Disasters with Unobservable Duration and Frequency: Intensified Responses and Diminished Preparedness*

Viral V. Acharya† Timothy Johnson‡ Suresh Sundaresan§ Steven Zheng¶

December 2024

Abstract

We study an economy subject to recurrent disasters when the frequency and duration of the disasters are unobservable parameters. Imprecise information about transition intensities increases the probability of the current state effectively lasting forever, with important consequences for agents' welfare and optimal behavior. In a disaster, uncertainty about duration makes disasters subjectively much worse and can make the welfare value of information extremely high. However, in advance of a disaster, uncertainty about the arrival rate can be welfare-increasing. Agents optimally invest less in mitigation than under full information and pay less for insurance against the next disaster.

JEL Codes: D6, D8, E21, E32, G10
Keywords : parameter uncertainty, welfare costs, disasters, mitigation

*We are grateful to Xavier Gabaix, Michael Gallmeyer, Laura Veldkamp and Stijn van Nieuwerburgh, two anonymous referees, anonymous Associate Editor, and Lukas Schmid (Editor), for thoughtful comments, and to Stefano Pastore for excellent research assistance. Viral Acharya, Suresh Sundaresan and Steven Zheng dedicate this article to their late co-author, Tim Johnson.

†Corresponding author. New York University, Stern School of Business, NBER, and CEPR: vacharya@stern.nyu.edu.

‡Formerly at University of Illinois at Urbana-Champaign: tcj@illinois.edu.

§Columbia University, Graduate School of Business: ms122@columbia.edu.

¶Columbia University, Graduate School of Business: steven.zheng@columbia.edu.

1 Introduction

This paper studies the real effects of parameter uncertainty in a model of repeated disasters. Among the many deeply alarming aspects of the COVID-19 pandemic was the realization of *how little anybody knew* about the right model for what would happen. Two dimensions of structural uncertainty seemed especially salient: uncertainty about the persistence (or duration) of the crisis, and uncertainty about the recurrence (or frequency) of future disasters. Such uncertainty pertains to disasters beyond pandemics. Uncertainty about the duration and frequency of recessions and financial crises is also realistic. Similar considerations apply in the context of climate-related disasters.

Our model depicts disasters as regimes in which the stock of wealth (potentially including human wealth) is subject to exogenous destruction. The economy transitions stochastically between these episodes and "normal times." Agents update their beliefs by observing the frequencies of transitions, and optimally solve their investment/consumption problem given that information. We derive closed-form expressions for belief dynamics, and we obtain the value function and optimal policies under generalized preferences up to a tractable system of difference/differential equations. We contrast agents' welfare, policies, and incentives in the partial-information setting to the full-information setting.

The main finding is that, when uncertainty about a transition intensity increases, the left tail of the distribution becomes the dominant factor in economic decision-making. That is, even holding the mean belief constant, agents may act as if the transition probability is near zero. In forming expectations and evaluating trade-offs, households rationally place increasing weight on the possibility of the *current* state effectively lasting forever. In a setting where regimes are either "good" or "bad", the result can look like overreaction to a regime change relative to the full-information benchmark.

The mechanism driving this is that Bayesian updating implies *negative duration dependence*, i.e., that the longer the current state lasts, the longer it is expected to last. This dynamic is quite general: the absence of a transition in a given observation interval shifts the posterior density for the transition intensity towards zero regardless of the form of prior beliefs. Negative duration dependence is equivalent to transition times subjectively exhibiting a decreasing hazard function. Unconditionally, this means that beliefs about future regime durations are described by heavy-tailed distributions.

As precision declines, negative duration dependence increases, and the unconditional expected waiting time for the next transition, T, can become unbounded even holding fixed the mean belief about the transition probability. This is depicted graphically in Figure 1 when observers have belief about the transition intensity that is described by a gamma distribution with mean 1 and variance $1/a$. In the figure, λ denotes the probability per unit time of a disaster ending. The figure illustrates how, as precision declines, the expected waiting time, $E[T]$, becomes increasingly determined by the possibility that the true value of the intensity is close to zero.[1]

Figure 1: Information Precision and Expected Transition Time

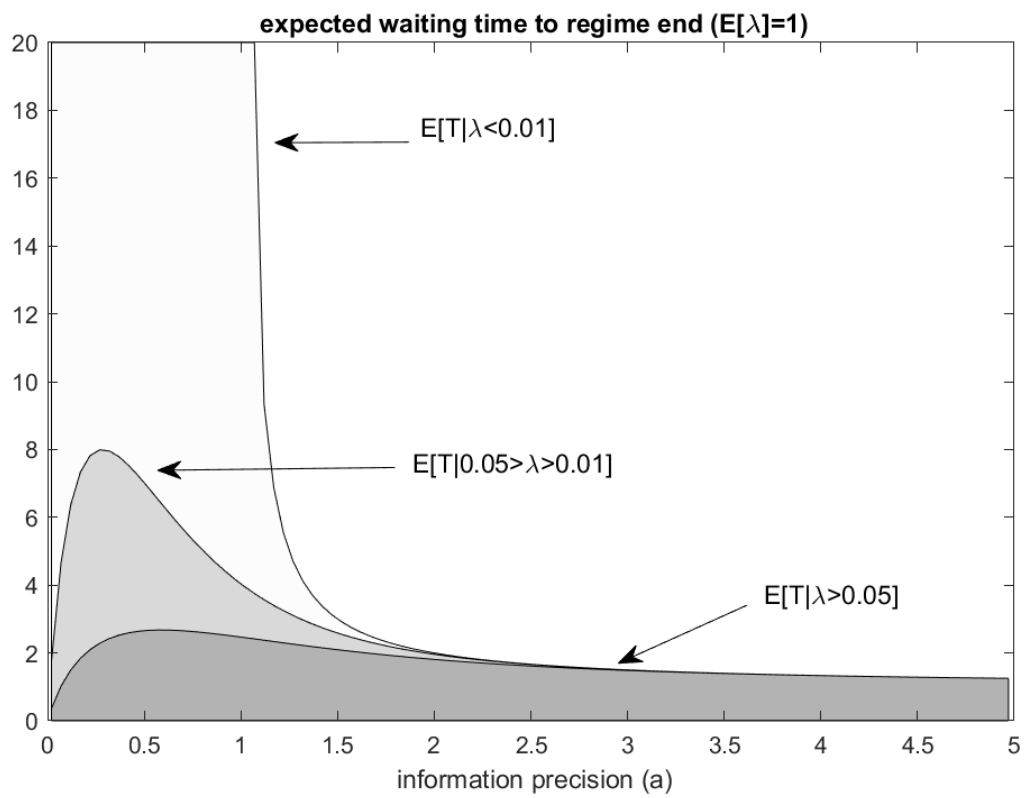

The top line plots the expected waiting time in years for the end of a regime when observers have belief about the intensity per unit time of a switch, λ, that are described by a gamma distribution with mean $E[\lambda] = 1$ and variance $1/a$, where a is the variable on the horizontal axis. The lower lines depict the contribution to this expectation of different components of the belief distribution.

A graphic example of a heavy-tailed waiting time distribution is seen in survey results

[1]These effects are also general: The pattern in Figure 1 is similar if the belief distribution is lognormal instead of gamma, for example.

from 2020 regarding the anticipated arrival time of an effective vaccine against the SARS-CoV-2 virus. The survey, conducted by Deutsche Bank, was sent to 800 global market participants eliciting their forecasts on four dates. As shown in Figure 2, as late as June of that year, and despite much positive clinical trial news, fully 4% of respondents thought the likely time to successful vaccine deployment would be *infinite.* This type of response can be rational under the structural uncertainty that we posit.

