the dependent variable is $\Delta = \hat{p}_w - \hat{p}_h$ (left panel) and where we separately examine changes in $\hat{p}_h$ and $\hat{p}_w$ (right panel). In all panels, the figure displays the results from sequentially expanding the set of controls. First, it presents results from a specification that only includes wave fixed effects and demographic controls, then reports estimates from consecutive specifications that add controls for age and having a health condition that puts the individual at a heightened risk of severe health outcomes, geographic risks (local infection and death rates, population density), and finally, county-level restrictions.

Figure 2: Change in beliefs about differential infection risk as a function of time left to return to the workplace

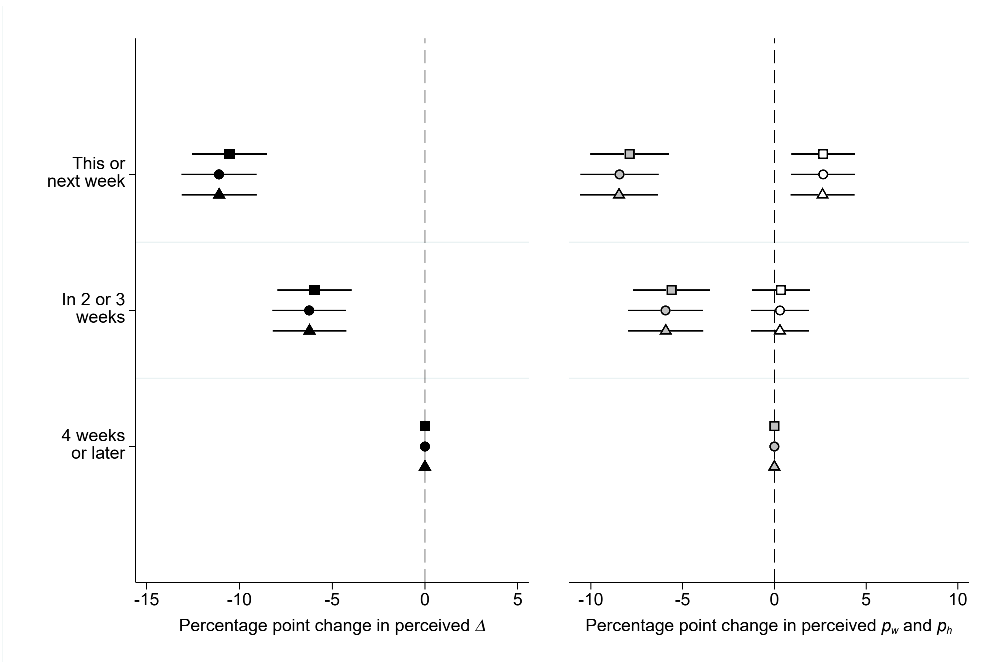

Notes: This figure provides the coefficient values and 95% confidence intervals for the effect of time horizon on the percentage point change in beliefs about infection risk. The black symbols (■, ●, ▲) in the left panel are the estimates for the perceived infection risk differential between returning to work and staying at home (Δ). The right panel presents the results separately for beliefs about the infection risk after returning to the workplace ($\hat{p}_w$) and the infection risk associated with staying at home ($\hat{p}_h$). The grey symbols (■, ●, ▲) are the estimates for $\hat{p}_w$, and the white symbols (□, ○, △) are the estimates for $\hat{p}_h$. The squares indicate estimates from regressions that control for wave fixed effects and demographics. The circles indicate estimates that additionally account for health and geographic risk factors. The triangles represent estimates from regressions with the full set of controls which also includes county-level restrictions.

The estimates are remarkably stable across specifications. Using all controls, we estimate the difference between $\hat{p}_w$ and $\hat{p}_h$ to be about 11 percentage points smaller for people who are about to return to work compared to those who expect to return in four or more weeks ($p < 0.001$). The estimate implies a reduction in the counterfactual belief differential of more than 50 percent, given a median belief differential of 20 percentage points. In other words, individuals who are about to return to the workplace think that the infection risk is more similar between work at home and going back to the office than those who return later. Individuals who expect to return within two or three weeks also adjust their beliefs, but to a lesser extent than those who return sooner: the belief differential is about 6 percentage points smaller for this group compared to those who return at a later date ($p < 0.001$). The coefficients are also reported in the first three columns of Table 1.

