Cognitive Dissonance in the Use of Scientific Estimates: Evidence from COVID-19 and Implications for Climate Policy

Abigail Peralta*

Jonathan B. Scott†

September 30, 2021

Abstract

Objective information from credible sources is critical for efficient policy that involves uncertain outcomes. This paper presents evidence that scientific projections are under-utilized in decisions involving dynamic risks. We leverage the COVID-19 pandemic, using plausibly exogenous updates to the dominant model on death projections and show revealed mitigation actions were largely driven by contemporary outcomes over scientific predictions. Further, we document behavior consistent with cognitive dissonance: agents favor scientific forecasts when they predict more optimistic outcomes. When taken to the context of climate policy, we demonstrate that these estimates would imply an undervaluing of carbon costs of 50 percent.

JEL Codes: D91; Q58; H23

*Economics Department, Louisiana State University; **aperalta@lsu.edu**.

†Economics Department, Binghamton University; **jbscott@binghamton.edu**.

Any error are our own.

1 Introduction

Credible information is a crucial component to solving complex problems in the realm of policy making. Information gathering is often outsourced to outside agencies made up of experts in the specific area of interest, conducting technical analyses, not usually feasible nor practical for politicians to conduct on their own. The ultimate output may include precise engineering estimates for proposed infrastructure projects, or tax elasticities for use in federal appropriations. This system creates an efficient means of making important decisions and overcomes the injection of ideology into a problem that can otherwise be solved objectively.

One of the most relevant situations in which policy hinges on sound scientific evidence relates to climate change. As current damages are not necessarily indicative of future damages, projecting into the future becomes a necessary task. While damages from an event such as an oil spill are both salient and immediately detectable, many environmental outcomes are slowly evolving, often playing out over many decades. Millner and Ollivier (2016) argue that the human detachment between certain actions and their long-run impacts on the environment often lead us to think abstractly, while generating skewed perceptions of cause and effect. While politicians may have access to more objective information than the average person, their objectives are often rooted in both their own ideological philosophies, and those of their constituents. The result is a population highly skeptical of the role that humans play in modern warming. A Yale study on climate opinions shows that only 57 percent of adults believe that climate change is caused by human activities (Marlon et al., 2020). Kahan, Jenkins-Smith, and Braman (2010) demonstrate that people's perception of the degree of scientific consensus on climate change varies by political preference. Such skepticism of scientific evidence will ultimately hinder any efforts to form efficient climate policy.

Most of the evidence on how people process and utilize scientific information from varying perspectives is anecdotal. Some survey and stated-preference studies have documented evidence that beliefs of individuals often diverge from scientific consensus (Kahan, Jenkins-Smith, and Braman, 2010; Howe et al., 2015; Leiserowitz et al., 2016; Millner and Ollivier, 2016), but identifying a causal relationship and a clear mechanism is a difficult task. This

paper measures the extent to which scientific estimates are ignored by individuals, while instead favoring a myopic belief based on the present state of the world. We study the COVID-19 pandemic, where estimates for future deaths were both salient and widely accessible for decision making. We leverage the rollout of COVID-19 death projections from the Institute for Health Metrics and Evaluation (IHME)—the sole source of estimates used by, and frequently cited by, the White House Coronavirus Task Force under the Trump Administration in 2020. The frequent, but unpredictable, discrete adjustments to the IHME forecasts across all U.S. states provides plausibly exogenous variation to credibly identify how this information was internalized in mitigation decisions. Combining this information with the widely broadcasted data on current deaths allows us to pin down the extent to which one metric was used over the other.

This paper incorporates quasi-experimental variation from both projected and contemporaneous deaths into an empirical model of risk mitigation. To overcome endogeneity from the direct use of broad mitigation measures, such as lockdowns, we examine the ultimate mobility outcomes of these policies.[1] Using data from SafeGraph, we jointly estimate the causal effects of both contemporary deaths and projected deaths on mobility outcomes, at a daily frequency. The ratio of the coefficients on these two variables is interpreted as a measure of the extent to which one metric was used over the other in risk mitigation strategies. To confirm that these two metrics are internalized into choices for dynamic risk mitigation purposes (i.e., to mitigate future risk), we additionally calculate this ratio on lagged variables, with the expectation that dynamic mitigation strategies will have longer run effects on mobility (e.g., through a lockdown). This is in contrast to the contemporaneous effects that may be observed from an individual mitigating personal risk, conditional only on current levels of deaths. By confirming robustness across these alternative specifications, we conclude that the majority of this effect is through the dynamic risk mitigation channel.

Estimates suggest that only 25 percent weight is placed on scientific estimates when making decisions to mitigate the risks of COVID-19. However, the data suggest that this

[1]The use of lockdowns measures as an outcome of responses to COVID deaths may produce biased estimates, as these types of policies are also likely to have a direct effect on the number of deaths. The use of individual mobility is less likely to create reverse causality concerns, as mobility today is not likely to impact deaths today in a significant manner.

effect is not merely a constant. To explore the potential mechanisms driving this result, we test whether the revealed utilization of these scientific forecasts varies based on their relative magnitude.[2] Cognitive dissonance theory (Festinger, 1957) suggests that, when mitigation is costly, agents may adapt their beliefs in favor of the forecast that is most consistent with their preferred, low-mitigating strategy. Results illustrate this point exactly. When the magnitude of a forecast decreases, relative to the current level of deaths, the emphasis placed on the forecast increases. These results suggests that agents strategically select which information to utilize, taking the scientific projection as the preferred forecast when it presents a more optimistic outlook.

To put this analysis of revealed preferences in the COVID context into broader perspective, we explore an analogous model of climate risk mitigation. Specifically, we apply the weighting parameter revealed from the COVID response to a similar model of climate change mitigation, using estimates generated from the Dynamic Integrated Climate-Economy (DICE) model (Nordhaus and Boyer, 2000; Nordhaus, 2008).[3] Weighting both the social cost of carbon estimates (a projection) and the contemporary damages of CO2 (a myopic outlook) in a similar manner to the main analysis allows for the calculation of a "perceived" social cost of carbon, which this paper argues, is the politically feasible carbon tax. This exercise illustrates that, under the same weighting revealed in the context of COVID, the perceived social cost of carbon would undervalue the cost of climate change by as much as 50 percent.[4]

The main result of this paper suggests that scientific projections may not be internalized in the decision making process to the degree to which they should when they represent the best information available. The results have significant implications for policy-making when the solution to a policy issue requires the use of scientific estimates.[5] In the setting of climate policy, when the cost of CO2 is not fully incorporated into carbon intensive activities, climate

[2]By allowing this weight parameter to vary based on the relative magnitude of projections compared to current death levels, this should further isolate responses through the dynamic risk mitigation mechanism. In Section 5.4, we estimate the effects that derive strictly from this differential effect by setting the estimated baseline effect equal to zero.

[3]Specifically, this paper uses the DICE-2016R model (Nordhaus, 2017, 2019).

[4]This paper does not claim that this is the *true* "perceived" social cost of carbon, as the weighting parameter used is directly derived from the response to COVID. This estimate should strictly be used to put the main analysis into a broader perspective.

[5]Oates and Portney (2003) survey the extensive literature on infusion of politics into environmental regulation.

damages will approach inefficient levels. Correcting these market failures hinges on accurate estimates of the social cost of carbon. Even in the case of uncertain climate outcomes (e.g., Millner, Dietz, and Heal, 2013; Heal and Millner, 2014; van den Bremer and van der Ploeg, 2021), the utilization of objective analyses likely dominates from an efficiency perspective, where alternative approaches may rely more on biased beliefs (Millner and Ollivier, 2016).

This paper relates to a large literature on the economics of information, begun by Stigler (1961). Recent work on consumer attention have focused on the cost of information search and the idea that the most salient features of a problem (e.g., most salient attributes of a product) yield the highest attention from consumers (Gabaix and Laibson, 2005; Gabaix, 2014; Sallee, 2014; Bordalo, Gennaioli, and Shleifer, 2016). While the results in this paper suggest that the most salient information component in mitigation efforts is the current state of the world, the focus on dynamic expectations contrasts it from the majority of this literature.

State-based expectations are central to the model proposed in this paper and are a feature in a number of behavioral models of belief formation. The mechanisms explored here broadly relate to the literature on projection bias (Loewenstein, O'Donoghue, and Rabin, 2003; Conlin, O'Donoghue, and Vogelsang, 2007; Busse et al., 2014), anchoring heuristics (Slovic, 1967; Tversky and Kahneman, 1974; Crandall and Graham, 1989; Furnham, 2011), and particularly, cognitive dissonance (Festinger, 1957). Cognitive dissonance theory (Festinger, 1957) argues that individuals seek psychological consistencey, for example, through their actions and their beliefs. Internal inconsistency leads to stressful behavior for agents, often causing them to generate consistency by altering beliefs in a manner that rationalizes their behavior, or by avoiding particular information that leads to this dissonance. This paper presents results consistent with this idea: the evidence suggests that individuals strategically align with the more optimistic projection when making mitigation choices. This strategic choice is consistent with the principle behind confirmation bias (Nickerson, 1998), arguing that individuals place too much weight on information that confirms their prior beliefs.

Of ultimate interest in this paper is the trade-offs between the utilization of reputable scientific information and one's own myopic expectations of events. Belief formation in the area of climate change has been empirically documented to a limited degree. Schlenker and

Taylor (2021) examine the role of temperature expectations in financial securities that depend on future weather. They are able to show that the price of these securities accurately incorporate weather expectations based on climatic model projections. These results differ substantially from those in this paper, which suggest that scientific projections are likely under-emphasized in decision making. The setting also differs; while Schlenker and Taylor (2021) examine financial markets offering strong monetary incentives, decisions in the COVID setting were often politically motivated (Fisman et al., 2021; Durante, Guiso, and Gulino, 2021; Pulejo and Querubín, 2021). This creates a more likely setting for the infusion of ideological preferences into one's beliefs.

The empirical literature on climate beliefs have mainly been based on survey and stated belief studies. The research has shown high geographic variability in climate beliefs across the U.S. (Howe et al., 2015; Leiserowitz et al., 2016; Marlon et al., 2020). While informative, these studies likely suffer from the same issues as those of stated preference studies (Carson et al., 1996; Hanley and Czajkowski, 2019; Mendelsohn, 2019). That is, these studies often rely on answers to hypothetical questions, where respondents' answers may not accurately reflect what their true behavior would be. Proponents of revealed preference methods argue that, when possible, market settings can provide a more accurate depiction of underlying preferences through prices and a consumer's incentive to maximize individual surplus. Similarly, in the context of this paper, political and economic incentives may affect how individuals incorporate different sources of information into their decisions.[6] Thus, as the literature suggests (e.g., Carson et al., 1996), stated and revealed beliefs may not converge in certain situations.