Figure 2: Duration Expectations: 2020

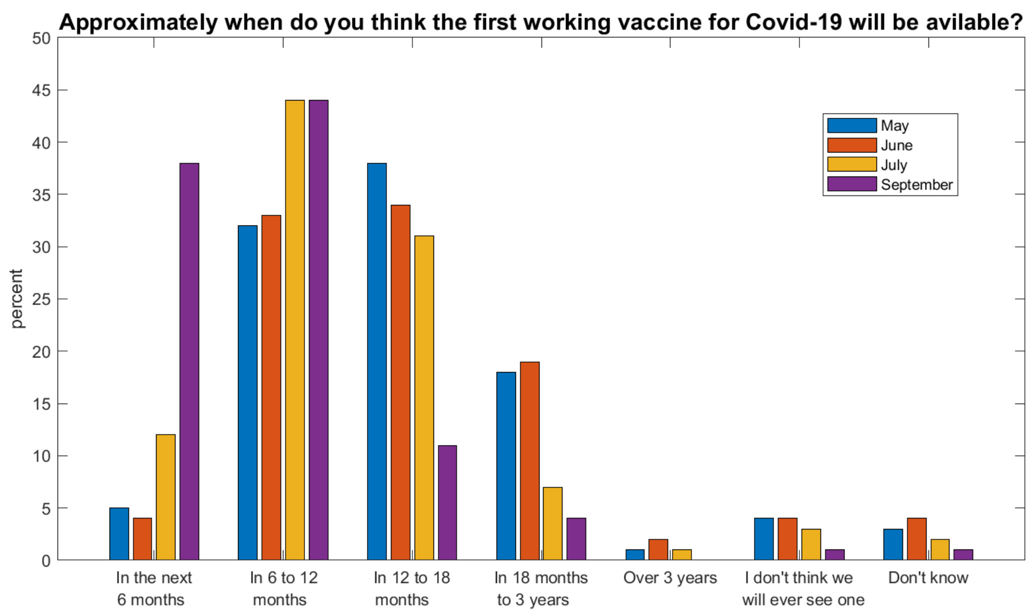

Source: Deutsche Bank, dbDIG.

The paper's analysis is based on comparisons of the representative agent's lifetime value function across states. We express these comparisons in terms of welfare costs, meaning the fraction of wealth the agent would be willing to pay to exchange one state for another. We first show that the benefit of ending a disaster is much higher with partial information compared to the full-information benchmark. The welfare gain from reducing the severity of an on-going disaster also increases strongly with parameter uncertainty. These welfare differences map directly to investment incentives. If the economy is augmented to include a mitigation technology, optimal investment in this activity at the onset of a disaster is increasing in parameter uncertainty.[2]

The observation that disasters are subjectively much worse under incomplete information raises the topic of the value of information. We show that the welfare gain from resolving parameter uncertainty – even without altering the disaster itself – can be as large as or larger than the benefit of ending the disaster. Imprecision acts as an amplification mechanism for perceived risk, leading agents to respond to a disaster with extreme conservatism in their investment/consumption decisions. As with mitigation incentives, our findings have implications for investment in information production. We consider endowing the economy with a technology to increase information precision and show that *marginal* value of information is very high when precision is low. From a policy perspective, reducing uncertainty about the evolution of disasters may be an important mechanism for alleviating their perceived harm.

The information dynamics of the model imply, however, that the welfare and incentive effects reverse prior to a disaster. Beliefs about the arrival rate also exhibit negative duration dependence, which increases when information is imprecise. Repeating the welfare computations in normal times, we show that information about the arrival rate can be welfare-destroying: agents may be subjectively better off with imprecise beliefs. This possibility occurs only for low risk-aversion, and high elasticity of intertemporal substitution. Turning again to incentives, when agents have the option to invest in a mitigation technology prior the onset of a disaster, we show that imprecise information about the disaster frequency induces less mitigation than under full information. For the same reason, agents with less precise information place lower value on disaster insurance.

Taken together, the model describes a belief dynamic across regimes that can lead to seemingly pessimistic behavior in bad times and optimistic behavior in good times. There is some empirical support for this implication. A well-established branch of behavioral economics documents the pattern that economic decision makers tend to ignore the risk of rare adverse events in good times and exaggerate them in bad times. The theory of *diagnostic expectations* has been advanced precisely to account for empirical evidence of this pattern. (See Bordalo et al. (2022) for a recent overview.) Our model presents a rational perspective also potentially consistent with this evidence.

[2] In a similar vein, Barnett et al. (2023) show that uncertainty about infectious parameters within a pandemic leads a central planner with ambiguity-averse preferences to impose stricter quarantine measures compared to the full-information benchmark.

1.1 Related Literature

The paper contributes a new insight to the literature that assesses the welfare costs of disaster risk (see Barro (2009), Pindyck and Wang (2013), Jordà et al. (2020), and Martin and Pindyck (2021)). Acharya et al. (2023) calibrate a version of the model studied here (with full information) to stock market responses to vaccine development news during 2020 in order to estimate the ex ante welfare cost of the pandemic.

A number of papers study learning problems in the context of models with disasters. Disasters are often parameterized as an exogenous shock process (hitting consumption or the capital stock) whose intensity is unobservable and possibly time-varying (for example, Benzoni et al. (2011) and Wachter and Zhu (2019)). We also have such a shock process, and its intensity varies over time: it is zero in normal times and positive in a disaster regime. However, we assume that agents do know the shock intensity.

In emphasizing uncertainty about persistence, our paper also shares similarities with Gillman et al. (2014) and Ghaderi et al. (2022) in which regimes of differing growth differ in their expected duration. These models assume the regime itself is unobservable. Another related work is Andrei et al. (2019) in which agents do not observe the mean-reversion speed of current consumption shocks and thus face persistence risk. In their model, as in ours, the persistence risk is asymmetric: information about persistence is positive in good times and negative in bad.

Collin-Dufresne et al. (2016) also study a 2-regime rare disaster economy with learning about the switching parameters. They show that, when risk-aversion exceeds the inverse of the elasticity of intertemporal substitution, even small amounts of persistence uncertainty can produce large effects on the equity premium and Sharpe ratio. The mechanism they highlight is the increase due to learning in the subjective volatility of consumption growth and marginal utility. While our setting is similar, the real effects we document are driven by the drift of the parameter estimates, not the volatility they induce.[3]

Most of the above papers focus on implications for asset pricing. An exception is Hong et al. (2023) who study implications of time-varying disaster beliefs for willingness to pay for mitigation efforts in the presence of externalities. Our focus too is on welfare effects. We highlight the interaction between unobservable persistence and the current state of the economy in determining the value of information and investment incentives.

[3] Also, on contrast, our main findings are *larger* in magnitude when the elasticity of intertemporal substitution is less than the inverse of the coefficient of risk-aversion.

Finally, our result that the welfare gain of information can be negative prior to a disaster for low enough risk-aversion and high enough intertemporal elasticity of substituon is seemingly at odds with the literature on parameter uncertainty in the context of Brownian motion shocks, notably **?** and **?**. This difference arises from the fact that in a model with learning about the risk of disasters or "regime changes," there is negative duration dependence, whereby the longer a disaster has not arrived, the higher is the expected waiting time. In general, this makes the distribution of beliefs about future regime durations heavy-tailed, inducing a behavior in agents as though disasters will almost never arrive. This, in turn, induces a preference not to resolve uncertainty as an adverse resolution of uncertainty can lead to a significant loss of utility from consumption smoothing between regimes.