The right panel of Figure 2 depicts how $\hat{p}_w$ and $\hat{p}_h$ change across time to return. As people get temporally closer to returning to work, they become more optimistic about the infection risk after returning to work (with all controls, $\beta_1 = -8.5$, $\beta_2 = -5.9$; both $p < 0.001$). On the other hand, they become less optimistic about the infection risk while staying at home (with all controls, $\beta_1 = 2.6$, $p = 0.003$, $\beta_2 = 0.3$; $p = 0.705$). Also, note that across all specifications, the belief adjustment in absolute terms is more pronounced for $\hat{p}_w$ than $\hat{p}_h$ (with all controls: $\left|\beta_{1,\hat{p}_w}\right| - \left|\beta_{1,\hat{p}_h}\right| = 5.8$, $p = 0.001$; $\left|\beta_{2,\hat{p}_w}\right| - \left|\beta_{2,\hat{p}_h}\right| = 5.6$, $p < 0.001$), which is consistent with the prediction of an extension of our model that assumes that the costs of adjusting beliefs are higher when beliefs are more certain (i.e., closer to zero). These results provide support for our hypotheses regarding how the temporal distance to work return affects the degree of belief adjustment.

Next, we empirically test for heterogeneity in temporal belief differences as a function of health stakes, by estimating the following OLS equation:

$$\Delta_{i(\tau)} = \beta_0 + \beta_1 1(\tilde{t}_i \in \{0,1\}) + \beta_2 1(\tilde{t}_i \in \{2,3\}) + \beta_3 m_i + \beta_4 1(\tilde{t}_i \in \{0,1\}) \cdot m_i + \beta_5 1(\tilde{t}_i \in \{2,3\}) \cdot m_i + \mathbf{X}_i + \mathbf{\Gamma}_{i(\tau)} + \eta_\tau + \varepsilon_{i(\tau)} \quad (2)$$

where m_i indicates whether a respondent is older than 60 or has any of the health conditions identified by the CDC as a risk factor. Recall that 40% of our sample fall into this category. Other notations are the same as in equation (1).

Table 1 reports results from this specification in Panel B, below the results from the main specification for comparison. Consistent with predictions, we find that belief adjustments are more pronounced for people who face a higher difference in payoffs between staying healthy and getting infected. Among the people who face higher health stakes, comparing respondents who are about to go back to those who have at least a month to do so, we find a belief differential of 13.5 percentage points (Column 3, with all controls $\mathbf{X}_i$ and $\mathbf{\Gamma}_{i(\tau)}$). Among those who are young and healthy, this difference is only 9.4 percentage points. Thus, the difference-in-differences between high- and low-stakes individuals is 4.1 percentage points ($p = 0.044$). The difference-in-differences is close to zero and not significant when comparing respondents who return to work in two or more weeks to those who return later (p = 0.865).

The survey also elicited beliefs regarding severe illness conditional on infection. Congruent with motivated optimism, we find directional evidence for individuals downplaying the possibility of severely negative health outcomes (i.e., being hospitalized or admitted to the ICU) as they get closer to going back to the workplace. In particular, individuals who are about to return assign a 1.4 percentage points lower probability to severely negative outcomes compared to those who have a month or longer (p = 0.089, Column 9). This implies a 8.5 percent reduction from an average baseline belief of about a 17% chance. Once again, belief adjustments are more pronounced among respondents who face greater health risks should they become infected. This group of individuals assigns a 4.16 percentage points lower probability to severely negative outcomes whereas those who are young and healthy slightly increase their subjective risk assessment by 0.5 percentage points (a difference of 4.66 percentage points, p = 0.010). This difference is smaller and not significant when comparing respondents who return to work in two or more weeks to those who return later (p = 0.197).

Overall, our results suggest that temporal belief adjustments are larger for people who have more at stake, and these differences arise when the temporal distance to the return date is very short. Supplemental Appendix C.1 reports additional results from tests of moderation by gender, risk aversion, education, household income, the impact of the pandemic on the household's income, individual's news consumption patterns, and political views. We do not find differences in belief adjustments across any of these other factors.

3.4 Robustness

It is natural to wonder whether there is an alternative explanation for these belief differences — perhaps some kind of learning occurred over the course of our experiment which was correlated with time to return. For example, it could be that people perceived the reopening of an industry as a signal of its relative safety (i.e., individuals infer that safer industries re-open earlier), generating an association between $\hat{p}_w$ and κ not due to motivated optimism, but to a channel of inference. In addressing this concern, we first note that a signal of an industry's relative safety should not influence individual's perceptions of the risk they face at home or the chance that they become severely ill, and therefore it cannot explain the entirety of our results. Second, we show that our main results replicate after eliminating the variation that arises from different reopening dates (by including reopening week fixed effects): Columns 4 to 6 of Table 1 report these results. These regressions only use the temporal variation across survey waves and therefore the coefficient

Table 1: Impact of time horizon and health stakes on beliefs about infection risk and severity of health outcomes.