Reducing the cost of climate change down to a socially efficient level means being able to incorporate true damages into the price of carbon intensive activities. However, implementation of a carbon tax has many political challenges. Ideological differences, self-interests, and skepticism to scientific models may all play a role. These factors affect the extent to which individuals and policy makers incorporate objective information into their choices. Reducing the role that these biases play in decisions is necessary for efficient policy-making to take

[6] See Mullainathan and Washington (2009) and Comin and Rode (2015) for the role of cognitive dissonance in political preferences.

place. However, understanding the manner in which objective information is internalized is a critical first step.

The paper proceeds as follows: Section 2 lays out a simple theoretical framework of decision making under different belief types. Section 3 presents the empirical approach used to identify how agents utilize different sources of information in the COVID setting. Section 4 discusses the data used in this paper. Section 5 presents the main estimation results of the paper. Section 6 applies the main results of this paper to an analogous climate policy setting. Section 7 concludes with a brief discussion of the findings in this paper and their implications.

2 Conceptual Framework

This paper begins with a stylized model of a community's decision to mitigate risks associated with environmentally damaging activities. In what directly follows, the basic model defers on the exact mitigation policy, which may include strategies such as carbon pricing or emissions mandates. One may interpret the level of mitigation taken as the extent to which an externality-correcting policy is imposed on agents, such as the level of Pigouvian tax to be implemented.

The decision to mitigate environmental risks depends both on how a community makes use of information on future damages and how they expect their mitigation strategies to affect those baseline beliefs. In the model that follows, let j represent a specific community from the set of possible communities, $1, 2, ...N$. Mitigation strategies are restricted to be made at the community level for convenience, which in practice, may represent regions such as states, countries, or a broader geographic region. One may interpret the decision-maker as a policy-maker for community j who determines mitigation policy implementation. In what follows, we introduce a simple two period model. For simplicity, we assume a discount rate of zero.

When making mitigation decisions, the community conditions its choice at time t on the future state of damages in period t'. Let $d_{jt'}$ represent the total damages in community j at time t' generated by a given externalitiy-producing activity at time $t < t'$. Mitigation today

limits the magnitude of these future damages, but does not directly affect outcomes today. The degree of mitigation taken depends on the policy-maker's expectations of its ultimate effects. Define D_j as cumulative damages between t and t'. That is,

$$D_j = d_{jt} + d_{jt'} \tag{1}$$

In the indeterminate case, only damages are know in the present period, and expectations are formed on the future period. Suppose that there are two approaches to deriving these expectations, through the scientific authority—who form projections of future damages based on all available scientific information—and through a baseline, myopic expectation, which conditions on the present state of damages. The former is assumed to be widely available to decision makers through periodic publication. For simplicity, we assume that the baseline belief is that damages remain constant at current levels. Accordingly, define two types of beliefs: *scientific* and *myopic*.

Let $\theta \in \{science, myopic\}$ be an agent's type, where type *science* takes the scientific projection as the truth, and makes mitigation decisions accordingly. Beliefs of type *myopic* disregard scientific projections and condition their decisions on the current state of damages. The scientific authority conditions its projections on all available information at time t (including their expectation of mitigation decisions), and derives the forecast $\mathbb{E}_t[d_{jt'}]$.

Let $\delta : \mathbb{R} \to \mathbb{R}^+$ be a mapping whose arguments are a mitigation level—defined by the set of real numbers—and responses are mitigation *effects*—defined as the set of nonnegative real numbers. Assume that decision makers take their baseline projections as given by their type, and make their mitigation decisions based on this independent mapping. That is, the function $\delta(m)$, with mitigation level $m \in \mathbb{R}$, determines the degree to which mitigation reduces future damages. Specifically, it holds that $\partial\delta/\partial m \leq 0$.

Dependent on type θ, an agent derives their expectation of cumulative damages as the following.

$$\mathbb{E}_t[D_j|\theta, m] = d_{jt} + \delta(m) \cdot \hat{d}_{jt'}(\theta) \tag{2}$$

where

$$\hat{d}_{jt'}(\theta) = \begin{cases} \mathbb{E}_t[d_{jt'}] & \text{if } \theta = science \\ d_{jt} & \text{if } \theta = myopic \end{cases} \tag{3}$$

This paper focuses on only a binary decision to mitigate (i.e., $m \in \{0, 1\}$). Results under this assumption will generalize to a chosen mitigation intensity on the closed unit interval, $[0, 1]$. When the choice is to mitigate, the action imposes a cost on society. Call this cost of mitigation, c. An agent weighs this cost against the reductions in damages. Let the marginal value of these damages (e.g., the social cost of carbon) be a constant at value $\tau > 0$. Further, the following normalizations are made: $\delta(0) = 1$ and $\delta(1) = \delta \in [0, 1]$. Under this framework, we can derive the following value function for community of type θ.

$$V_i(\theta) = \max \Big\{ -\tau \cdot \hat{d}_{jt'}(\theta),\ -c - \tau \cdot \delta \cdot \hat{d}_{jt'}(\theta) \Big\} \tag{4}$$

where the leftmost term inside the brackets represents the loss incurred when the community does not mitigate, and the rightmost term is the loss under mitigation. From Equation 4, the policy-maker will choose a mitigation strategy when the following inequality holds.

$$\tau \cdot (1 - \delta) \cdot \hat{d}_{jt'}(\theta) - c > 0 \tag{5}$$

Whether Equation 5 holds will depend not only on the expected effect of mitigation, δ, but also on an agent's beliefs about future deaths, defined in Equation 3. Of course, this depends on the agent's belief type, θ. The main interest of this paper is in identifying the distribution of belief types. Section 3, presents the empirical approach used in this paper to

pin down this estimate of interest.

Cognitive Dissonance: The preceding framework presents the mitigation strategy under an exogenous endowed belief type, θ. We now allow types to be endogenously chosen. When an agent considers the efficient strategy for each θ, the cost of mitigation, c, may justify mitigation under one belief, but not in the other. For example, the agent may choose to believe that true damages are the minimum of each metric. The chosen belief type will justify the strategy, and in when choosing non-mitigation choice, will rationalize the avoidance of cost, c. This of course does not circumvent the true external costs to be incurred in future periods as a result of this bias.

When an agent exhibits cognitive dissonance, their belief type will be that which maximizes the expected utility across mitigation strategies—defined by Equation 4. The following describes the adapted value function under these endogenous beliefs.

$$V_i = \max_{\theta} \Big\{ \max \big\{ -\tau \cdot \hat{d}_{jt'}(\theta),\ -c - \tau \cdot \delta \cdot \hat{d}_{jt'}(\theta) \big\} \Big\} \tag{6}$$

3 Empirical Design

3.1 Model

The empirical approach in this paper proceeds in two phases. First, a simplified version of Equation 4 is estimated, which allows us to examine the extent to which communities might place more emphasis on current damages versus projected future damages. This allows us to examine what we interpret as a baseline revealed belief or relative emphasis placed on one potential metric over another in making mitigation decisions. Importantly, the extent to which the estimated parameter is interpreted strictly as beliefs may be confounded by the salience of one metric or other constraints. We do not claim that a general underutilization of scientific forecasts is solely due to one's biased beliefs.

The second phase allows the revealed utilization of each forecast estimate to vary based on their relative magnitude. When scientific projections are pessimistic, implying a lower relative expected utility of mitigation, agents will place more weight on contemporary outcomes

to rationalize their choices.

Taking the model to the COVID-context, the community chooses mitigation, in the form of staying home when the marginal damages abated by doing so exceeds the marginal costs. We measure mitigation on a daily frequency as the share of a 24 hour day the residents of a community stay at home. Section 4 describes how this measure is derived. The agent can use two sources of information to make mitigation choices: current death levels and forecasts from the scientific authority—in this paper, the Institute for Health Metrics and Evaluation (IHME). Let the marginal utility of mitigation for an agent i from state j on day t be described by the following function.

$$u_{ijt} = c_0 + \beta \cdot \left(\alpha_i \cdot \mathbb{E}_t[d_{jt'}] + (1 - \alpha_i) \cdot d_{jt}\right) + \xi_{jt} + \varepsilon_{ijt} \tag{7}$$

where β encompasses both the damages and the mitigation effect (δ and τ in Equation 5). $\mathbb{E}_t[d_{jt'}]$ represents the scientific projection for the next period and d_{jt} represents current damages—the myopic beliefs. c_0 represents nonzero costs of mitigation. ξ_{jt} represents unobserved, state-level utlity, and ε_{ijt} represents unobserved, individual-level utility.

The two terms, $\beta \cdot \alpha_i$ and $\beta \cdot (1 - \alpha_i)$, represent the implied marginal disutilities from an increase in scientifically projected and contemporary damages, respectively. The agent-specific parameter, $\alpha_i \in \{0, 1\}$, parameterizes an agent's belief type. In the framework of Section 2, $\alpha_i = 1$ represents the type who internalize the scientific projection and $\alpha_i = 0$ represents the type who holds myopic beliefs.

Given the high-frequency, daily observations in the data, we extend the model beyond the two-period framework to allow for decisions to be made on expectations over a pre-established time-horizon. We observe daily scientific projections as of a current date over the course of a significant number of days into the future. Therefore, these longer run forecasts can be incorporated into the model. For myopic beliefs, expectations from Equation 3 generalize to expected cumulative damages of $T \cdot d_{jt}$ over the course of a T-day time-horizon. Thus, define the alternative utility representation under a T-day time-horizon as follows.

$$
\begin{aligned}
u_{ijt} &= c_0 + \beta \cdot \left(\alpha_i \cdot \mathbb{E}_t[d_j|T] + (1-\alpha_i) \cdot T d_{jt}\right) + \xi_{jt} + \varepsilon_{ijt} \\
&= c_0 + \beta \cdot \hat{d}_{jt}(\alpha_i) + \xi_{jt} + \varepsilon_{ijt}
\end{aligned} \tag{8}
$$

where $\mathbb{E}_t[d_j|T] = \sum_{s=1}^{T} \mathbb{E}_t[d_{jt+s}]$ and $\hat{d}_{jt}(\alpha_i)$ is defined in a similar matter as Equation 3. In the context of COVID-19 deaths, the chosen time-horizon should be consistent with expectations on how long an infected individual may be contagious for. As the time of infection is estimated to be roughly two weeks, we shouldn't expect mitigation today to affect deaths past this point. Thus, for the paper's main specification we choose $T = 15$, but present estimates for a range of values.