2 Model

In this section, we introduce a regime-switching model of disasters under partial information. The goal is to study how the representative agent's value and policy functions vary with information precision.

2.1 Disaster Dynamics

Following Nakamura et al. (2013), we consider the state of the economy to be either in a "non-disaster" regime or in a "disaster" regime, and denote the state as $s \in \{0,1\}$. Let η denote the probability per unit time (or, intensity) of a disaster arrival, and let λ denote the probability per unit time of a disaster ending.

The model's depiction of the disaster consists of a state-specific stochastic process for the accumulation of wealth. Specifically, let q denote the quantity of productive capital of an individual household (which could be viewed as both physical and human capital). We assume that the stock of q is freely convertible into a flow of consumption goods at rate C per unit time. Then our specification is that q evolves according to the process

$$dq = \mu(s)qdt - Cdt + \sigma(s)qdB_t - \chi(s)qdJ_t \tag{1}$$

where B_t is a standard Brownian Motion and J_t is a Poisson process with intensity $\zeta(s)$. We set $\chi(0) = 0$ and $\chi(1) > 0$ for the disaster state. The Poisson shock captures the risk of an economic loss to the household. While we refer to the occurrence of the state $s = 1$ as the "disaster" (i.e., independent of whether or how many wealth shocks actually occur), somewhat more common in the literature would be to label these dJ shocks themselves as the "disasters", in which case our model maps to a particular specification of time-varying disaster risk, being either "on" or "off" depending on the regime. In Section 3 we will consider augmenting the economy to include real options to mitigate the disaster or acquire information.

2.2 Information Structure

Within a disaster there is likely to be deep uncertainty about *all* the governing parameters. Our focus on the timing parameters is motivated by the experience of COVID-19 in which the likely duration of the pandemic and the frequency of future pandemics were especially urgent questions to resolve. To model this, we assume the switching intensities η and λ are unobservable. While, formally, all disasters have the same parameters, this is not essential. Our main economic conclusions apply as well to the case in which parameter uncertainty re-sets with each new regime.

We will assume that at time zero the agent has beliefs about the two intensity parameters that are described by independent gamma distributions. Each distribution has a pair of non-negative hyperparameters, a^η, b^η and a^λ, b^λ, that are related to the first and second moments via

$$\mathbb{E}[\eta] = \frac{a^\eta}{b^\eta}, \qquad \text{Std}[\eta] = \frac{\sqrt{a^\eta}}{b^\eta}, \tag{2}$$

and likewise for λ. The *relative precision* about η, defined as its mean divided by its standard deviation, is $\sqrt{a^\eta}$.

By Bayes' rule, under this specification, as the agent observes the switches from one regime to the next, her beliefs remain in the gamma class with the hyperparameters updating as follows

$$\begin{aligned} a_t^\eta &= a_0^\eta + N_t^\eta \\ b_t^\eta &= b_0^\eta + t^\eta \end{aligned}$$

where t^η represents the cumulative time spent in state 0 and N_t^η represents the total number of observed switches from 0 to 1. Analogous expressions apply for a^λ and b^λ. Thus, while in $s = 0$, the only information that arrives (on a given day, say) is whether or not we have switched to $s = 1$ on that day. If that has occurred, the counter N^η increments by one and the clock t^η turns off and t^λ turns on. The system is assumed to start in the state $s = 0$ with $N^\eta = N^\lambda = 0$.

The model thus pastes together two linked learning regimes. In each regime, we have a finite dimensional filter in the sense that the two updated parameters fully characterize beliefs about that regime. Further, $\hat{\eta}_t \equiv \mathbb{E}_t[\eta] = a^\eta / b^\eta$, and it remains the case that the agent views this number as the probability per unit time of an instantaneous switch from $s = 0$ to $s = 1$ (again with equivalent expressions for the other regime.)

The gamma-exponential conjugate system is well studied in stochastic process theory (e.g., see Harris and Singpurwalla (1968) and Rubin (1972)). Under these beliefs, the measure for the switching time is a Lomax distribution (Lomax (1954)), whose expectation (in the $s = 0$ regime) is $1/\hat{\eta}$ times $a^\eta/(a^\eta - 1)$. This can be infinite when the relative precision of knowledge of η is low (as illustrated in Figure 1), reflecting the fat-tailed distribution of beliefs about future regime durations. As we will see, this has important consequences for agents' welfare and optimal behavior, even though it is not necessary for our results that the fat-tailedness lead to unbounded waiting times.[4]

2.3 Preferences

We assume the economy has a unit mass of identical agents (households). Each agent has stochastic differential utility or Epstein-Zin preferences Duffie and Epstein (1992) based on consumption flow rate C, given as

$$\mathbb{J}_t = \mathbb{E}_t \left[\int_t^\infty f(C_{t'}, \mathbb{J}_{t'}) dt' \right] \tag{3}$$

and aggregator

$$f(C, \mathbb{J}) = \frac{\rho}{1 - \psi^{-1}} \left[\frac{C^{1-\psi^{-1}} - [(1-\gamma)\mathbb{J}]^{\frac{1}{\theta}}}{[(1-\gamma)\mathbb{J}]^{\frac{1}{\theta}-1}} \right] \tag{4}$$

where ρ is the discount factor, γ is the coefficient of relative risk-aversion, ψ is the elasticity of intertemporal substitution (EIS), and $\theta \equiv \frac{1-\gamma}{1-\psi^{-1}}$.The representative agent's problem is to choose optimal consumption $C(s)$ that maximizes the objective function $\mathbb{J}(s)$.[5]

[4]As explained in the Introduction, the idea of negative duration dependence relies only on Bayesian updating and as such extends beyond the assumption of gamma distribution for beliefs about the intensity parameters. While the gamma distribution offers analytical tractability, we conjecture that our results hold more generally.

[5]We recognize the limitations of using a utility specification driven by consumption goods, particularly within a crisis when other considerations (e.g., health) may affect well-being. However, using a familiar formulation ensures that our findings are not driven by non-standard assumptions about utility.

2.4 Solution

Under the model's information setting, the economy is characterized by a six-dimensional state vector consisting of the stock of wealth, q, $a^\eta, b^\eta, a^\lambda, b^\lambda$ and the regime indicator S. However this six-dimensional space can be reduced to three when solving the agent's optimization problem.

Since the switches between states alternate, we can define an integer index M_t to be the total number of switches $N_t^\eta + N_t^\lambda$ and then $N_t^\eta = M_t/2$ when M is even, and $N_t^\lambda = (M_t+1)/2$ when M is odd. Knowing M (along with the priors a_0^η and a_0^λ) is equivalent to knowing a_t^η and a_t^λ. Given these values, specifying the current mean estimates $\hat{\eta}_t$ and $\hat{\lambda}_t$ is equivalent to specifying the remaining hyperparameters b_t^η and b_t^λ.