	Belief Differental (Δ)						Chance of severe symptoms					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel (a) Main Effects												
This or next week	-10.541	-11.101	-11.093	-6.873	-7.577	-7.607	-1.251	-1.343	-1.425	-0.067	-0.403	-0.417
	(1.030)	(1.034)	(1.032)	(1.443)	(1.441)	(1.438)	(0.880)	(0.834)	(0.837)	(1.225)	(1.158)	(1.158)
Two or three weeks	-5.954	-6.239	-6.223	-5.259	-5.494	-5.484	0.093	0.087	0.055	-0.112	0.024	0.010
	(1.021)	(1.014)	(1.014)	(1.089)	(1.085)	(1.083)	(0.918)	(0.849)	(0.850)	(0.999)	(0.930)	(0.931)
Panel (b) Moderation												
This or next week	-9.144	-9.361	-9.387	-5.314	-5.690	-5.759	0.321	0.600	0.500	1.572	1.562	1.530
	(1.290)	(1.288)	(1.286)	(1.648)	(1.643)	(1.641)	(0.845)	(0.853)	(0.856)	(1.211)	(1.210)	(1.210)
Two or three weeks	-6.112	-6.368	-6.378	-5.311	-5.554	-5.573	0.936	1.029	1.002	0.950	0.986	0.977
	(1.291)	(1.290)	(1.289)	(1.349)	(1.345)	(1.343)	(0.815)	(0.819)	(0.821)	(0.890)	(0.893)	(0.894)
High Stakes	3.108	-2.274	-2.485	3.194	-1.934	-2.146	15.704	-3.523	-3.555	15.720	-3.440	-3.464
	(1.296)	(2.509)	(2.511)	(1.291)	(2.486)	(2.489)	(1.190)	(3.305)	(3.306)	(1.191)	(3.298)	(3.300)
This or next week × High Stakes	-3.496	-4.244	-4.147	-3.811	-4.561	-4.461	-4.192	-4.709	-4.657	-4.251	-4.748	-4.705
	(2.076)	(2.057)	(2.057)	(2.076)	(2.055)	(2.056)	(1.840)	(1.811)	(1.811)	(1.840)	(1.811)	(1.811)
Two or three weeks × High Stakes	0.369	0.272	0.343	0.164	0.071	0.144	-2.095	-2.409	-2.420	-2.214	-2.514	-2.525
	(2.041)	(2.021)	(2.022)	(2.034)	(2.013)	(2.015)	(1.932)	(1.874)	(1.877)	(1.936)	(1.877)	(1.880)
Controls												
Wave FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Demographics	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Health Risk Factors	-	✓	✓	-	✓	✓	-	✓	✓	-	✓	✓
Geographic Risk Factors	-	✓	✓	-	✓	✓	-	✓	✓	-	✓	✓
County Interventions	-	-	✓	-	-	✓	-	-	✓	-	-	✓
Return Week FE	-	-	-	✓	✓	✓	-	-	-	✓	✓	✓
N	3312	3312	3312	3312	3312	3312	3312	3312	3312	3312	3312	3312

Note: The dependent variable in Columns (1)-(6) is the difference between subjective beliefs about infection risks associated with staying at home and infection risks associated with returning to the workplace. The dependent variable in Columns (7) - (12) is the subjective probability of being sick enough to be hospitalized or needing treatment in the ICU conditional on getting infected. Panel (a) reports estimated coefficients (and associated standard errors) from Equation 1 and Panel (b) reports results from Equation 2, staggering different sets of control variables and fixed effects.

estimates are smaller in magnitude, but our conclusions are materially unchanged.

Another concern may be that people who are about to return to work pay more attention to the news or other information about COVID-19. Acquiring more information, by itself, is unlikely to be a driver for the totality of our empirical results for at least two reasons. First, we observe that people become more optimistic about the infection risk associated with going back to work, but less optimistic about the infection risk of staying at home. If initial beliefs were systematically biased (either positive or negative), we would expect home and work beliefs to move in the same direction as individuals get more informed. It seems unlikely that people acquire information that would lead them to simultaneously believe the infection risk associated with staying at home is worse and the infection risk associated with going back to work is better than they previously thought.[29] Second, if acquiring more information is making individuals more knowledgeable, either about infection risks or the pandemic, we would expect to see similar shifts in beliefs regarding the impact of the pandemic on aggregate outcomes (i.e., total number of COVID-19 deaths, unemployment rate, GDP growth). In contrast, if individuals manipulate their beliefs to manage their anticipatory emotions regarding personal risks, we would expect differences in those beliefs, but not beliefs regarding systemic risks. In the survey, we asked respondents about their expectations of how many people in the U.S. will die from COVID-19 by the beginning of July 2020, which is roughly one month after they completed the survey. The average expectation was 133,800 deaths. We also elicited beliefs about how the pandemic will affect the U.S. economy by asking respondents about their expectations of GDP growth in 2020 as well as the rate of unemployment by the beginning of July 2020. The average respondent expected a 15.4% unemployment rate (the realization was 10.2%) and a 1.67% GDP growth (the realization was -3.5%). Detailed statistics for all elicited beliefs are summarized in Table A.5 in the Supplemental Appendix.