When unobserved individual utility, ε_{ijt}, is an *iid* draw from the Extreme Value, type I distribution, the expected maximum utility derived from i has the following form.

$$
IV_{jt}(\alpha) = log(1 + exp(c_0 + \beta \cdot \hat{d}_{jt}(\alpha) + \xi_{jt})) \tag{9}
$$

This term is often referred to as the agent's inclusive value. Under cognitive dissonance, an agent will endogenously choose α which maximizes their expected utility. Whether $\alpha = 1$ then depends directly on the relative magnitude of $IV(1)$ over $IV(0)$. We parameterize this probability using a logistic regression by estimating the following equation.

$$
Pr(\alpha = 1) = \frac{exp(\tilde{\alpha}_0 + \tilde{\alpha}_1 \cdot IV(1))}{exp(\tilde{\alpha}_0 + \tilde{\alpha}_1 \cdot IV(1)) + exp(\tilde{\alpha}_1 \cdot IV(0))} \tag{10}
$$

Given this framework, we estimate a nested logit model to identify the parameters in Equation 10. For identification, the model leverages both level variation in each measure of damages—contemporary damages and scientific forecasts—in addition to their *relative* values. Their relative values are assessed as the difference $IV(1) - IV(0)$, implicitly included in Equation 10. The observed mitigation shares are parameterized according to the following function.

$$
\begin{aligned}
Pr(mitigate)_{jt} &= Pr(\alpha = 1) \cdot Pr(c_0 + \beta \cdot \hat{d}_{jt}(1) + \xi_{jt} + \varepsilon_{ijt} > 0) \\
&\quad + Pr(\alpha = 0) \cdot Pr(c_0 + \beta \cdot \hat{d}_{jt}(0) + \xi_{jt} + \varepsilon_{ijt} > 0) \\
&= \frac{exp(\tilde{\alpha}_0 + \tilde{\alpha}_1 \cdot IV(1))}{exp(\tilde{\alpha}_0 + \tilde{\alpha}_1 \cdot IV(1)) + exp(\tilde{\alpha}_1 \cdot IV(0))} \cdot \frac{exp(c_0 + \beta \cdot \hat{d}_{jt}(1) + \xi_{jt})}{1 + exp(c_0 + \beta \cdot \hat{d}_{jt}(1) + \xi_{jt})} \\
&\quad + \frac{exp(\tilde{\alpha}_1 \cdot IV(0))}{exp(\tilde{\alpha}_0 + \tilde{\alpha}_1 \cdot IV(1)) + exp(\tilde{\alpha}_1 \cdot IV(0))} \cdot \frac{exp(c_0 + \beta \cdot \hat{d}_{jt}(0) + \xi_{jt})}{1 + exp(c_0 + \beta \cdot \hat{d}_{jt}(0) + \xi_{jt})}
\end{aligned} \tag{11}
$$

3.2 Base framework

As the data begin on January 1, 2020—before the first COVID death was reported—this establishes two different baselines for the outcome variable: pre-COVID deaths and pre-IHME reports. State-level reported deaths have staggered start times, however, all state-level IHME projections report at the same dates. While we are able to exploit variation in deaths across states and over time, identifying variation for projections will mainly derive from the abrupt projection updates, and heterogeneity in these updates across states. The following equation is estimated.

$$
y_{jt} = \omega_0 \cdot projection_{jt} \times post\ report_t + \omega_1 \cdot current\ deaths_{jt} + \mu_t + \gamma_j + \xi_{jt} \tag{12}
$$

where $post\ report_t$ is an indicator for the timing of the first IHME report (common across all states) and $projection_{it}$ represents the projected deaths for state i as of day t, for a pre-established T-day time horizon. The main specification uses $T = 15$, but other horizons are examined in Section 5. If only one projection were released, this value should be expected to evolve smoothly over time. However, regular adjustments to the IHME model produce discrete jumps in its value, as illustrated in Figure 1. $deaths_{jt}$ represents reported deaths for state j on day t. To maintain similar scales for each regressor, regardless of T, we specifically use deaths per 10,000, where projections represent the mean predicted deaths per-day over the T-day horizon.[7] The estimates are also reported with, and without state-

[7] This is in contrast to the specification in Equation 8, which multiplies current deaths by T, rather than

level trends. Finally, the base specification controls for a full set of state and day-of-sample fixed effects—γ_j and μ_t, respectively.

A binary choice model is estimated for the decision to stay home. The parameters are estimated within a logit framework and, as observed at-home shares are always between zero and one, we can linearize this model by inverting the function. Therefore, the outcome variable of interest becomes the logit-inverse share of time spent at home, or $log(\%home) - log(1 - \%home)$.

The direct coefficients of interest in the model are ω_0 and ω_1. These are analogous to $\beta \cdot \alpha_i$ and $\beta \cdot (1 - \alpha_i)$ in Equations 7 and 8. The paper focuses less on the direct interpretation of these coefficients, as interpretation would be contingent on the imposed time-horizon. Of key interest is the derived parameter $\alpha = \omega_0/(\omega_0 + \omega_1)$. All specifications present standard errors clustered at the state-level. Standard errors for α are approximated by the delta method.[8]

Estimates from Equation 12 allow us to infer something about the relative utilization of each measure of damages. This is our *baseline* measure of α. While informative, the interpretation of this estimate is convoluted due to the role of salience in contemporary versus projected damages. Thus, a low estimate of α from this specification may not be strictly due to skepticism of scientific estimates or a myopic bias. However, it should properly reveal how much mitigation efforts hinged on leading scientific estimates, which we believe has its own value. Of particular interest in this paper is whether α changes when one measure of damages becomes more favorable than the other. Thus, our primary specification will predominately leverage the relative differences between the two. In this paper, this relationship is derived from a nested logit specification, presented in Equation 11. In Appendix C we present an alternative approach that estimates a varying α within a standard logit framework, synonymous to Equation 12. This alternative framework allows each damage term to enter as a quadratic such that there is a non-constant marginal effect on each term and, thus, a non-constant derived α.

dividing projections by T. Results for α will be identical, due to scaling.

[8] The asymptotic variance of α is approximated by the following.

$$Var(\hat{\alpha}) \approx \left(\frac{1-\hat{\alpha}}{\hat{\omega}_0 + \hat{\omega}_1}\right)^2 \cdot Var(\hat{\omega}_0) + \left(\frac{\hat{\alpha}}{\hat{\omega}_0 + \hat{\omega}_1}\right)^2 \cdot Var(\hat{\omega}_1) - 2 \cdot \frac{\hat{\alpha} \cdot (1-\hat{\alpha})}{(\hat{\omega}_0 + \hat{\omega}_1)^2} \cdot Cov(\hat{\omega}_0, \hat{\omega}_1)$$

4 Data

This paper estimates the impact of projected and actual COVID-19 deaths on mitigation efforts. Mitigation outcomes are proxied for using metrics of social distancing, mainly, physical mobility data. Thus, the primary data sources used in this paper include (1) geolocation data from SafeGraph Inc., (2) actual COVID-19 death counts from the John Hopkins University (JHU), and (3) projected COVID-19 deaths from the Institute for Health Metrics and Evaluation (IHME). The study period for this paper is the entire year of 2020, where an observation is a state-day. The analysis is conducted on all 50 states, plus the District of Columbia.

SafeGraph data

The ultimate goal is to measure revealed utilization of scientific projections by observing communities' actions in terms of mitigation. At least early on in the pandemic, social distancing, was one of the most obvious types of mitigation efforts. To derive a meaningful metric of social distancing, we make use of anonymized geolocation data produced by SafeGraph Inc. SafeGraph partners with mobile application services that have opt-in consent from users to collect location data. These partnerships allow SafeGraph to track location data from approximately 35 million unique devices in a given month. The company makes these data available for free to researchers as part of their COVID-19 Data Consortium. Access to the data are available to free users at the census block group level. Block groups with fewer than five observations are omitted. Give the unit of observation for the IHME data, we aggregate these data to the state-level, weighting block groups by population levels. Importantly, the data allow us to track mobility at a daily frequency.

The primary mitigation metric used in this paper is the block group level, median home dwelling time, reported by SafeGraph. The main metric of interest in this paper is the share of time at home. The share of time at home is defined as follows.

$$\%\ at\ home_{jt} = \frac{minutes\ at\ home_{it}}{1,440} \tag{13}$$

where $minutes\ at\ home_{it}$ is the home dwelling time in minutes for state j on day t, and $1,440$ is the number of minutes in a day. This measure lets us generalize mitigation beyond a simple binary metric, allowing us to measure responses on an intensive margin.

It is important to point out several limitations of these SafeGraph data. SafeGraph provides, arguably, some of the most comprehensive data on mobility—observed behavior essential for this paper. However, the data are representative only for the subgroup of individuals who 1) own a smartphone and 2) have consented to location data collection. As interest is in measuring the impact of deaths and projected deaths on mitigation, this will only bias the estimates when mitigation responses to these deaths coincide with underrepresented states.[9] Additionally, the data is collected through regular "pings" to the devices, rather than continuously monitored throughout the day. This implies that frequent short trips may be missed in the calculation of time at home. Similarly, this will only bias the estimates when more frequent short trips coincide with mitigation responses to deaths. State fixed effects are included to account for any of these potential confounders that may remain constant within states; for example, if denser populations produce shorter trips and a stronger response to deaths. Finally, the state-level measure of home dwelling time is the average of block group-level, median estimates, produced by SafeGraph. This should not directly bias estimates; although, it produces a slightly different interpretation than if SafeGraph were to present estimates for mean dwelling times. This distinction should only make a meaningful difference when the distribution of dwelling time is substantially skewed.

John Hopkins University COVID-19 data

This paper uses deaths as its preferred measure of the damages associated with COVID-19, and as a primary driver of mitigation strategies taken by communities. Deaths are a salient measure, widely reported, and were easily accessible during the pandemic. COVID-19 Data Repository by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University is one such source of deaths data (Dong, Du, and Gardner, 2020). The CSSE publishes live and historical COVID-19 cumulative cases and counts on their GitHub

[9] It's important to note, as ultimate interest is the ratio of two estimates, any bias that is constant across the two estimates on deaths and projected deaths will cancel out.

account. Data were tracked beginning in January 2020, and are gathered from a multitude of sources, including news conferences, local government databases, and official government and public health reports.

In their published cumulative deaths, the CSSE includes both confirmed and probable deaths. Local governments will often report a COVID-19 death as probable when COVID-19 deaths are listed on the death certificate as the cause of death or a significant contributing condition, but where there has been no positive confirmatory laboratory test. As this paper is primarily interested in the response to *information*, we do not expect the inclusion of probable deaths to confound the estimates.

Institute for Health Metrics and Evaluation projections

The IHME model (IHME and Murray, 2020) is arguably the most reputable source for projections of COVID-19 deaths. This is the model favored by the White House and most notably cited by the Coronavirus Task Force in multiple national press conferences. For this reason, it is likely the most salient source of estimates for local communities.

The institute makes its projections on a variety of outcomes, including hospitalizations, ventilator utilization, rates of infection, and deaths. The projections are partitioned by state and produced in an effort to inform local policy and health system responses.

The model has changed rapidly as the dynamics of the pandemic evolved. Updates to the model were made on a frequent basis. The average frequency of release between March 25 and June 27, 2020 was once every 3.4 days, with a range of 1 to 11, and an interquartile range of 1 to 5. The distribution of report-to-report percentage changes in *remaining* deaths, 15 days out, is presented in Figure 2. The distribution of inter-report estimates are reported at the state-report level and are based on projections for the next 15 days out. Occasionally, a new model released presents drastically different projections from those of a day prior. For example, the May 4 release produced national estimates 66 percent higher than the previous report, released 6 days earlier. The April 5 model reported estimates 45 percent higher on average than the previous report, 4 days earlier. National averages for the two death metrics—current deaths and projected deaths—are illustrated in Figure 1. Throughout the paper, deaths and projected deaths are reported as the average number of deaths per 10,000.