Within each regime the only changes to the state (apart from q) come through variation in the estimates $\hat{\eta}_t$ and $\hat{\lambda}_t$ which change deterministically with the respective clocks t^η and t^λ. Holding M fixed, the dynamics of $\hat{\eta}_t$ are given by

$$\begin{aligned} d\hat{\eta}_t &= d\frac{a_t^\eta}{b_t^\eta} = a_t^\eta d\frac{1}{b_t^\eta} \\ &= -\frac{a_t^\eta}{(b_t^\eta)^2}\,dt \\ &= -\frac{(\hat{\eta}_t)^2}{a_t^\eta}\,dt. \end{aligned} \tag{5}$$

The latter expression says that, until new information arrives, $\hat{\eta}$ decays quadratically and deterministically to zero at a rate that is faster when a^η is small. This dynamic defines the negative duration dependence of the system and illustrates its dependence on the degree of information precision.[6]

The agent's Hamilton-Jacobi-Bellman equation links the value functions for states with successively more history. For large M, the estimation errors for both η and λ go to zero:

$$\frac{\mathrm{Std}[\eta]}{\mathbb{E}[\eta]} = \frac{1}{\sqrt{a^\eta}} = \frac{1}{\sqrt{a_0^\eta + M_t/2}}.$$

[6]The ODE in (5) has the exact solution

$$\hat{\eta}_t = \frac{1}{\frac{1}{\hat{\eta}_0} + \frac{t}{a_0^\eta}}$$

where t is the time since the regime began.

Thus the system converges to the full-information solution, which is characterized by two coupled algebraic equations. Online Appendix establishes the following:

Proposition 1. *Let $H(\hat{\eta}, \hat{\lambda}, M)$ denote the solutions to system of coupled first-order differential equations in the Online Appendix. Assuming these are positive, optimal consumption is*

$$C = \rho^{\psi} \; (H)^{-\frac{\psi}{\theta}} q, \tag{6}$$

and the value function of the representative agent is

$$\mathbb{J} \equiv \frac{H(s) q^{1-\gamma}}{1-\gamma}. \tag{7}$$

The Online Appendix also describes an efficient solution algorithm for the system, and discusses conditions for existence of a unique positive solution.

3 Results

We now turn to numerical analysis to illustrate the model's effects. Our baseline calibration fixes the growth rate $\mu(s)$ and volatility $\sigma(s)$ across regimes to be 0.04 and 0.05. (The values are chosen to approximately capture the growth rate and volatility of aggregate dividends in non-disaster times.) The disaster shock size is set to $\chi = 0.04$ ($\chi(1) = 0.04$ and $\chi(0) = 0$, to be notationally precise). We fix the disaster shock intensity to be 1.0 in order to interpret χ as the expected loss of wealth per year. We use baseline preference parameters ($\gamma = 3.8$, $\psi = 2.0$, $\rho = 0.04$) that are broadly consistent with the macro-finance literature. The risk-aversion value we use is the one chosen by Gourio (2012). A number of papers, including Benzoni et al. (2011), Collin-Dufresne et al. (2016), and Andre, et.al (2018) choose a higher value of gamma of 10. The EIS parameter value that we have chosen is consistent with Gourio (2012), Benzoni et al. (2011), Collin-Dufresne et al. (2016) and Andrei et al. (2019). The crucial role of variation in preference parameters around these benchmark choices in the welfare calculations and the value of information is explored below.

3.1 Information Precision in a Disaster

To start, consider the welfare consequences of parameter uncertainty during a disaster. Since Lucas (1987), a large literature has analyzed the welfare costs of aggregate risks in business cycle models in order to quantify incentives to reduce them. Here, we extend this line of research to encompass the *perceived* risk that stems from parameter unobservability. We address two main questions. First, comparing partial information to full information, how much worse is the disaster compared to the non-disaster state? Second, how much would agents pay to gain information about the unknown parameters?

For any pair of economies or states, $\{i, j\}$, we report the fraction of wealth that the representative agent would be willing to pay for a one-time transition from the worse (j) to the better state (i). The welfare gain is computed as the certainty equivalent change in the representative agent's lifetime value function :

$$1 - \left(\frac{H(j)}{H(i)}\right)^{\frac{1}{1-\gamma}}$$

This definition is standard in the literature.

3.1.1 Welfare Gain from Curtailing a Disaster

To quantify the severity of disasters under our base parameterization, Table 1 reports the welfare gain from ending a disaster, that is, to transitioning from $s = 1$ to $s = 0$ holding everything else fixed. In the context of a pandemic, this could be viewed as the value of a perfectly effective cure or vaccine. Each cell of the table shows this gain for three values of $\hat{\lambda}$ and two values of $\hat{\eta}$. The top-left entries of Panel (A) shows the result when there is no uncertainty about the parameters. In this benchmark case, agents would be willing to pay between roughly 4% and 18% of wealth to return to the normal economic state. The values are intuitively reasonable in the sense that, for $\eta = 0.01$ say, they are not too far from just the expected length of the disaster ($1/\lambda$) times the expected loss of wealth per year, $\chi = 0.04$. In Panel (B), the top row shows that the preference parameters do not have large effects on the full-information values.

The top-right, bottom-left and bottom-right entries of Panel A show the same computation when agent has uncertainty about the timing parameters. These make the following assumptions about agent's current uncertainty: the posterior standard deviation of η is equal to its mean belief so that the relative precision of η is 1.0, whereas the posterior standard deviation of λ takes on three values such that the relative precision of λ is 10,000 or 10, or 1.0.[7]

Compared to the full information case (top-left entries of Panel A), the partial information situation is subjectively much worse. Adding parameter uncertainty greatly

[7] In the case where both relative precisions are one, the gamma prior is an exponential distribution.

increases the resources that the economy would be willing to expend to find a cure or otherwise limit the damage. An analogous computation (omitted for brevity) shows that the welfare benefit from lowering the disaster severity (χ) is also much larger under partial information. In Panel (B), the bottom row shows that the preference parameters can have large effects on the partial-information values, unlike the case for full-information values.

To see how these welfare differentials map into investment incentives, suppose now that the economy is endowed with a real option to undertake a lump-sum expenditure, I, to reduce the severity according to $\chi = g(I/q)$ for an arbitrary function $g > 0$ with $g' < 0$. By an argument that we formalize in the Online Appendix, the sensitivity of the welfare function H to χ effectively pins down the marginal benefit of I. Hence, for standard parameterizations of the mitigation technology, the optimal investment will be strictly greater under partial information than under full information.

3.1.2 Welfare Gain from Resolving Parameter Uncertainty in a Disaster

The results above immediately raise the question of how much agents would be willing to pay to resolve parameter uncertainty, even without curtailing the current disaster. Panel (A) of Table 2 answers this question by varying a_0^λ, the relative precision of λ, the likelihood of exiting the disaster, from 1.0 to 10, and for the case of relative precision of 1.0, varying separately the risk-aversion parameter γ to a lower value of 2.0 and the elasticity of intertemporal substitution ψ to a lower value of 0.2. For instance, when the relative precision of λ equals 1.0, the value of resolving the parameter uncertainty during a disaster is comparable in magnitude or even larger than the value of resolving the ongoing disaster (bottom-right entries of Panel (A) of Table 1).[8]

It is perhaps not surprising that risk-averse agents would be willing to pay to resolve parameter uncertainty. However, as we will see below, this need not always be the case *prior to a disaster*. Moreover, here, it is the *magnitude* of the value of information in a disaster that is surprising. The numbers are much larger than typically found in analogous calculations in the literature for other types of risk. In a similar setting, Collin-Dufresne et al. (2016) also show that, using a myopic utility benchmark, uncertainty about the persistence of the bad state is an order of magnitude more important than uncertainty about other parameters, e.g., growth rates and volatilities in the two regimes.