We assess the plausibility of an information-based explanation by examining temporal differences in general beliefs about the impact of the pandemic using the same regression approach as before. Whether these beliefs differ across temporal distance to return serves as a placebo test for our theory. Table 2 presents the results. The even columns display the results from estimating a specification that additionally includes reopening-week fixed effects. Consistent with our hypotheses, we do not see systematic differences in placebo beliefs as a function of $\tilde{t}$. The results indicate that respondents who are about to go back to work and those who return in four or more weeks have very similar expectations about the impact of the pandemic on GDP growth ($p = 0.434$), the

[29] Another information exposure difference may arise from people receiving reassuring information from their employers (e.g,. about new safety measures at the workplace) shortly before they go back to work. However, such information cannot explain the change in risk beliefs associated with staying at home.

Table 2: Differences in placebo beliefs across time horizons to work return

	GDP		COVID-19 Deaths		Unemployment Rate	
	(1)	(2)	(3)	(4)	(5)	(6)
Main Effects						
This or next week	0.727	0.029	0.285	-0.134	-0.018	-0.044
	(0.929)	(1.371)	(2.688)	(3.637)	(0.486)	(0.670)
Two or three weeks	0.903	0.589	-5.049	-7.644	-0.974	-0.695
	(0.872)	(0.931)	(2.450)	(2.577)	(0.466)	(0.505)
Controls						
Wave FE	✓	✓	✓	✓	✓	✓
Demographics	✓	✓	✓	✓	✓	✓
Health Risk Factors	✓	✓	✓	✓	✓	✓
Geographic Risk Factors	✓	✓	✓	✓	✓	✓
County Interventions	✓	✓	✓	✓	✓	✓
Return Week FE	-	✓	-	✓	-	✓
N	3312	3312	3312	3312	3312	3312

Note: Each column reports estimated coefficients (and associated standard errors) from a regression of the dependent variable indicated in column heading on the main regression specification plus indicated controls. Beliefs about GDP and the unemployment rate are measured in percentage points; beliefs about COVID-19 deaths are reported in thousands.

total number of COVID-19 related deaths (p = 0.915), and the unemployment rate (p = 0.971). Interestingly, respondents who return to work within two or three weeks have generally more optimistic beliefs than either of the other two groups. For example, shown in column 3, they expect about 5,000 fewer COVID-19 related deaths compared to respondents who return to work in four weeks or later (p = 0.039). Despite some differences in general beliefs across groups of respondents with different horizons to work return, we do not find that respondents who have to return immediately have the most optimistic beliefs. Thus, the overall pattern is inconsistent with an information-based explanation.

In summary, an alternative hypothesis based on information acquisition would have a hard time explaining the overall pattern in the data, which includes: (i) beliefs about workplace infection becoming more optimistic, (ii) beliefs about home infection becoming more pessimistic (and so move in the opposite direction), (iii) belief movements are mediated by pre-existing health concerns, (iv) a reduction in beliefs about severity of infection (conditional on catching COVID19), and (v) no systematic movement in placebo beliefs. The survey also asked individuals to list news sources individuals usually rely on. We check whether the number of news sources people consume responds to the time horizon before work-return. Time trends in news source choices do not indicate an increase in the variety of news consumed. We also check whether people who consumer more news, or certain types of news (e.g., Fox, NYT) show larger temporal adjustments, but do not find a mediating relationship of the type of the number of news sources on the degree of temporal belief differences.

In Supplemental Appendix B.1, we also present results from analyses that show the robustness of our results to (i) a log-likelihood transformation of beliefs to account for their truncated nature, (ii) accounting for outliers in dependent variables by winsorizing 5% of the tails, (iii) including all participants who were admitted to our survey, and excluding (iv) 140 respondents who return with a delay due to exogenous reasons, (v) 248 respondents who suspect they have been infected with the virus, and (6) 934 respondents who are not sure about when the state will reopen their industry and therefore are placed in the LATE group. Overall, our conclusions do not change based on those robustness checks.

4 Conclusion

Using a natural experiment, we test the idea that people manage their anticipatory emotions by adjusting their beliefs over time. After the initial wave of COVID-19 lockdowns in the spring of 2020, many workers had to return to their workplaces despite an elevated risk of infection compared to staying at home. We conducted a series of surveys among workers who were required to return to their workplaces and find that they became more optimistic about the risk on infection at the workplace and less optimistic about the risk at home as they approached their return date. These belief distortions were larger for people who were expected to face more adverse health outcomes if infected.