In the case of the projections, this average is taken over the next T days. Key summary statistics and baseline values of the key variables in this paper are reported in Table 1.

5 Estimating Revealed Utilization of COVID-19 Projections

5.1 Event study specification

The variation leveraged in the projection metric makes it unique from a standard difference-in-differences setting in that there are multiple events, defined by the publication of a new projection. Illustrating the estimates in a standard event study figure is difficult as the dynamic effects from one projection cannot be easily disentangled from the publication of an updated projection. Disentangling these two effects is not of particular interest in this paper, however, illustrating parallel trends is important for the credibility of the paper's empirical design.

An event study specification is estimated using the timing of the first forecast as the main event time. Doing so allows for a clearly defined *pre*-treatment period, when no known forecast was available. The following equation is estimated, where the pre-treatment periods serve as placebos.

$$y_{jt} = \sum_{s=-L}^{L} \omega_{0s} \cdot projection_{jt^0} \times \mathbb{1}(t = t^0 + s) + X_{jt}\beta + \mu_t + \gamma_j + \xi_{jt} \tag{14}$$

where $projection_{jt^0}$ are projected deaths for state j reported at the initial publication date, t^0: March 25, 2020. As all states share a common forecast publication timing, identification in this paper relies on intensity of these projections levels. A death projection level of zero is analogous to a full control group. $\mathbb{1}(t = t^0 + s)$ indicate dates lagging or leading initial publication and ω_{0s} represents the estimated marginal effect at lag or lead s, where $s < 0$ represent placebo effects. Importantly, X_{jt} includes dynamic controls for actual deaths in order to isolate the response to the projections.

Results are presented in Figure 3 as the coefficients on at-home percentage (see Equation 13). Lags and leads are examined on a weekly scale, where the total number of lags/leads is symmetric at $L = 4$ weeks. The first lead is normalized to zero. Note that the dynamic increase of treatment may not be due to the lagged response to the initial forecast, but likely encompasses responses to updated forecasts.

5.2 Standard logit estimates

The main results from Equation 12 are reported in Table 2. The specifications use projected deaths on a 15-day time horizon. As described in Section 4, a 7-day moving average version of reported deaths is used. By doing so, we seek to smooth through irrelevant noise in reporting error (as observed in Figure 1), and better proxy for the measure which is actually internalized on a daily basis by individuals. Rather than aggregating our data to weekly observations, this approach will maintain the variation in the projection updates. Results for alternative moving average lengths are examined in Appendix B.

Columns 1 and 2 present the reduced form effects of projected deaths and current deaths, respectively, in two separate specifications. These two variables are certainly correlated with each other. The IHME projections are not only based on assumptions of future behavior, but also, prior trends in deaths. Thus, to isolate variation in these projections, independent of the current state, we include the two in the same specification. As expected, doing so reduces the magnitude of each coefficient. This estimate is presented in Column 3 and controls for state and day fixed effects. Column 4 introduces state time trends, further isolating the discrete changes in updated projections. Doing so results in a slight increase in the magnitude of the response to both projections and actual deaths, and a moderately larger estimate on α; though, not statistically different from the base specification. For robustness, we allow for time-varying differences across states with different baseline projections. Doing so allows trends to vary based on a state's *initial* projection levels on March 25, 2020, and isolates variation solely in the *updates* to these projections. This specification is presented in Column 5, which results in very similar estimates.

As results in Table 2 are from a logit specification, the direct estimates on deaths and projected deaths can be interpreted as the marginal effect of deaths per 10 thousand on the

log-odds ratio of staying at home. Appendix A reports estimates from a linear probability model. Results yield very similar estimates for α, but with a simpler to interpret result for each death metric. The marginal effect of one additional death per 10 thousand is estimated at a 20 percentage point increase in stay-at-home probability. This is from a baseline stay-at-home average of 46 percent (see Table 1). As for projections, an average increase of one per 10 thousand projected deaths per day, for the next 15 days, is estimated to increase stay-at-home probability by about 10 percentage points.

Estimates from Table 2 use a 15-day time-horizon for the projection length. While a 15-day window is consistent with the time an infected individual is expected to be contagious for, one would expect the magnitude of this estimate to vary based on the time-horizon chosen. As less informative variation is added to the variable by expanding the time window, we should expect the estimate to attenuate. Further, as variation close to the present period is likely more informative, the magnitude of the estimate would be expected to increase. This is illustrated in Figure 4, which estimates the base specification for several different time-horizons in 5-day increments, between 5 and 40 days out. Not surprisingly, the largest estimate is for the narrowest time-window estimated of $T = 5$. This produces an estimate for α of about 0.34. Further, while the estimates seem to decay as larger time-windows are used, it seems to level off at a rapid pace at around 0.2. Thus, while there is variability in estimates for α across the chosen time-horizon, it seems that the range of estimates is relatively low.

An alternative mechanism for these results is that individuals respond directly to contemporary damages to mitigate their own private risk. This is not a dynamic decision and it is reasonable to suspect that current damages are a strong indicator of current risks. Examining lockdown measures would overcome this problem, however, lockdowns are susceptible to reverse causality issues. How contemporary and projected damages affect lockdowns are of interest in this paper, but now how lockdowns might affect damages. Using only individual actions overcomes this problem.

To test whether private risk mitigation is a factor, we incorporate lagged damages into Equation 12. While private risk mitigation should only be a response to contemporaneous changes in damages, lagged responses are likely a function of inertia in policies. An example

would be a lockdown measure implemented in response to a current surge in deaths that has long-term effects on future social distancing, due to the time-frame in which the lockdown applied to.[10] Table 3 reports estimates from a standard logit specification, while including a 1-week lag of deaths and projected deaths. If we expect the lagged responses to be relevant only through a dynamic response, then the ratio of these estimates should provide consistent estimates for α. Using the lagged death measures to calculate α produces similar results. Column 2, which includes state time trends produces the largest upward adjustment of about 0.1, relative to the unlagged version of this specification in Column 4 of Table 2. In either case, we fail to reject that these estimates are statistically different from those produced in the unlagged version.

5.3 Mixed logit estimates

When agents generating the data are each endowed with a given $\alpha_i \in \{0, 1\}$, the observed, state-level outcomes are a mixture of the actions of each type. This is in contrast to the *representative* utility function estimated in Section 5.2, which incorporates the parameter $\alpha \in [0, 1]$ into its form. To account for the mixing of different strategies performed by each type, this section of the paper presents estimates from a mixed logit specification.

The first specification estimates α and β directly, allowing α to operate as a mixing parameter. An analytical solution to the inverse probability function is no longer feasible, thus, we implement the fixed point iteration procedure proposed by Berry, Levinsohn, and Pakes (1995). Doing so allows us to easily control for a full set of state and day fixed effects. Estimates are reported in the first column of Table 4. Standard errors are bootstrapped on state-level clusters. Of main interest is the estimate for α, which is nearly identical to those of Column 3 in Table 2. This suggests that modeling the mixing of strategies directly does not contribute significantly to the identification of α. An estimate of β can be derived from the standard logit estimates in Table 2 by dividing the estimate on projected deaths by the estimate for α. In Column 3 this produces an estimate of about 1.2, similar to that estimated in Table 4.

[10] This is just one of many examples. Other examples may include new bar and restaraunt restrictions, or even private businesses going remote.

5.4 Cognitive dissonance estimates

Cognitive dissonance theory (Festinger, 1957) argues that individuals seek psychological consistencey, for example, through their actions and their beliefs. Internal inconsistency leads to stressful behavior for agents, often causing them to generate consistency by altering beliefs in a manner that rationalizes their behavior, or by avoiding particular information that leads to this dissonance. The theory suggests that by manipulating one's own beliefs in such a manner, individuals reduce mental stress. Such behavior can ultimately lead people to believe whatever is most compatible with their preferred behavior.

When mitigating is costly, individuals may choose to alter their beliefs in order to rationalize less costly behaviors. Intuitively, this is more likely to occur when making the decision to mitigate external costs as opposed to private costs; as in the latter case the individual may have an incentive to rationally incorporate the most accurate information into their decisions. In the context of this paper, someone might selectively choose to form their beliefs based on current deaths over scientific projections ($\alpha_i = 0$) if the outlook seems more optimistic. By leaning in favor of the more optimistic forecast, an agent can maintain consistency with their preferred, low-mitigating behavior.

As opposed to the baseline estimates for α, which reveal the average utilization of each metric, the specification illustrated in Equation 10 incorporates the relative difference in disutility from each damage estimate into the parameter. Parameter estimates from Equation 11 are presented in Table 4. The coefficient estimate for $\tilde{\alpha}_1$ provide evidence of cognitive dissonance. The negative sign suggests that utilization of the scientific projections of damages decreases as their magnitude increases, relative to contemporary damages.

Figure 5 illustrates these estimates graphically. The estimated function for α is plotted over its primary argument—the relative difference in inclusive values, scaled to the mean *contemporary* damage levels. This provides an interpretation of the relative value of the projection, in percentage terms. The function is evaluated for values outside the range of the vast majority of our data. This is indicated by the histogram, which plots the frequency of the relative inclusive values in our data. Plotting it outside of this range illustrates the large differences in internalization of projections that can occur when distant damage estimates

significantly differ from contemporary states. For easier interpretation, we calculate the slope of the α function with respect to projected damages. The marginal effect is evaluated across all observations in our data. Estimates suggest that a doubling in the value of the projected damages, all else equal, produces a 4 percent reduction in the value of α.

6 Application to Climate Policy

6.1 A framework for calculating a feasible Pigouvian tax on CO2

In this section, the model that has been discussed thus far in the context of the COVID-19 pandemic is applied to the context of climate change. We begin by generalizing the framework in Section 2 to allow damages to occur over an infinite time-horizon. Let the marginal value of a ton of carbon emitted at time t on damages in period $t+s$, $s \geq 0$, be described by the function $d_t(s)$, where the i subscript is omitted to indicate a global damage. The present discounted value of a ton of carbon emitted today is as follows.

$$D_t = \sum_{s=0}^{\infty} \left(\frac{1}{1+r}\right)^s \cdot d_t(s) \tag{15}$$

where r is the social discount rate. Social leaders who take science as the truth condition their mitigation decision on the scientific forecast. These forecasts establish the best available estimate for the future cost of carbon. Thus, we define the social cost of carbon in this paper as the following.

$$SCC_t = \sum_{s=0}^{\infty} \left(\frac{1}{1+r}\right)^s \cdot \mathbb{E}_t[d_t(s)] \tag{16}$$

where $\mathbb{E}_t$ denotes the expectation, conditional on information available at time t.[11] In contrast, the climate skeptic does not form their expectations according to scientific forecasts,

[11] This notation omits a third time dimension for simplicity. Of course, the social cost of carbon can be evaluated for future periods, conditional on expectations today.

but conditions on contemporaneous damages. We assume those with myopic beliefs take current damages and projects the same state into the future, indefinitely. Further, we assume that they discount these expected damages according to the same social discount rate, r. This creates a two-type analogue to Equation 3. Define expectations over the cost of carbon as the following.

$$\mathbb{E}_t[D_i|\theta] = \begin{cases} SCC_t & \text{if } \theta = science \\ \frac{1+r}{r} \cdot d_t(0) & \text{if } \theta = myopic \end{cases} \tag{17}$$

where the myopic beliefs represent the discounted sum of constant damages at time t, $d_t(0)$. Similar to the framework in Section 2, the policy-maker forms their mitigation strategy conditional on its expectations of future damages. Here, we discuss mitigation strategies in the context of a Pigouvian tax on carbon emissions. When marginal damages are constant over consumption, the efficient tax is equal to the cost of damages. However, given uncertainty over future damages, the ultimate mitigation strategy is the tax which conforms with the policy maker's belief type.