[8] The welfare gain is an understatement in that it excludes any value from, for example, information helping agents' ability to avert future disasters. The model contains no mechanism by which *knowing more about* λ and η allows agents to affect them.

Comparing the results in Panel (A) of Table 2 across preference specifications, the value of resolving parameter uncertainty increases with higher risk-aversion (γ). The γ effect is intuitive: parameter risk increases the subjective volatility of wealth, which agents dislike.[9] The largest effects come from lowering ψ, the elasticity of intertemporal substitution. This is noteworthy because there is a common understanding of Epstein-Zin preferences under which agents with $\psi \leq 1/\gamma$ can be viewed as having a preference for "later resolution of uncertainty," which might suggest that they value information *less* than high EIS agents, whereas here the result is the opposite.[10]

To understand this, note that, with recursive preferences, agents with low EIS cut consumption when the economy enters the disaster state. This is because a low EIS implies strong consumption smoothing motives, and the prospect of lower future wealth motivates a sharp increase in savings. By contrast, a higher EIS implies relatively more concern with the volatility of wealth than consumption smoothing. Agents with a high EIS therefore decrease investment in a disaster. However, the differing consumption responses do not make disasters worse *per se* for agents with low EIS: the top row of Panel (B) in Table 1 shows little effect of the EIS under full information. Instead, it is the extreme *decrease* in consumption as information precision declines that leads to the large welfare losses for these agents. This is again due to the time horizon effect. With low precision of information about λ, there is a chance that the withdrawal of consumption will be effectively permanent.[11]

As with mitigation, there is a direct mapping from the welfare costs of information to investment incentives. The findings above imply that the ability to produce information about the underlying determinants of disaster duration could be an extremely valuable real option. From a policy perspective, the implication of the large *marginal* value of information in a disaster is that fundamental research can crucially complement (or perhaps even substitute for) efforts to directly affect the course of the disaster.

[9]Likewise, we find that agents with a longer time horizon (lower subjective discount rate ρ) care strongly about persistence.

[10]See Epstein et al. (2014) for an analysis of the welfare consequences of varying the timing of the resolution of uncertainty.

[11]In Van Nieuwerburgh and Veldkamp (2006) and Kozlowski et al. (2020) learning effects within downturns endogenously cause the downturns to last longer. In our case, the uncertainty-induced investment and consumption distortions do not affect the length of the disaster. However, negative duration dependence implies that the *perceived* duration lengthens the longer the episode goes on.

3.2 Parameter Uncertainty Prior to a Disaster

The analysis above immediately suggests a corollary: all of the conclusions may be *reversed* prior to a disaster. Low precision of information about the disaster intensity in normal times could cause agents to give increasing weight to *best* case scenarios, namely, that a disaster will never materialize. This fat-tailed nature of expected waiting times interacts in crucial ways with preferences. When risk-aversion is low, agents do not save much by way of precaution ahead of a bad state. With high EIS, agents do not cut consumption much upon entering the bad state. Hence, the loss of utility associated with parameter uncertainty upon entering the bad state has a lower effect with low risk-aversion and high EIS (as shown in Panel (A) of Table 2). Then, if deadweight losses upon arrival of a disaster are not too high, then the agent can be better off prior to disaster without resolving uncertainty if the expected waiting time to a disaster is sufficiently high. This is because if information uncertainty resolves, then a resolution with high likelihood of a disaster arrival will induce significantly greater consumption smoothing and loss of agent's welfare. The agent is therefore better off with a high expected waiting time to a disaster.

We now show numerically that, indeed, this can be the case. Moreover, we will see that *both* types of effects – seemingly pessimistic in a disaster and optimistic beforehand – may co-exist.

3.2.1 Value of Information

We start by examining the welfare effect of uncertainty about η when $s = 0$. This effect can be isolated by setting the prior precision for λ to be very high, so that, effectively agents know its value. Top-left, bottom-left and bottom-right entries in Panel (B) of Table 2 show the value of information by varying preference parameters when a_0^λ, the relative precision of λ, is set to 10,000 (for sake of comparison, top-right entries set the relative precision of λ to 10). In the baseline case (top-left), the value of information about η indeed can be negative, although the magnitude is not always large. With $\gamma = 2$ (bottom-left) the effect can be economically significant: when the point estimate $\hat{\eta}$ is 0.05 the representative agent would be willing to give up to 2% of wealth to *not* learn the true disaster frequency. However, this counter-intuitive effect is not observed in the right set of entries in Panel (B) of Table 2 where either the relative precision of λ or the elasticity of intertemporal substitution ψ is low.

When information about both λ and η is imprecise, the former typically matters more in the sense that full information is overall welfare improving in both states. Intuitively, the worst-case scenarios still loom large prior to a disaster. However, we can vary the degree to which duration dependence operates in each regime by observing that the percentage drift in the means (which drives the effect) scales with the ratio of the mean to the precision. Thus, when $\hat{\eta}/a_0^{\eta}$ and $\hat{\lambda}/a_0^{\lambda}$ are similar, we obtain similar belief dynamics in the two states. The top panel of Figure 3 illustrates this co-existence of pessimism and optimism in terms of expectations of growth rate in wealth over different time horizons.

3.2.2 Disaster Mitigation Incentives

We saw above that, when information about disaster duration was imprecise, agents had stronger incentive to end or curtail the disaster. But that logic would also be expected to flip prior to a disaster. When agents place more weight on best-case scenarios, their incentives to invest in mitigation are weaker. To make this explicit, again consider endowing the economy with a one-time real option to expend resources to lower the disaster severity, χ. But now the investment decision is made prior to the onset of a disaster. We argued above that, for any given mitigation technology, the optimal amount invested will scale with the sensitivity of the value function to χ.

The lower panel Figure 3 plots the difference between full information and partial information cases in $\log H$ as a function of the disaster severity under full and partial information.[12] When $s = 1$ (right panel) we verify our assertion above that the slope of $\log(H)$ in χ is steeper under partial information. However, with these parameters, when $s = 0$ (left panel) the relation is reversed for low values of χ, and the magnitude of the difference in slope is numerically smaller throughout. Hence, for standard specifications of the mitigation technology, lower precision of information will result in underinvestment or underpreparedness relative to full-information in advance of a disaster.

Continuing the example, the slope differences between full information and partial information cases can imply economically important differences for the effect of information precision on mitigation. Consider the mitigation function $\chi(i) = \chi_0\, e^{-bi}$ where i is the fraction of wealth invested in mitigation technology. The real option is available one time for a small fixed cost c of wealth that represents the exercise price of the real option, i.e., the total cost of mitigation is $(i + c)$ of wealth.[13] The left panel of Figure 4 shows optimal investment prior to the disaster when the exercise price of the real option is 0.11% of wealth, $b = 3$ and $\chi_0 = 0.06$. For low precision, a^{η}, agents will not engage in

[12]Recall the full value function is negative, so higher values of H are worse.

[13]The mitigation technology and its costs are meant to merely be illustrative. Their calibration to parameters relevant in practice, and indeed the suitability of their functional forms, are left for future research.

mitigation at all. However, once a precision threshold is crossed, the option is exercised and investment jumps to over 2.25% of aggregate wealth resulting in a 40 basis points (or 6.67% from χ_0) reduction in the disaster severity (right panel). Hence small changes in parameter uncertainty may have important consequences for disaster mitigation incentives.[14]

3.2.3 Pricing of Disaster Insurance

Another way of capturing preparedness incentives is via willingness to pay for insurance against a disaster. Consider the price of a financial contract which pays 1 upon the arrival of the next disaster. This contract is in net zero supply and does not affect real outcomes. However, its price provides a measure of agents' assessment of the likelihood and timing of a disaster, as well as its consequences in marginal utility terms.