Our paper, like others that document belief distortions, raises the normative question of whether these beliefs are "optimal" in a welfare sense. In other words, is there potentially scope for interventions that would improve the well-being of individuals who motivate their beliefs? One can construct models to support either yes or no. We know of no work, including our own, which allows to distinguish between normative approaches.[30] Our evidence can weigh in on one particular formulation of distorted beliefs — that they are mistakes which can be reduced via costly effort or attention. We find that individuals who have the most to lose from distorting their beliefs are in fact the ones to exhibit the largest time path of distortions. Such a finding goes against this interpretation of motivated beliefs.

Despite the lack of clear welfare criteria in this literature, our results can still inform policy

[30] Of course, even if individually optimal, belief distortions may exacerbate externalities, such a reducing preventative measures, and so preventing them may be socially beneficial.

makers, if they want to intervene, about when and for whom interventions will be most effective. These interventions may try to directly correct beliefs, or simply adjust the environment to help protect individuals from the consequences of their distorted beliefs. For example, it may be best to correct beliefs shortly before actions must be taken, or a situation changes, as this allows for benefits from distorted beliefs prior to this time period, but prevents individuals from acting in utility-reducing ways.[31]

Our findings have also implications for other areas of research. First, distorting beliefs about future risk is just one possible strategy to regulate emotions. A large body of literature in psychology suggests that people can learn to control the impact of their emotions through cognitive reappraisal of a situation (e.g., Gross, 2015). An early example of this in economics is Becker and Rubinstein (2004), who present a model in which individuals can invest in their ability to control their fears and then empirically test it in the context of terrorism. A different literature suggests that individuals may try to control their emotions through selective access to information (e.g., Caplin and Leahy, 2001 make the point theoretically, while Oster et al., 2013 and Ganguly and Tasoff, 2017 provide evidence in the medical domain). One may view these mechanisms as potentially substitutes (because by avoiding information, people reduce the need to distort beliefs) or complements (because avoiding information may make distorting beliefs easier). Although relatively under-powered, our study did not reveal differences in information avoidance across individuals with different time horizons to work (see Table C.6 in Supplemental Appendix C.2). Thus, more research is needed to understand the relative importance and interplay of these distinct emotion regulation strategies.

Second, wage differentials across industries with varying health risks are often used to estimate the value of a statistical life (VSL) (see Viscusi, 2018 for a recent survey). However, using data from the COVID-19 pandemic, recent work by Anelli and Koenig (2021) shows that using subjective beliefs rather the objective probabilities about workplace risks can lead to significantly different estimates of VSL. Our work suggests that the problem may run even deeper: individuals may manipulate their subjective beliefs depending on their circumstances — whether they feel they have to be working in a particular industry, rather than choosing between a risky industry and a

[31] How beliefs influence actions is unclear, both in general and in the context of the pandemic. For example, whether people should take more or fewer precautions depends on their perception of the marginal benefit of such actions, which may or may not rise as the level of infection risk increases. The empirical evidence is also mixed. For example, in a correlational study, Bundorf et al. (2021) document that people who believed they faced higher risks of infection also avoided everyday activities such as grocery shopping and taking public transport. Heffetz and Ishai (2021) find that risk perceptions correlate only with some kinds of activities. In contrast, Akesson et al. (2020) find that concerns about higher risks led to fewer protective actions, potentially due to a countervailing effect related to fatalism. Supplemental Appendix C.2 presents results from our survey – we do not find a meaningful change in preventative actions, hypothetical insurance uptake or the degree to which individuals engage in risky economic activities as a function of time left before individuals return to their workplace.

less risky industry without any current employment. This implies that we may end up with very different estimates of the VSL depending on the particular details of workers' situations. Different VSLs may result not from preference heterogeneity, but rather from the degree of belief distortion about risks across different groups of workers. Understanding how distorted beliefs may impact our estimates of policy-relevant parameters like the VSL is an important, and as of yet, unexplored topic.

The COVID-19 pandemic distributed proverbial lemons in the form of severe health risks to people throughout the world. Our results suggest that individuals, when drawing near to facing a heightened risk, mentally sweeten the lemons, adjusting their beliefs to make the the world seem safer and more palatable. As highlighted by the discussion above, more research is needed to understand the broader impact of such belief distortions.

References

Adams, M. L., D. L. Katz, and J. Grandpre (2020). Population-based estimates of chronic conditions affecting risk for complications from coronavirus disease, united states. *Emerging Infectious Diseases 26*(8), 1831.