Given a distribution of belief types, we can calculate a *perceived* social cost of carbon, which places relevant weights on each type's expectations. The *perceived* social cost of carbon is a *feasible* tax, in that it takes into account the mitigation strategies each type finds optimal. Thus, of ultimate interest is the following expression.

$$Perceived\ SCC_t = \alpha \cdot SCC_t + (1-\alpha) \cdot \frac{1+r}{r} \cdot d_t(0) \tag{18}$$

6.2 The (perceived) social cost of carbon

The ultimate interest in this section of the paper is the value of Equation 18, which may be interpreted as a feasible tax on CO2 emissions. An obvious caveat of this portion of the paper is the contrast between the model estimated in this paper in the context of COVID and the model applied to the context of climate change. While damages associated with the virus

may be converted to similar scales as those from CO2 emissions (e.g., economic welfare), the manner in which policy makers internalize the damages from each source may differ. Further, it is possible that the distribution of belief types for COVID are quite different from belief types for climate change estimates. These are starkly different mechanisms which may influence information attainment in contrasting ways. Therefore, caution should be made in directly interpreting the estimated α in Section 5 as exactly equal to α in the climate setting. Rather, one should interpret the results that follow as expressing an analogue to the policy response to COVID. That is, as the estimated α expresses implied myopia in the COVID response, this section asks what this would mean for climate policy, if the response to climate change were similar.

To calculate Equation 18, we make use of the Dynamic Integrated Climate-Economy (DICE) model (Nordhaus and Boyer, 2000; Nordhaus, 2008). Specifically, we use the DICE-2016R model (Nordhaus, 2017, 2019) to gather information about the estimated social cost of carbon, in addition to estimated contemporaneous damages—both terms in Equation 18. This integrated assessment model simulates the impact of an incremental increase in CO2 emissions at a present period on future climate states. The model then applies a damage function to estimate its impact on social welfare. The ultimate result is a counterfactual path of economic welfare, for which inference is made against a baseline.

Marginal damages change nonlinearly as the predicted stock of carbon in the atmosphere changes over time. Thus, contemporaneous damages, $d_t(0)$, change over t. The DICE model is used to calculate these damages for each t between 2020 and 2100. We use the model's specification which incorporates optimal carbon taxes into its projections and uses a 5 percent social discount rate. Ultimately, the social cost of carbon estimates vary to a degree from those used by the EPA, as DICE is only one of several that the EPA incorporates into their estimates.[12]

Figure 6 plots the path of SCC, as predicted by DICE. The figure also plots two different measures of the perceived SCC in Equation 18. The first is derived using the base estimate of $\alpha = 0.26$, presented in Section 5. This is the constant value of α, which this paper argues, illustrates the implied weighting of projections versus contemporary damages in response to

[12]See Interagency Working Group on Social Cost of Greenhouse Gases (2016).

the COVID-19 pandemic. When applied to the climate context, this level of attention to scientific estimates would suggest a significant undervaluing of the cost of carbon. While DICE estimates a SCC equal to $37 per ton in 2020, the implied valuation is half that, at about $19. The implied valuation approaches 70 percent of full valuation in 2100.

The second valuation measure applies the non-constant version of α, estimated in Section 5.4. A major caveat here is the substantial difference between contemporary carbon damages and the present, discounted values in the setting of predicted climate damages. While the range illustrated in Figure 5 approaches 50 percent, in the climate setting this difference begins at 105 percent in 2020. This generates a predicted value of α near zero. Using such a weight implies that agents place full emphasis on current damages in 2020—which is one-third of the SCC. This measure of α does not approach 5 percent until 2285—when the difference between contemporary and expected future damages is 16 percent. By 2100 this measure generates an implied valuation of only 60 percent of the full SCC.

7 Conclusion

Beliefs about the inevitable risks of climate change are a determining factor for whether efficient policy is feasible to implement. In order to internalize the costs of damaging actions today, we must first understand the true value of these costs. While activities that impose contemporaneous damages are salient, actions that produce significantly lagged responses make it difficult to properly infer cause and effect. This is most prominant in the domain of climate change, where carbon emitted today can stay in the atmosphere for hundreds of years into the future. Suffice to say that the realized damages at the time of emission are not representative of its long-run impact on the environment. Understanding these costs often means relying on the experts, which some may not be willing to do.

The main result of this paper suggests that scientific projections may not be internalized in the decision making process to the degree to which they should be when they represent the best available information. The mechanism is likely some mixture of skepticism, inattention, and cognitive dissonance. The paper demonstrates a general underweighting of scientific estimates in decision-making, opting instead for a heuristic projection based on the current

state. In the context of the COVID-19 pandemic, where projections were both salient—through means such as daily White House press conferences—and readily available through the IHME, this paper show that mobility outcomes are instead driven by current measures of deaths, on a scale of three to one. Evidence is provided to suggest that these effects are largely driven by the dynamic decision to mitigate societal damages, rather than an individual's decision to mitigate personal risk.

Finally, this paper presents evidence that agents strategically choose more favorable forecasts when making mitigation decisions. In the trade-off between a naïve expectation of future damages that is strictly based on the current state, versus a more objective outlook based on scientific analyses, we show that, while agents on average prefer the myopic expectation, they will shift in favor of a scientific perspective when the projection is relatively more optimistic. This evidence is suggestive of an inherent cognitive dissonance among individuals, where beliefs adapt to be most consistent with the mitigation strategy that an agent ultimately prefers. Unfortunately, this behavior will not produce a socially efficient outcome in policy design.

In the context of climate change, when emissions today have costly outcomes for future generations, incorporating these damages into carbon intensive activities requires a good understanding of the social cost of carbon. When environmental impacts have uncertain outcomes, the efficient strategy is that which incorporates all available information—the crux of rational expectations theory. In practice, the manner in which agents derive their expectations is not as simple, and will often come about through ideological, rather than objective means.

While this paper looks to apply its findings to the contexts of climate change, this exercise is merely an extrapolation for illustrative purposes. The idea is to better understand the costs of this revealed subjective belief formation on a broader scale. In the context of climate change, the relevance of cognitive dissonance may be exacerbated due to the disconnect people have from these future events. However, better understanding the role this behavior plays in the realm of climate policy requires further research. A better understanding of how climate beliefs are formed is essential in properly guiding efficient policy design in the future.

References

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4):841–890.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2016. "Competition for Attention." *Review of Economic Studies* 83:481–513.

Busse, Meghan R., Devin G. Pope, Jaren C. Pope, and Jorge Silva-Risso. 2014. "The Psychological Effect of Weather on Car Purchases." *Quarterly Journal of Economics* 131 (1):371–414.

Carson, R. T., N. E. Flores, K. M. Martin, and J. L. Wright. 1996. "Contingent valuation and revealed preference methodologies: comparing the estimates for quasi-public goods." *Land Economics* 72:80–99.

Comin, Diego and Johannes Rode. 2015. "From Green Users to Green Voters." *NBER Working Paper 19219* .

Conlin, Michael, Ted O'Donoghue, and Timothy J. Vogelsang. 2007. "Projection Bias in Catalog Orders." *American Economic Review* 97 (4):1217–1249.

Crandall, Robert W. and John D. Graham. 1989. "The Effects of Fuel Economy Standards on Automobile Safety." *Journal of Law and Economics* 32:97–118.

Dong, Ensheng, Hongru Du, and Lauren Gardner. 2020. "An interactive web-based dashboard to track COVID-19 in real time." *The Lancet Infectious Diseases* 20 (5):533–534.

Durante, Ruben, Luigi Guiso, and Giorgio Gulino. 2021. "Asocial capital: Civic culture and social distancing during COVID-19." *Journal of Public Economics* 194:104342. URL https://www.sciencedirect.com/science/article/pii/S0047272720302061.

Festinger, Leon. 1957. "A Theory of Cognitive Dissonance." *California: Stanford University Press* .

Fisman, Raymond, Hui Lin, Cong Sun, Yongxiang Wang, and Daxuan Zhao. 2021. "What motivates non-democratic leadership: Evidence from COVID-19 reopenings in China." *Journal of Public Economics* 196:104389. URL https://www.sciencedirect.com/science/article/pii/S0047272721000256.

Furnham, Adrian. 2011. "A Literature Review of the Anchoring Effect." *The Journal of Socio-Economics* 40 (1):35–42.

Gabaix, Xavier. 2014. "A Sparsity-Based Model of Bounded Rationality, Applied to Basic Consumer and Equilibrium Theory." *The Quarterly Journal of Economics* 129 (4):1661–1710.

Gabaix, Xavier and David Laibson. 2005. "Bounded Rationality and Directed Cognition." *Unpublished manuscript. New York University, Stern School of Business, New York.* .

Hanley, N. and M. Czajkowski. 2019. "The role of stated preference valuation methods in understanding choices and informing policy." *Review of Environmental Economics and Policy* 13:248–66.

Heal, Geoffrey and Antony Millner. 2014. "Uncertainty and decision making in climate change economics." *Review of Environmental Economics and Policy* 8 (1):120–137.

Howe, Peter D., Matto Mildenberger, Jennifer R. Marlon, and Anthony Leiserowitz. 2015. "Geographic variation in opinions on climate change at state and local scales in the USA." *Nature Climate Change* 5:596–603.

IHME, COVID-19 Health Service Utilization Forecasting Team and Christopher JL Murray. 2020. "Forecasting COVID-19 impact on hospital bed-days, ICU-days, ventilator-days and deaths by US state in the next 4 months." *medRxiv* URL https://www.medrxiv.org/content/early/2020/03/30/2020.03.27.20043752.

Interagency Working Group on Social Cost of Greenhouse Gases, United States Government. 2016. "Addendum to Technical Support Document on Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866: Application of the Methodology to Estimate the Social Cost of Methane and the Social Cost of Nitrous Oxide." URL https://www.epa.gov/sites/production/files/2016-12/documents/addendum_to_sc-ghg_tsd_august_2016.pdf.