Proposition 2. *The price, P, in the non-disaster state of the claim which pays 1 upon the arrival of the next disaster, satisfies the equation*

$$-\frac{(\hat{\eta})^2}{a^{\eta}}\,\frac{\partial P}{\partial \hat{\eta}} \;+\; \hat{\eta}\,\frac{H(\hat{\eta},\hat{\lambda},M+1)}{H(\hat{\eta},\hat{\lambda},M)}\,(1-P) \;-\; r_0\,P = 0 \tag{8}$$

where r_0 is the riskless rate.[15]

Given the value function solutions, this is a first-order differential equation in $\hat{\eta}$, with boundary condition $P(0) = 0$. Figure 5 plots the solutions for the parameter set we have been considering. Since partial information entails longer expected waiting times, we see that, counterintuitively, the contract is substantially underpriced relative to its full-information value.

Summarizing, this section has shown that information about disaster frequency can be welfare-reducing prior to a disaster because, with less information, agents rationally believe a disaster may almost never materialize (the expected waiting time becomes unbounded or very large). This negative value of information may shed light on failure to prepare adequately for disasters and on "don't look up" behavior of seemingly willful ignorance towards their threat.[16]

[14] Imprecision may also exacerbate collective action problems. In this example, incorporating investment externalities can dramatically expand the no-action region.

[15] The rate and the pricing kernel are derived in terms of the model primitives in the Online Appendix.

[16] Models with costly information processing have also been used to explain failure to prepare for disasters. See Maćkowiak and Wiederholt (2018). Aversion to information is explicitly modelled in the preference specification of Andries and Haddad (2020).

4 Conclusion

This paper considers the economic effects of uncertainty about state transition probabilities. The main finding is that, as uncertainty increases, the left tail of the distribution becomes the dominant factor in decision-making and welfare, even holding the mean constant. Fully rational agents may act as if the current state will never end. In a setting where regimes are either "good" or "bad", the result can look like overreaction relative to the full-information benchmark, or pessimism in bad times and optimism in good times.

References

Viral V Acharya, Timothy Johnson, Suresh Sundaresan, and Steven Zheng. Vaccine progress, stock prices, and the value of ending the pandemic. *Management Science*, 2023. forthcoming.

Daniel Andrei, Michael Hasler, and Alexandre Jeanneret. Asset pricing with persistence risk. *The Review of Financial Studies*, 32(7):2809–2849, 11 2019.

Marianne Andries and Valentin Haddad. Information aversion. *Journal of Political Economy*, 128(5):1901–1939, 2020.

Michael Barnett, Greg Buchak, and Constantine Yannelis. Epidemic responses under uncertainty. *Proceedings of the National Academy of Sciences*, 120(2), 2023.

Robert J Barro. Rare disasters, asset prices, and welfare costs. *American Economic Review*, 99(1):243–64, 2009.

Luca Benzoni, Pierre Collin-Dufresne, and Robert S. Goldstein. Explaining asset pricing puzzles associated with the 1987 market crash. *Journal of Financial Economics*, 101(3): 552–573, 2011.

Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Overreaction in macroeconomics and diagnostic expectations. *Journal of Economic Perspectives*, 36(3):223–244, 2022.

Pierre Collin-Dufresne, Michael Johannes, and Lars A. Lochstoer. Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review*, 106(3): 664–98, March 2016.

Darrell Duffie and Larry G. Epstein. Asset pricing with stochastic differential utility. *Review of Financial Studies*, 5:411–436, 1992.

Larry G Epstein, Emmanuel Farhi, and Tomasz Strzalecki. How much would you pay to resolve long-run risk? *American Economic Review*, 104(9):2680–97, 2014.

Mohammad Ghaderi, Mete Kilic, and Sang Byung Seo. Learning, slowly unfolding disasters, and asset prices. *Journal of Financial Economics*, 143(1):527–549, 2022.

Max Gillman, Michal Kejak, and Michal Pakoš. Learning about Rare Disasters: Implications For Consumption and Asset Prices*. *Review of Finance*, 19(3):1053–1104, 05 2014.

Carl M Harris and Nozer D Singpurwalla. Life distributions derived from stochastic hazard functions. *IEEE Transactions on Reliability*, 17(2):70–79, 1968.

Harrison Hong, Neng Wang, and Jinqiang Yang. Mitigating disaster risks in the age of climate change. *Econometrica*, 91(5):1763–1802, 2023.

Òscar Jordà, Moritz Schularick, and Alan M Taylor. Disasters everywhere: The costs of business cycles reconsidered. Working Paper 26962, National Bureau of Economic Research, August 2020.

Julian Kozlowski, Laura Veldkamp, and Venky Venkateswaran. The tail that wags the economy: Beliefs and persistent stagnation. *Journal of Political Economy*, 128(8):2839–2879, 2020.

Kenneth S Lomax. Business failures: Another example of the analysis of failure data. *Journal of the American Statistical Association*, 49(268):847–852, 1954.

Robert Lucas. *Models of Business Cycles*. Blackwell, Oxford, 1987.

Bartosz Maćkowiak and Mirko Wiederholt. Lack of preparation for rare events. *Journal of Monetary Economics*, 100:35–47, 2018.

Ian WR Martin and Robert S Pindyck. Welfare costs of catastrophes: lost consumption and lost lives. *The Economic Journal*, 131(634):946–969, 2021.

Emi Nakamura, Jón Steinsson, Robert Barro, and José Ursúa. Crises and recoveries in an empirical model of consumption disasters. *American Economic Journal: Macroeconomics*, 5(3):35–74, 2013.

Robert S. Pindyck and Neng Wang. The economic and policy consequences of catastrophes. *American Economic Journal: Economic Policy*, 5(4):306–39, November 2013.

Izhak Rubin. Regular point processes and their detection. *IEEE Transactions on Information Theory*, 18(5):547–557, 1972.

Stijn Van Nieuwerburgh and Laura Veldkamp. Learning asymmetries in real business cycles. *Journal of Monetary Economics*, 53(4):753–772, 2006.

Jessica A Wachter and Yicheng Zhu. Learning with rare disasters. Wharton Working Paper, September 2019.

Table 1: Welfare Gain to Ending Disaster

(A) Benchmark preference parameters ($\gamma = 3.8, \psi = 2.0$)

Full Information

		$\hat{\lambda}$		
		0.2	0.5	1.0
$\hat{\eta}$	0.01	0.1789	0.0873	0.0451
	0.05	0.1414	0.0788	0.0431

Partial Information: $a_0^\lambda = 10{,}000$

		$\hat{\lambda}$		
		0.2	0.5	1.0
$\hat{\eta}$	0.01	0.1791	0.0874	0.0451
	0.05	0.1434	0.0790	0.0431

Partial Information: $a_0^\lambda = 10$

		$\hat{\lambda}$		
		0.2	0.5	1.0
$\hat{\eta}$	0.01	0.2194	0.1147	0.0624
	0.05	0.1715	0.1054	0.0640

Partial Information: $a_0^\lambda = 1$

		$\hat{\lambda}$		
		0.2	0.5	1.0
$\hat{\eta}$	0.01	0.3449	0.3045	0.2583
	0.05	0.2655	0.2545	0.2343

(B) Varying preference parameters

Full Information ($\gamma = 2.0, \psi = 2.0$)