Akerlof, G. A. and W. T. Dickens (1982). The economic consequences of cognitive dissonance. *The American Economic Review 72*(3), 307–319.

Akesson, J., S. Ashworth-Hayes, R. Hahn, R. D. Metcalfe, and I. Rasooly (2020). Fatalism, beliefs, and behaviors during the covid-19 pandemic. NBER Working paper, No. 27245.

Allcott, H., L. Boxell, J. Conway, M. Gentzkow, M. Thaler, and D. Yang (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics 191*, 104254.

Anand, S., M. Montez-Rath, J. Han, J. Bozeman, R. Kerschmann, P. Beyer, J. Parsonnet, and G. M. Chertow (2020). Prevalence of SARS-CoV-2 antibodies in a large nationwide sample of patients on dialysis in the USA: A cross-sectional study. *The Lancet 396*(10259), 1335–1344.

Anelli, M. and F. Koenig (2021). Willingness to pay for workplace safety. https://www.andrew.cmu.edu/user/fkoenig/The_willingness_to_pay21jul26.pdf [Accessed: 2021-10-30].

Babcock, L., G. Loewenstein, S. Issacharoff, and C. Camerer (1995). Biased judgments of fairness in bargaining. *The American Economic Review 85*(5), 1337–1343.

Barseghyan, L., F. Molinari, T. O'Donoghue, and J. C. Teitelbaum (2013). The nature of risk preferences: Evidence from https://www.overleaf.com/project/602d702e8db355b8ca5706e0insurance choices. *American Economic Review 103*(6), 2499–2529.

Becker, G. S. and Y. Rubinstein (2004). Fear and the response to terrorism: An economic analysis.

Ben-Horin, D. (1979). Dying to work: Occupational cynicism plagues chemical workers. *In These Times 3*, 24.

Bénabou, R. (2015). The economics of motivated beliefs. *Revue d'économie politique 125*(5), 665–685.

Bénabou, R. and J. Tirole (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics 117*(3), 871–915.

Bénabou, R. and J. Tirole (2011). Identity, morals, and taboos: Beliefs as assets. *The Quarterly Journal of Economics 126*(2), 805–855.

Bénabou, R. and J. Tirole (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives 30*(3), 141–64.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics: Applications and Foundations 1 2*, 69–186.

Benzell, S. G., A. Collis, and C. Nicolaides (2020). Rationing social contact during the COVID-19 pandemic: Transmission risk and social benefits of US locations. *Proceedings of the National Academy of Sciences 117*(26), 14642–14644.

Bordalo, P., K. B. Coffman, N. Gennaioli, and A. Shleifer (2020). Older people are less pessimistic about the health risks of COVID-19. Technical report, NBER Working paper No. 27494.

Bracha, A. and D. J. Brown (2012). Affective decision making: A theory of optimism bias. *Games and Economic Behavior 75*(1), 67–80.

Brunnermeier, M. K. and J. A. Parker (2005). Optimal expectations. *American Economic Review 95*(4), 1092–1118.

Bundorf, M. K., J. DeMatteis, G. Miller, M. Polyakova, J. L. Streeter, and J. Wivagg (2021). Risk Perceptions and Protective Behaviors: Evidence from COVID-19 Pandemic. NBER Working Paper 28741.

Burks, S. V., J. P. Carpenter, L. Goette, and A. Rustichini (2013). Overconfidence and social signalling. *Review of Economic Studies 80*(3), 949–983.

Caplin, A. and J. Leahy (2001). Psychological expected utility theory and anticipatory feelings. *The Quarterly Journal of Economics 116*(1), 55–79.

Caplin, A. and J. V. Leahy (2019). Wishful thinking. NBER Working Paper, No. 25707.

Charness, G., A. Rustichini, and J. Van de Ven (2018). Self-confidence and strategic behavior. *Experimental Economics 21*(1), 72–98.

Chew, S. H., W. Huang, and X. Zhao (2020). Motivated false memory. *Journal of Political Economy 128*(10), 3913–3939.

Coutts, A. (2019). Testing models of belief bias: An experiment. *Games and Economic Behavior 113*, 549–565.

Di Tella, R., S. Galiant, and E. Schargrodsky (2007). The formation of beliefs: evidence from the allocation of land titles to squatters. *The Quarterly Journal of Economics 122*(1), 209–241.

Eil, D. and J. M. Rao (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics 3*(2), 114–38.

Engelmann, J., M. Lebreton, P. Schwardmann, J. J. van der Weele, and L.-A. Chang (2019). Anticipatory anxiety and wishful thinking. Tinbergen Institute Discussion Paper 2019-042/I.