Kahan, Dan M., Hank Jenkins-Smith, and Donald Braman. 2010. "Cultural cognition of scientific consensus." *Journal of Risk Research* 14 (2):147–174.

Leiserowitz, Anthony, Edward Maibach, Connie Roser-Renouf, Seth Rosenthal, and Matthew Cutler. 2016. "Climate Change in the American Mind: November 2016." *Yale University and George Mason University* .

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility." *The Quarterly Journal of Economics* 118 (4):1209–1248.

Marlon, Jennifer, Peter Howe, Matto Mildenberger, Anthony Leiserowitz, and Xinran Wang. 2020. "Yale Climate Opinion Maps 2020." URL https://climatecommunication.yale.edu/visualizations-data/ycom-us/.

Mendelsohn, R. 2019. "An examination of recent revealed preference valuation methods and results." *Review of Environmental Economics and Policy* 13:267–82.

Millner, Antony, Simon Dietz, and Geoffrey Heal. 2013. "Scientific Ambiguity and Climate Policy." *Environmental and Resource Economics* 55:21–46.

Millner, Antony and Hélène Ollivier. 2016. "Beliefs, Politics, and Environmental Policy." *Review of Environmental Economics and Policy* 10 (2):226–244.

Mullainathan, Sendhil and Ebonya Washington. 2009. "Sticking with Your Vote: Cognitive Dissonance and Political Attitudes." *American Economic Journal: Applied Economics* 1 (1):86–111.

Nickerson, R. S. 1998. "Confirmation bias: A ubiquitous phenomenon in many guises." *Review of General Psychology* 2:175–220.

Nordhaus, William D. 2008. "A Question of Balance: Weighing the Options on Global Warming Policies." *New Haven, CT: Yale University Press* .

———. 2017. "Revisiting the social cost of carbon." *Proceedings of the National Academy of Sciences* 114 (7):1518–1523.

———. 2019. "Climate Change: The Ultimate Challenge for Economics." *American Economic Review* 109 (6):1991–2014.

Nordhaus, William D. and J. Boyer. 2000. "Warming the World: Economic Models of Global Warming." *Cambridge, MA: MIT Press* .

Oates, Wallace and Paul R. Portney. 2003. "The political economy of environmental policy." *Chapter 08 in Handbook of Environmental Economics* 1:325–354.

Pulejo, Massimo and Pablo Querubín. 2021. "Electoral concerns reduce restrictive measures during the COVID-19 pandemic." *Journal of Public Economics* 198:104387. URL https://www.sciencedirect.com/science/article/pii/S0047272721000232.

Sallee, James M. 2014. "Rational Inattention and Energy Efficiency." *The Journal of Law and Economics* 57 (3):781–820.

Schlenker, Wolfram and Charles A. Taylor. 2021. "Market expectations of a warming climate." *Journal of Financial Economics* .

Slovic, Paul. 1967. "The Relative Influence of Probabilities and Payoffs upon Perceived Risk of a Gamble." *Psychonomic Science* 9:223–224.

Stigler, George J. 1961. "The Economics of Information." *Journal of Political Economy* 69 (3):213–225.

Tversky, Amos and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* :1124–1131.

van den Bremer, Ton S. and Frederick van der Ploeg. 2021. "The Risk-Adjusted Carbon Price." *American Economic Review (Forthcoming)* .

8 Figures

Figure 1: Daily Deaths and 15-Day Out Projected Average Daily Deaths (per 10,000)

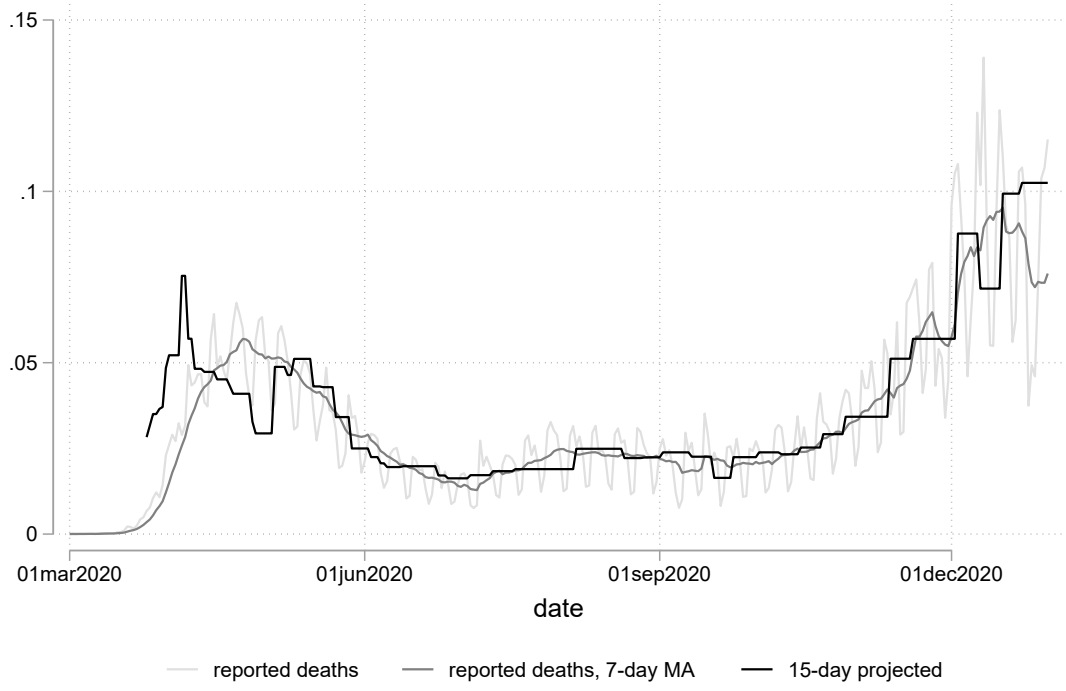

Notes: The 15-day projected risk, in black, takes the daily 15-day out projection from the given date, and computes the average daily number deaths per 10,000. In light gray represents the current level of daily reported deaths and dark gray represents a 7-day moving average of reported daily deaths, each per 10,000.

Figure 2: Distribution of Report Update Percent Changes in Projected Deaths

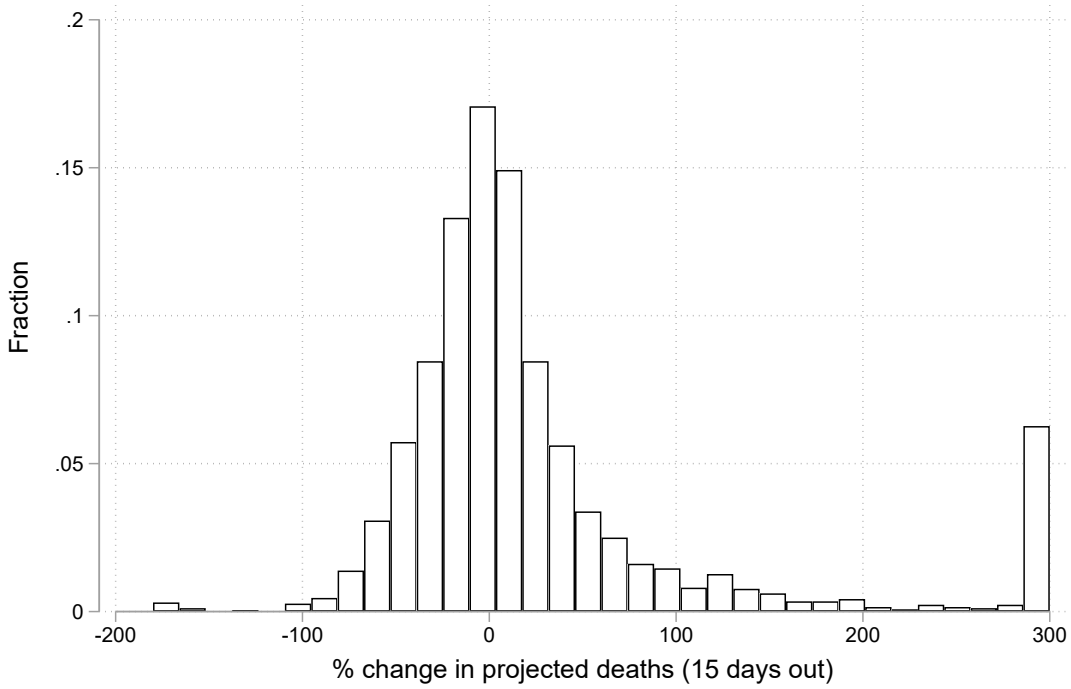

Notes: The above reports the distribution of report-to-report percentage change in the 15-day projected outlook. All changes are mainly a function of adjustments to the IHME model. All observations above 300% are included into the 300% bin.

Figure 3: Event Study of First Projection Release

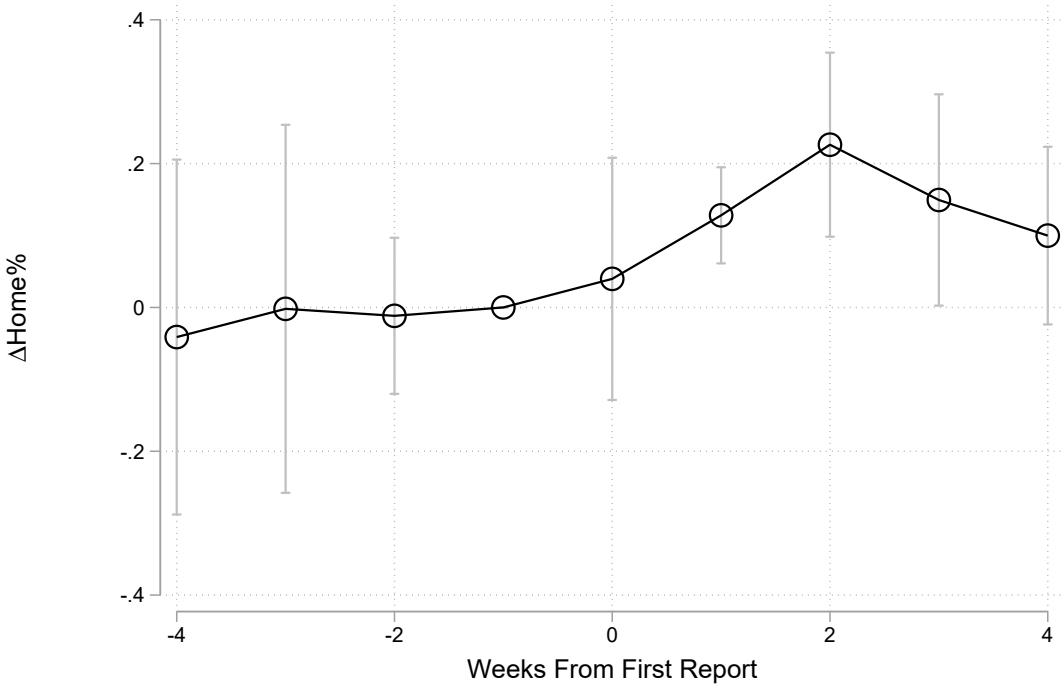

Notes: The above reports the coefficient estimates from the event study specification in Equation 14. All periods along the horizontal axis are weeks, relative to the publication of the intial forecast: March 25, 2021. The marginal effects are for a 1 per 10,000 increase in average daily projections for the 15-day forecast. The coefficient for the one-period lead is normalized to zero.