		$\hat{\lambda}$		
		0.2	0.5	1.0
$\hat{\eta}$	0.01	0.1575	0.0781	0.0416
	0.05	0.1322	0.0716	0.0398

Full Information ($\gamma = 3.8, \psi = 0.20$)

		$\hat{\lambda}$		
		0.2	0.5	1.0
$\hat{\eta}$	0.01	0.2003	0.0915	0.0459
	0.05	0.1599	0.0840	0.0441

Partial Information: $a_0^\lambda = 1$ ($\gamma = 2.0, \psi = 2.0$)

		$\hat{\lambda}$		
		0.2	0.5	1.0
$\hat{\eta}$	0.01	0.3243	0.2624	0.2065
	0.05	0.2673	0.2364	0.2021

Partial Information: $a_0^\lambda = 1$ ($\gamma = 3.8, \psi = 0.20$)

		$\hat{\lambda}$		
		0.2	0.5	1.0
$\hat{\eta}$	0.01	0.8103	0.8363	0.8505
	0.05	0.6949	0.7455	0.7791

The table shows the fraction of wealth the agent would be willing to surrender for a one-time transition out of the disaster state. In Full information case, agents in the economy know the parameters λ and η. In Partial Information case, they have posterior standard deviation equal to their point estimates for η, the probability of switching into a disaster, i.e., $a_0^\eta = 1$, whereas for λ, the probability of switching out of a disaster, it is driven by the relative precision a_0^λ, as specified in the panels. The benchmark parameters are employed in Panel (A) as given in Section 3. Preference parameters are varied in Panel (B) as shown.

Figure 3: Optimism and Pessimism

(A) Parameter Uncertainty and Growth Forecasts

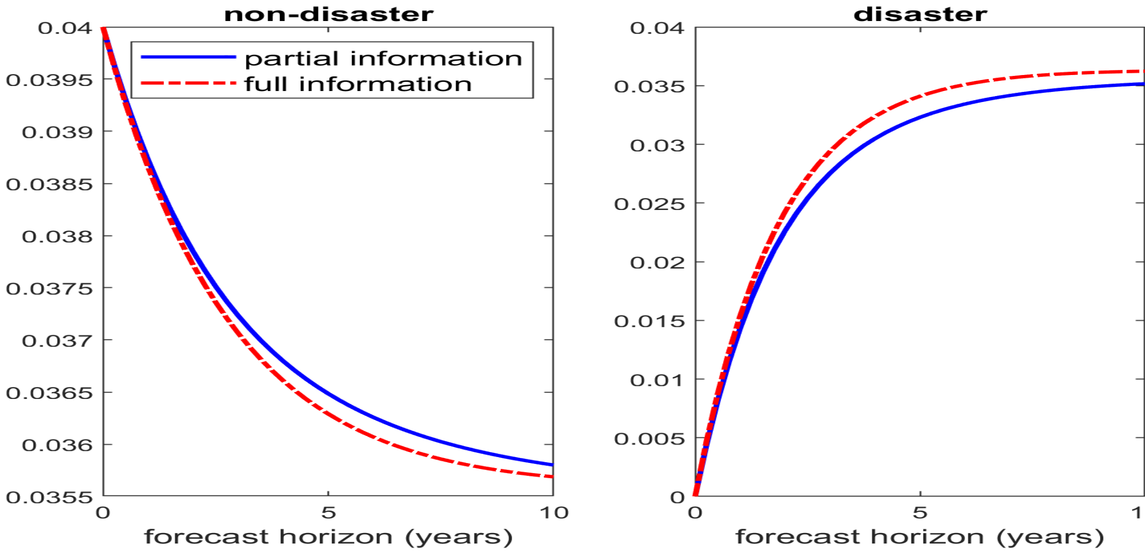

(B) Parameter Uncertainty and Mitigation Incentives

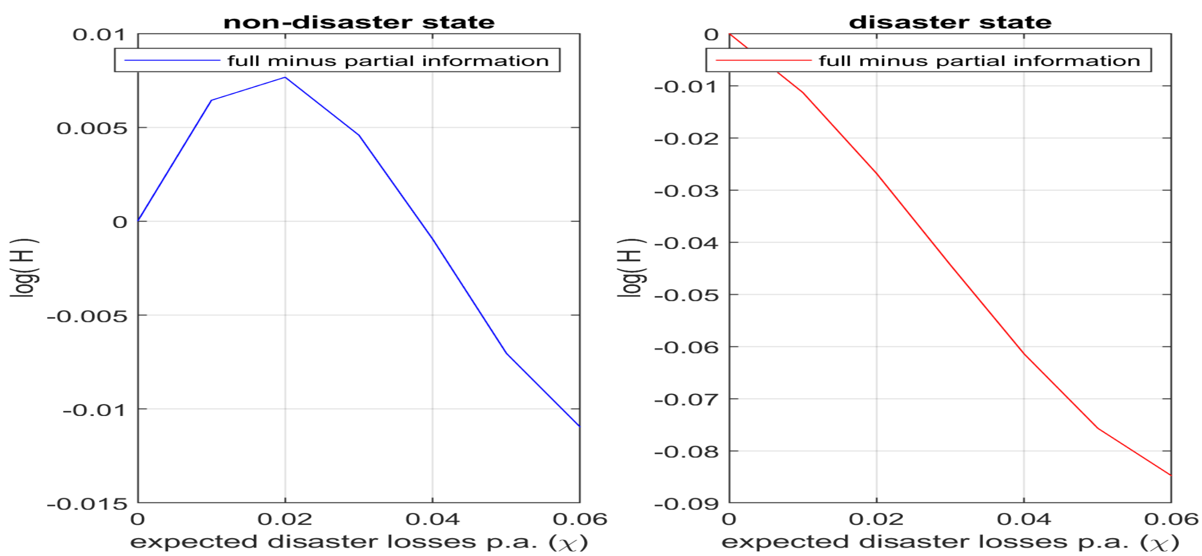

Panel (A) plots subjective expectations for the growth of wealth to different horizons, T. The left panel shows agents' forecasts when in normal times. The right panel shows forecasts during a disaster. The plots take the agent's posterior expected switching intensities for the two states to be $(0.05, 0.5)$ with respective posterior standard deviations of $(0.05, 0.05)$, so that the relative precisions are $a_0^\eta = 1.0$ and $a_0^\lambda = 10$. Panel (B) plots the log value function multiplier, H, as a function of the disaster severity χ also within the disaster (right) and nondisaster (left) states, assuming $a_0^\eta = 1.0$ and $a_0^\lambda = 10$. For each plot, the full-information economy's values are plotted as dotted (red) lines and the partial information ones as solid (blue) lines. Both panels use the benchmark preference parameters given in the text.