Falk, A., A. Becker, T. Dohmen, B. Enke, D. Huffman, and U. Sunde (2018). Global evidence on economic preferences. *The Quarterly Journal of Economics 133*(4), 1645–1692.

Fan, Y., A. Y. Orhun, and D. Turjeman (2020). Heterogeneous actions, beliefs, constraints and risk tolerance during the COVID-19 pandemic. NBER Working Paper 27211.

Fisher, K. A., S. M. Olson, M. W. Tenforde, L. R. Feldstein, C. J. Lindsell, N. I. Shapiro, D. C. Files, K. W. Gibbs, H. L. Erickson, M. E. Prekker, et al. (2020). Telework before illness onset among symptomatic adults aged ≥ 18 years with and without covid-19 in 11 outpatient health care facilities – united states, july 2020. *Morbidity and Mortality Weekly Report 69*(44), 1648. Centers for Disease Control and Prevention.

Galasso, V., V. Pons, P. Profeta, M. Becher, S. Brouard, and M. Foucault (2020). Gender differences in COVID-19 attitudes and behavior: Panel evidence from eight countries. *PNAS 117*(44), 27285–27291.

Ganguly, A. and J. Tasoff (2017). Fantasy and dread: The demand for information and the consumption utility of the future. *Management Science 63*(12), 4037–4060.

Gerber, A., M. Hoffman, J. Morgan, and C. Raymond (2020). One in a million: Field experiments on perceived closeness of the election and voter turnout. *American Economic Journal: Applied Economics 12*(3), 287–325.

Gödker, K., P. Jiao, and P. Smeets (2020). Investor memory. Available at SSRN 3348315.

Golman, R., D. Hagmann, and G. Loewenstein (2017). Information avoidance. *Journal of Economic Literature 55*(1), 96–135.

Gottlieb, D. (2014). Imperfect memory and choice under risk. *Games and Economic Behavior 85*, 127–158.

Gross, J. J. (2015). Emotion regulation: Current status and future prospects. *Psychological Inquiry 26*(1), 1–26.

Heffetz, O. and G. Ishai (2021). Which Beliefs? Behavior-Predictive Beliefs are Inconsistent with Information-Based Beliefs: Evidence from Covid-19. NBER Working Paper, No. 29452.

Herstein, J. J., A. Degarege, D. Stover, C. Austin, M. M. Schwedhelm, J. V. Lawler, J. J. Lowe, A. K. Ramos, and M. Donahue (2021). Characteristics of SARS-CoV-2 transmission among meat processing workers in Nebraska, USA, and effectiveness of risk mitigation measures. *Emerging infectious diseases 27*(4), 1032. Centers for Disease Control and Prevention.

Hoffman, M. and S. V. Burks (2020). Worker overconfidence: Field evidence and implications for employee turnover and firm profits. *Quantitative Economics 11*(1), 315–348.

Huffman, D., C. Raymond, and J. Shvets (2019). Persistent overconfidence and biased memory: Evidence from managers. Unpublished.

Iuliano, A. D., H. H. Chang, N. N. Patel, R. Threlkel, K. Kniss, J. Reich, M. Steele, A. J. Hall, A. M. Fry, and C. Reed (2021). Estimating under-recognized Covid-19 deaths, United States, March 2020-May 2021 using an excess mortality modelling approach. *The Lancet Regional Health-Americas 1*, 100019.

Jehi, L., X. Ji, A. Milinovich, S. Erzurum, A. Merlino, S. Gordon, J. B. Young, and M. W. Kattan (2020). Development and validation of a model for individualized prediction of hospitalization risk in 4,536 patients with COVID-19. *PloS One 15*(8), e0237419.

Killeen, B. D., J. Y. Wu, K. Shah, A. Zapaishchykova, P. Nikutta, A. Tamhane, S. Chakraborty, J. Wei, T. Gao, M. Thies, and M. Unberath (2020). County-level Socioeconomic Data for Predictive Modeling of Epidemiological Effects. https://github.com/JieYingWu/COVID-19_US_County-level_Summaries.

Kim, J. K. and E. M. Crimmins (2020). How does age affect personal and social reactions to COVID-19: Results from the national Understanding America Study. *PloS One 15*(11), e0241950.

Klinowski, D. and J. S. Paulsen (2013). Impaired awareness in Huntington disease: Medical evidence in relation to the optimal expectations model.

Köszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association 4*(4), 673–707.

Krizan, Z. and P. D. Windschitl (2007). The influence of outcome desirability on optimism. *Psychological Bulletin 133*(1), 95.

Krizan, Z. and P. D. Windschitl (2009). Wishful thinking about the future: Does desire impact optimism? *Social and Personality Psychology Compass 3*(3), 227–243.