Figure 4: Estimate of α Across Different Time Horizons

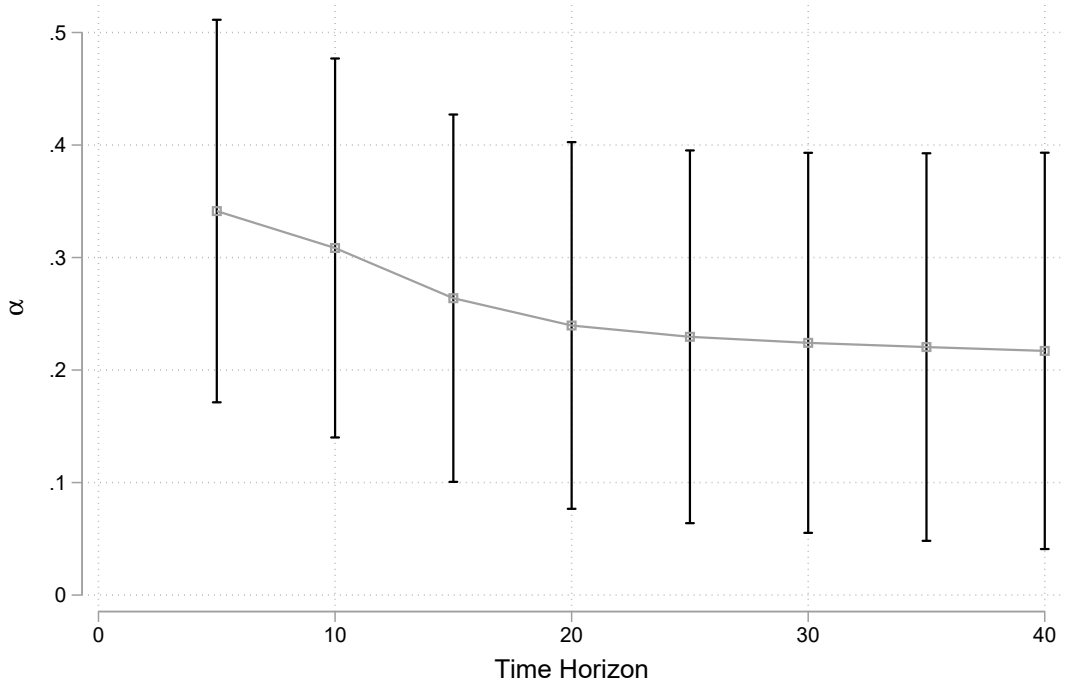

Notes: The above reports the coefficient estimates for projected deaths from the base specification in Equation 12 across different projection time-horizon. Estimates are reported for $T = 5$ to $T = 40$, in increments of 5.

Figure 5: Non-Constant α by Relative Projection Magnitude

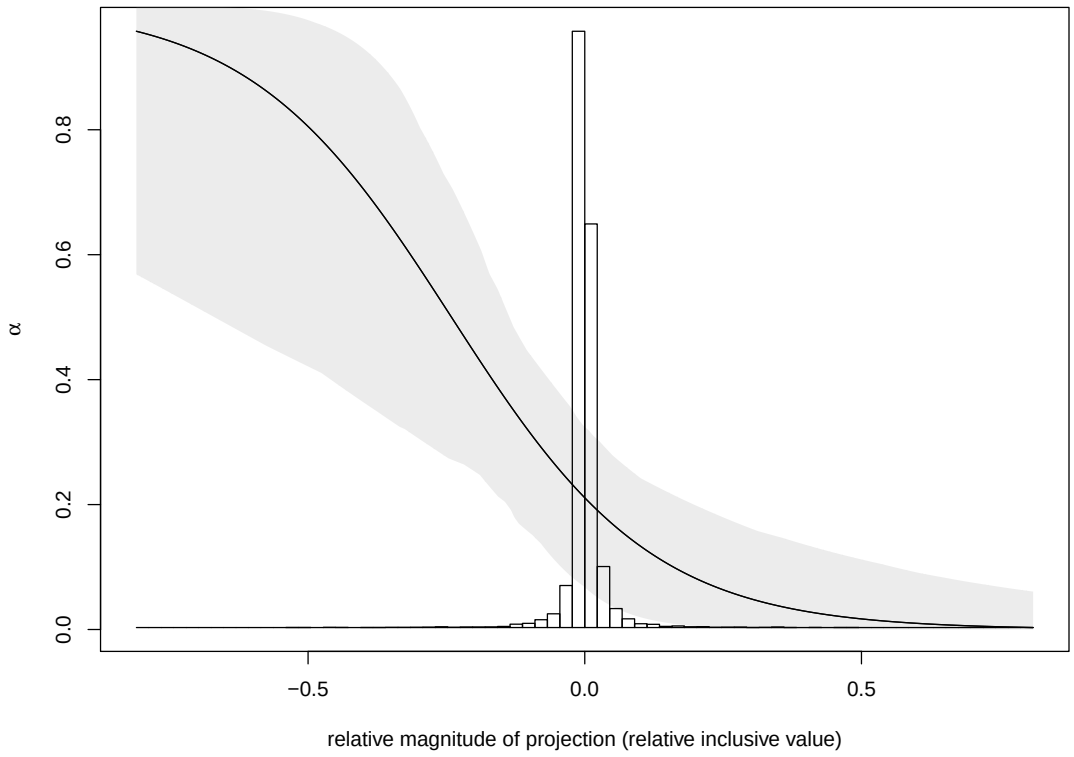

Notes: In solid black are the estimates from Equation 10 with the corresponding bootstrapped 95% confidence intervals. This is a smooth function of the relative inclusive values for each damage metric. Scatter points are the model-predicted estimates plotted against the log differences between the two damage metrics. Estimates suggest an average 10.7 percentage point reduction in α for a 10 percent relative increase in the scientific projection.

Figure 6: Implied Valuations of Carbon using α Estimates

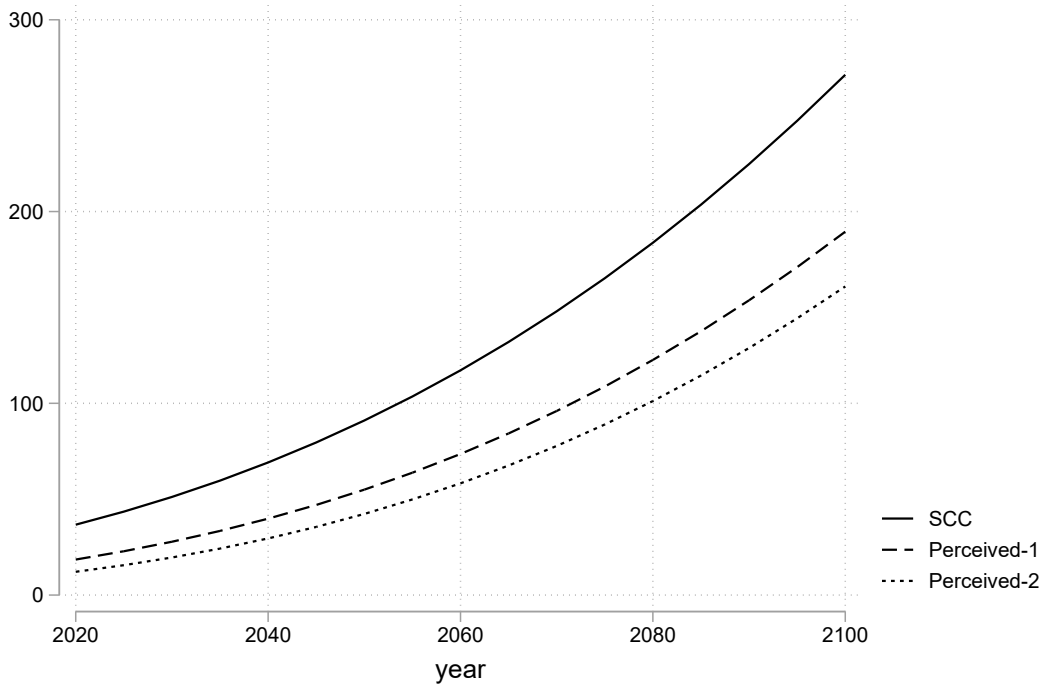

Notes: In solid black represents the social cost of carbon estimates generated from the DICE-2016R model (Nordhaus, 2017, 2019). Perceived-1 represents the analogous weighted average of contemporary damages from the DICE model and social cost of carbon (Equation 18) using an estimate of α of 0.26. Perceived-2 uses the non-constant α estimates (Equation **??**).

9 Tables

Table 1: Summary Statistics

	Mean	St. Dev.	Interquartile Range
At Home Percent	0.461	0.072	[0.421, 0.501]
Daily Deaths, per 10,000	0.028	0.052	[0.000, 0.033]
Projected 5-Day Avg, per 10,000	0.027	0.040	[0.000, 0.034]
Projected 15-Day Avg, per 10,000	0.028	0.041	[0.000, 0.036]
Projected 30-Day Avg, per 10,000	0.028	0.039	[0.000, 0.037]
IHME Update Frequency (days)	5.5	3.9	[2.0, 7.0]

Notes: At Home Percent is the fraction of time spent at home during the day (Equation 13). Daily deaths and projected deaths are reported in deaths per 10,000, where projected deaths are average daily deaths over the T-day forecast. IHME Update Frequency is the number of days between forecast updates.

Table 2: Relative Mitigation Responses to Deaths and Projected Deaths

	(1)	(2)	(3)	(4)	(5)
Projected Deaths	0.98***		0.32***	0.40***	0.36***
	(0.11)		(0.11)	(0.10)	(0.11)
Current Deaths		1.10***	0.88***	0.95***	0.99***
		(0.12)	(0.13)	(0.16)	(0.14)
α			0.264***	0.297***	0.266***
			(0.084)	(0.076)	(0.076)
N	18253	18253	18253	18253	18253
State FEs	Yes	Yes	Yes	Yes	Yes
Day FEs	Yes	Yes	Yes	Yes	Yes
State Trends				Yes	Yes
Day × First Projections					Yes

Notes: * $p < 0.1$, ** $p < .05$, *** $p < .01$. Standard errors, clustered by state, in parentheses. Standard errors for α are approximated using the delta method. Column 3 is the base specification from Equation 12. Projected Deaths include the 15-day out forecasts for average daily deaths per 10,000. Current deaths is a 7-day rolling average of deaths per 10,000.