Figure 4: Parameter Uncertainty and Mitigation

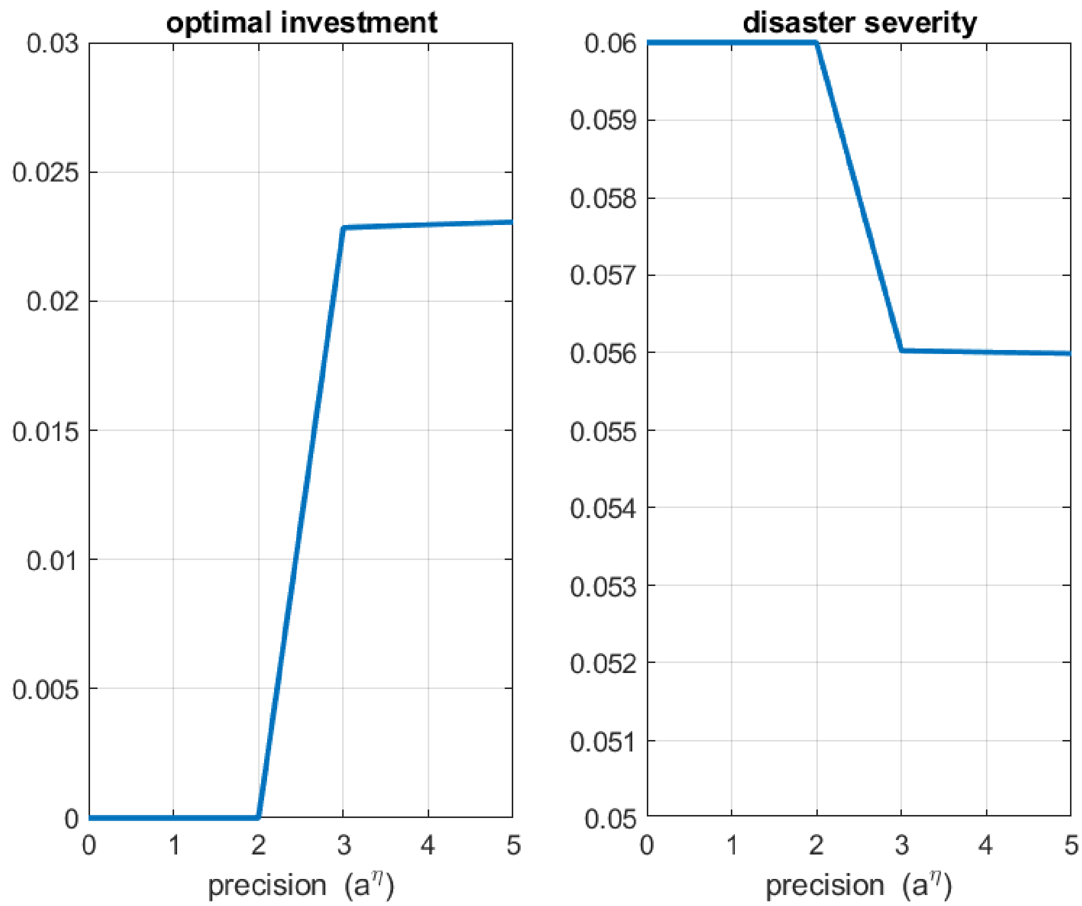

The figure plots the optimal investment mitigation (left) and resulting expected annual disaster losses (right) as a function of the precision of information about the arrival intensity, a^η. The economy has a one-time real option to spend a fraction $i + c$ of wealth to lower the disaster severity via $\chi = \chi_0 e^{-bi}$ where $\chi_0 = 0.06, b = 3$ and the fixed costs is $c = 0.11\%$. Parameter values are the same as in Panel (B) of Figure 3.

Figure 5: Disaster Insurance Pricing

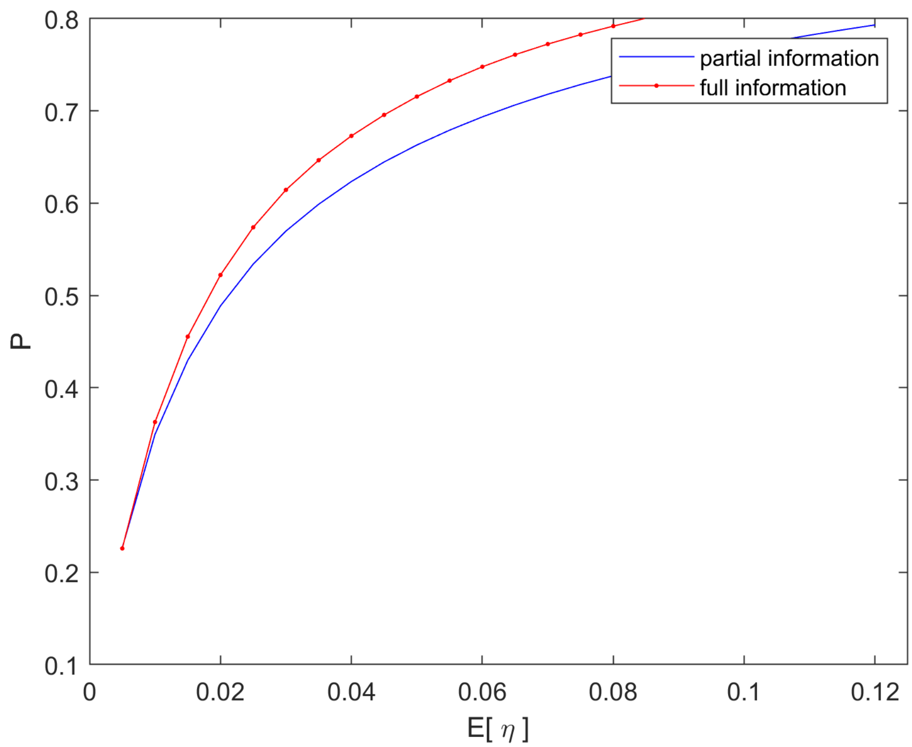

The figure plots the price of a contract paying 1 upon the arrival of the next disaster as a function of the mean arrival intensity, $\hat{\eta}$. Other parameter values are the same as in Panel (B) of Figure 3.

Table 2: The Value of Information

(A) During a Disaster

$a_0^\lambda = 1, \gamma = 3.8, \psi = 2.0$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.1716	0.1937	0.1686
	0.05	0.1278	0.1798	0.1776

$a_0^\lambda = 10, \gamma = 3.8, \psi = 2.0$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.0384	0.0201	0.0121
	0.05	0.0244	0.0247	0.0214

$a_0^\lambda = 1, \gamma = 2, \psi = 2.0$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.1432	0.1411	0.1149
	0.05	0.1125	0.1401	0.1298

$a_0^\lambda = 1, \gamma = 3.8, \psi = 0.20$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.9111	0.9043	0.8888
	0.05	0.8967	0.9010	0.8911

(B) Prior to Disaster

$a_0^\lambda = 10{,}000, \gamma = 3.8, \psi = 2.0$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	-0.0001	0.0003	0.0000
	0.05	-0.0121	-0.0026	-0.0003

$a_0^\lambda = 10, \gamma = 3.8, \psi = 2.0$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.0091	0.0049	0.0027
	0.05	0.0013	0.0109	0.0114

$a_0^\lambda = 10{,}000, \gamma = 2, \psi = 2.0$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	-0.0014	-0.0003	-0.0001
	0.05	-0.0200	-0.0073	-0.0026

$a_0^\lambda = 10{,}000, \gamma = 3.8, \psi = 0.20$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.0098	0.0017	0.0005
	0.05	0.0664	0.0211	0.0060

Panel (A) shows the fraction of wealth that the representative agent would be willing to surrender for a transition from partial information to full information (as defined in Table 1) while remaining in the disaster state. Panel (B) shows the fraction of wealth the agent would surrender for a transition from partial information to full information about the intensities η and λ. Under partial information, agents have posterior standard deviation equal to their point estimates for η, the probability of switching into a disaster, i.e., $a_0^\eta = 1$, whereas for λ, the probability of switching out of a disaster, it is driven by the relative precision a_0^λ, as specified in the panels. Benchmark parameters are given in Section 3 with variations as indicated in the table.