Kunda, Z. (1987). Motivated inference: Self-serving generation and evaluation of causal theories. *Journal of Personality and Social Psychology 53*(4), 636.

Kunreuther, H., R. Ginsberg, L. Miller, P. Sagi, P. Slovic, B. Borkan, and N. Katz (1978). *Disaster insurance protection: Public policy lessons.* Wiley New York.

Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *The Economic Journal 97*(387), 666–684.

Loewenstein, G. F. (1985). *Expectations and intertemporal choice.* Ph. D. thesis, Yale University.

Macera, R. (2014). Dynamic beliefs. *Games and Economic Behavior 87*, 1–18.

Malmendier, U. and T. Taylor (2015). On the verges of overconfidence. *Journal of Economic Perspectives 29*(4), 3–8.

Mayraz, G. (2011). Wishful thinking. Available at SSRN 1955644.

Mayraz, G. (2018). Priors and desires: A bayesian model of wishful thinking and cognitive dissonance. Unpublished.

Mijović-Prelec, D. and D. Prelec (2010). Self-deception as self-signalling: a model and experimental evidence. *Philosophical Transactions of the Royal Society B: Biological Sciences 365*(1538), 227–240.

Möbius, M. M., M. Niederle, P. Niehaus, and T. S. Rosenblat (2014). Managing self-confidence. NBER Working paper, No. 17014.

Molnar, A. and G. Loewenstein (2021). Thoughts and Players: An Introduction to Old and New Economic Perspectives on Beliefs. *The Science of Beliefs: A multidisciplinary Approach (provisional title, to be published in October 2021). Cambridge University Press. Edited by Julien Musolino, Joseph Sommer, and Pernille Hemmer*.

Mulligan, C. B. (2021). The backward art of slowing the spread? Congregation efficiencies during COVID-19. NBER Working paper, No. 28737.

Oster, E., I. Shoulson, and E. Dorsey (2013). Optimal expectations and limited medical testing: Evidence from Huntington disease. *American Economic Review 103*(2), 804–30.

Park, Y. J. and L. Santos-Pinto (2010). Overconfidence in tournaments: Evidence from the field. *Theory and Decision 69*(1), 143–166.

Reese, H., A. D. Iuliano, N. N. Patel, S. Garg, L. Kim, B. J. Silk, A. J. Hall, A. Fry, and C. Reed (2021). Estimated incidence of coronavirus disease 2019 (covid-19) illness and hospitalization - united states, february - september 2020. *Clinical Infectious Diseases 72*(12), e1010–e1017.

Rosenberg, E. S., J. M. Tesoriero, E. M. Rosenthal, R. Chung, M. A. Barranco, L. M. Styer, M. M. Parker, S.-Y. J. Leung, J. E. Morne, D. Greene, et al. (2020). Cumulative incidence and diagnosis of SARS-CoV-2 infection in New York. *Annals of Epidemiology 48*, 23–29.

Schwardmann, P. (2019). Motivated health risk denial and preventative health care investments. *Journal of Health Economics 65*, 78–92.

Schwardmann, P., E. Tripodi, and J. J. Van der Weele (2019). Self-persuasion: Evidence from field experiments at two international debating competitions. CESifo Working Paper, No. 7946.

Shepperd, J. A., J. A. Ouellette, and J. K. Fernandez (1996). Abandoning unrealistic optimism: Performance estimates and the temporal proximity of self-relevant feedback. *Journal of Personality and Social Psychology 70*(4), 844.

Sloman, S. A., P. M. Fernbach, and Y. Hagmayer (2010). Self-deception requires vagueness. *Cognition 115*(2), 268–281.

Sood, N., P. Simon, P. Ebner, D. Eichner, J. Reynolds, E. Bendavid, and J. Bhattacharya (2020). Seroprevalence of SARS-CoV-2–specific antibodies among adults in Los Angeles County, California, on April 10-11, 2020. *JAMA 323*(23), 2425–2427.

Sutton, D., K. Fuchs, M. D'alton, and D. Goffman (2020). Universal screening for sars-cov-2 in women admitted for delivery. *New England Journal of Medicine 382*(22), 2163–2164.

Sweeny, K. and Z. Krizan (2013). Sobering up: A quantitative review of temporal declines in expectations. *Psychological Bulletin 139*(3), 702.

Viscusi, W. K. (2018). Best estimate selection bias in the value of a statistical life. *Journal of Benefit-Cost Analysis 9*(2), 205–246.

Weinstein, N. D., S. E. Marcus, and R. P. Moser (2005). Smokers' unrealistic optimism about their risk. *Tobacco Control 14*(1), 55–59.

Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review 110*(2), 337–61.