Table 3: Relative Mitigation Responses to Deaths and Projected Deaths

	(1)	(2)	(3)
Projected Deaths	0.26***	0.32***	0.24***
	(0.10)	(0.09)	(0.08)
Current Deaths	0.48***	0.49***	0.55***
	(0.16)	(0.17)	(0.17)
Lagged Projected Deaths	0.32**	0.41***	0.33**
	(0.16)	(0.15)	(0.15)
Lagged Current Deaths	0.93***	0.96***	1.05***
	(0.17)	(0.19)	(0.18)
α	0.254**	0.390***	0.309***
	(0.120)	(0.111)	(0.109)
N	17840	17840	17840
State FEs	Yes	Yes	Yes
Day FEs	Yes	Yes	Yes
State Trends		Yes	Yes
Day × First Projections			Yes

Notes: * $p < 0.1$, ** $p < .05$, *** $p < .01$. Standard errors, clustered by state, in parentheses. Standard errors for α are approximated using the delta method. Projected Deaths include the 15-day out forecasts for average daily deaths per 10,000. Current deaths is a 7-day rolling average of deaths per 10,000. Specifications are analogous to Columns 3-5 in Table 2, but including a one week lag in deaths/projected deaths. Estimates of α are derived from the coefficients on the lagged terms.

Table 4: Mixed Logit Results

	Base		Nested Logit
α	0.263	$\tilde{\alpha}_0$	-1.320
	(0.076)		(0.510)
β	1.196	$\tilde{\alpha}_1$	-8.877
	(0.111)		(4.200)
		β	1.300
			(0.104)
N	18253		18253
State FEs	Yes		Yes
Day FEs	Yes		Yes

Notes: Wild bootstrapped standard errors, clustered by state, in parentheses. The column labeled "Base" is from a mixed logit specification, where α is a (constant) mixing parameter. The nested logit specification models α as a function of the mitigation inclusive values (Equation 11). $\tilde{\alpha_0}$ and $\tilde{\alpha_1}$ are logit parameters. Each specification includes state and day fixed effects.

Appendices

A Linear Probability Model

Equation 12 is estimated as a linear probability model, using fraction of time at home as the outcome variable. The results are presented in Table A.1. Estimates for α are both quantitatively and qualitatively similar to the results from the logit specification in Table 2. The results suggest that the marginal effect of a one additional death per 10 thousand is a 20 percentage point increase in stay-at-home probability. This is from a baseline stay-at-home average of 46 percent. As for projections, an average increase of one per 10 thousand projected deaths per day, for the next 15 days, is estimated to increase stay-at-home probability by about 10 percentage points.

Table A.1: Relative Mitigation Responses to Deaths and Projected Deaths (Linear Probability Model)

	(1)	(2)	(3)	(4)	(5)
Projected Deaths	0.24***		0.08***	0.10***	0.09***
	(0.02)		(0.03)	(0.02)	(0.02)
Current Deaths		0.26***	0.21***	0.23***	0.24***
		(0.03)	(0.03)	(0.04)	(0.03)
α			0.267***	0.299***	0.266***
			(0.083)	(0.076)	(0.073)
N	18253	18253	18253	18253	18253
State FEs	Yes	Yes	Yes	Yes	Yes
Day FEs	Yes	Yes	Yes	Yes	Yes
State Trends				Yes	Yes
Day × First Projections					Yes

Note: * $p < 0.1$, ** $p < .05$, *** $p < .01$. Standard errors, clustered by state, in parentheses.

B The Use of Alternative Moving Average Lengths for Deaths

Tables A.2, A.3, and A.4 present the standard logit results for Equation 12 using alternative moving average lengths of 3, 5, and 10 days, respectively. The estimates in the main text use a moving average of 7 days. By doing so, uninformative day-to-day variation based on reporting error should be smoothed through. This approach should properly discard variation that may otherwise attenuate the estimates on current deaths, thereby artificially

increasing the magnitude of α.

Table A.2 demonstrates that when using a smaller time-window, estimates for α increase due to the smaller estimated coefficient on current deaths. Estimates from the 5-day and 7-day moving averages, from Tables A.3 and A.4, respectively, produce estimates similar to those reported in the paper.

Table A.2: Relative Mitigation Responses to Deaths and Projected Deaths (3-day MA)

	(1)	(2)	(3)	(4)	(5)
Projected Deaths	0.98***		0.52***	0.62***	0.62***
	(0.11)		(0.10)	(0.10)	(0.11)
Current Deaths		0.92***	0.60***	0.65***	0.63***
		(0.09)	(0.08)	(0.09)	(0.10)
α			0.461***	0.490***	0.497***
			(0.063)	(0.058)	(0.066)
N	18457	18457	18457	18457	18457
State FEs	Yes	Yes	Yes	Yes	Yes
Day FEs	Yes	Yes	Yes	Yes	Yes
State Trends				Yes	Yes
Day × First Projections					Yes

Note: * $p < 0.1$, ** $p < .05$, *** $p < .01$. Standard errors, clustered by state, in parentheses.

Table A.3: Relative Mitigation Responses to Deaths and Projected Deaths (5-day MA)

	(1)	(2)	(3)	(4)	(5)
Projected Deaths	0.98***		0.38***	0.48***	0.45***
	(0.11)		(0.10)	(0.10)	(0.11)
Current Deaths		1.04***	0.79***	0.85***	0.86***
		(0.11)	(0.11)	(0.13)	(0.12)
α			0.327***	0.360***	0.346***
			(0.076)	(0.069)	(0.074)
N	18355	18355	18355	18355	18355
State FEs	Yes	Yes	Yes	Yes	Yes
Day FEs	Yes	Yes	Yes	Yes	Yes
State Trends				Yes	Yes
Day × First Projections					Yes

Note: * $p < 0.1$, ** $p < .05$, *** $p < .01$. Standard errors, clustered by state, in parentheses.

Table A.4: Relative Mitigation Responses to Deaths and Projected Deaths (10-day MA)

	(1)	(2)	(3)	(4)	(5)
Projected Deaths	0.98***		0.31***	0.38***	0.33***
	(0.11)		(0.11)	(0.10)	(0.10)
Current Deaths		1.12***	0.90***	0.97***	1.02***
		(0.13)	(0.14)	(0.16)	(0.15)
α			0.254***	0.284***	0.243***
			(0.084)	(0.077)	(0.074)
N	18100	18100	18100	18100	18100
State FEs	Yes	Yes	Yes	Yes	Yes
Day FEs	Yes	Yes	Yes	Yes	Yes
State Trends				Yes	Yes
Day × First Projections					Yes

Note: * $p < 0.1$, ** $p < .05$, *** $p < .01$. Standard errors, clustered by state, in parentheses.

C Non-Constant α from Standard Logit

The following modified version of Equation 12 is estimated to allow for non-constant marginal effects on each term.

$$\begin{aligned} y_{jt} =& \omega_{01} \cdot projection_{jt} \times post\ report_t + \omega_{02} \cdot projection_{jt}^2 \times post\ report_t \\ &+ \omega_{11} \cdot current\ deaths_{jt} + \omega_{12} \cdot current\ deaths_{jt}^2 + \lambda_t + \gamma_j + \xi_{jt} \end{aligned} \tag{19}$$

From this quadratic specification, a non-constant value for α is derived as the implied weight placed on the partial effects of each term. Specifically, an alternative form for α_{jt} is specified as follows.

$$\alpha_{jt} = \frac{\omega_{01} + 2\omega_{02} \cdot projection_{jt}}{\omega_{01} + 2\omega_{02} \cdot projection_{jt} + \omega_{11} + 2\omega_{12} \cdot current\ deaths_{jt}} \tag{20}$$

The most significant distinction of this form for α from that in Equation **??** is that it is simply a function of the levels of each variable, rather than their log difference. Estimates for each term in this framework are presented in Table A.5, under the base specification that includes state and day fixed effects.

Table A.5: Quadratic Specification of Standard Logit

	(1)
Projected Deaths - Linear	0.58***
	(0.18)
Projected Deaths - Quadratic	-1.57***
	(0.48)
Current Deaths - Linear	1.51***
	(0.16)
Current Deaths - Quadratic	-2.33***
	(0.36)
State FEs	Yes
Day FEs	Yes
Observations	18253

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

To plot this modified α in one dimension, I evaluate it at the mean value of current deaths. The result is now strictly a function of projected deaths. This is illustrated in Figure A.1. While the interpretation is different from that in the main text, since the domain differs, the graph illustrates a similar result. As projections near zero, the weight placed on them increases. This suggests that, all else equal, agents favor the metric which produces a lower level of damages.

Figure A.1: Non-Constant α from Standard Logit

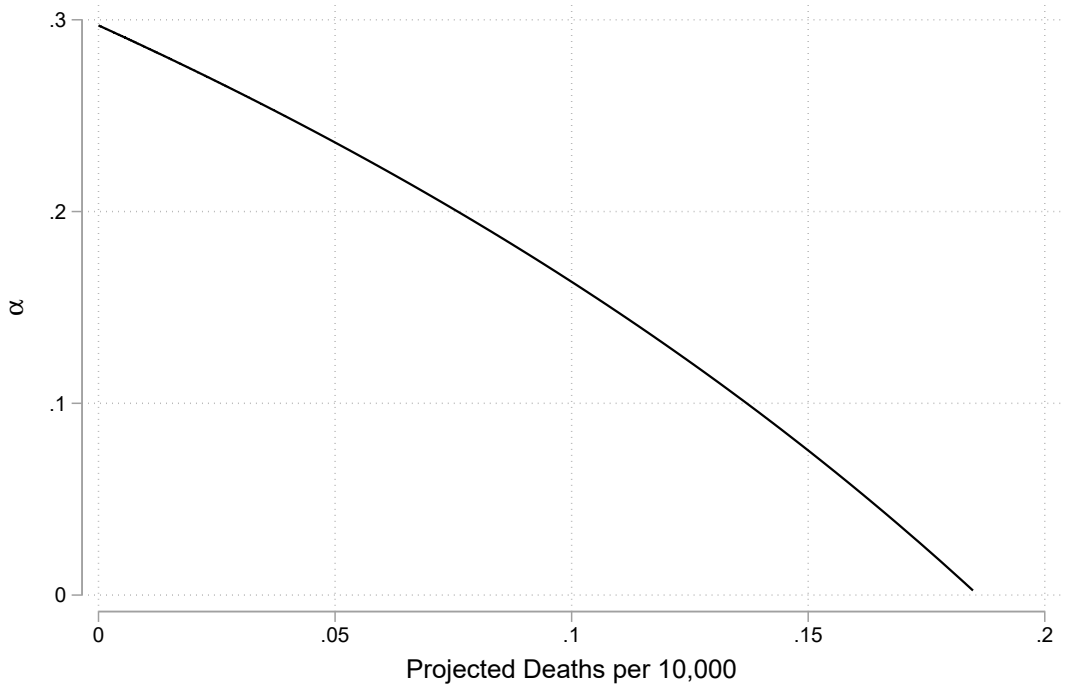

Notes: This function is graphed outside of the range of *most* of the data. See Table 1 for distribution statistics for projected deaths.