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

February 11, 2023

Abstract

We study an economy subject to recurrent disasters when agents have imprecise information about the frequency and duration of the disasters. Uncertainty about the persistence of states can lead to seemingly pessimistic behavior in bad times and optimistic behavior in good times. In a disaster, uncertainty about duration acts as an amplification mechanism. Agents alter their optimal investment and consumption more intensely relative to the full-information benchmark, and the welfare cost of parameter uncertainty can be extreme. However, in advance of a disaster, uncertainty about the arrival rate can be welfare-increasing and agents exhibit diminished preparedness: they 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 : welfare costs, rare disasters, parameter uncertainty, mitigation, overreaction

*We are grateful to Xavier Gabaix, Laura Veldkamp and Stijn van Nieuwerburgh for thoughtful comments.

1 Introduction

This paper studies the real effects of parameter uncertainty in a model of repeated disasters. Among the many terrifying aspects of the COVID-19 pandemic was the realization of *how little we knew* about what would happen. Structural uncertainty about the forces at work encompassed many dimensions, of course. We focus on two of these that seemed especially salient: uncertainty about the persistence (or duration) of the crisis, and uncertainty about its recurrence (or frequency) in future.[1] Such uncertainty appears pervasive and spans economic disasters beyond pandemics. For instance, structural uncertainty about the duration and frequency of recessions and financial crises is also realistic and potentially important, and likely to be as – if not more – relevant in the context of climate-related disasters.

Our model depicts disasters in reduced form as simply 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 optimally solve their investment/consumption problem, whose solution depends on both the current state and on current information about the unobserved switching parameters. We derive closed-form expressions for the information dynamics, and we obtain the value function and optimal policies under generalized preferences up to a tractable system of difference/differential equation. We contrast the solution to a representative agent's lifetime value maximization in the partial-information or parameter-uncertainty setting to the full-information setting.

Our primary aim is to understand when and how uncertainty about persistence affects agents' welfare and economic decisions. While economic intuition suggests that any source of uncertainty will be generically bad for risk-averse agents, we find a surprising range of effects, for which we provide an explanation. In particular, uncertainty about the persistence of states can lead to seemingly pessimistic behavior in bad times and optimistic behavior in good times. Hence, in a disaster, uncertainty about duration leads to an intensified response, or in other words, acts as an amplification mechanism. Agents alter their optimal investment and consumption conservatively relative to the full-information

[1] To be clear, when we refer to structural uncertainty we are distinguishing *not knowing the model* (specifically, the parameters of the data generating process) from the (ordinary) uncertainty of simply not knowing the outcomes, which are random variables.

benchmark, and the welfare cost of parameter uncertainty can be extreme. However, in advance of a disaster, uncertainty about the arrival rate can be welfare-increasing. Finally, agents may exhibit diminished preparedness against future disasters, i.e., optimally invest less in mitigation than under full-information, which also manifests as them willing to pay less for insurance against the next disaster.

Consider first the situation within a disaster. As a benchmark for how bad the disaster is, we report the amount of wealth an agent would be willing to give up to immediately end it and return to normal times. Then we perform the same calculation for the willingness to pay to resolve the parameter uncertainty, starting from a baseline in which agent's standard deviation of beliefs about the parameters is equal to their expectation. For a wide range of preference specifications, the welfare gain from removing parameter uncertainty is of the same order of magnitude as the gain from ending the disaster, and can even be substantially larger. The scale of the effect is much greater than the gains from eliminating other sources of uncertainty. Note that the calculation varies only the precision of the information, leaving mean beliefs about the switching intensities – and indeed the current state of being in the disaster – unchanged. The calculation is also an understatement in the sense that it includes no direct benefit to increased information precision in terms of, for example, improved design of interventions or mitigation strategies to cope with the disaster. Rather, the information effects that we identify stem only from changing agents' perception of the risks of their future consumption and wealth.

The paper thus contributes a new observation to the literature that assesses the welfare costs of disaster risk (see Barro (2009), Martin (2008), Pindyck and Wang (2013), Martin and Pindyck (2015), Jordà et al. (2020), Martin and Pindyck (2021), and Hong et al. (2022)). While modeling information production is beyond the scope of our work, there is a clear normative implication that fundamental research on the evolution of disasters (e.g., long-term epidemiological modeling in the case of pandemics or understanding the drivers of climate risks) may be even more valuable than is already appreciated.

To understand the mechanism behind the welfare benefit of information, we next describe the effect that parameter uncertainty has on agents' choices. Compared to the full information economy, we show that imprecision acts as an amplification mechanism for perceived risk, leading agents to respond to a disaster with extreme conservatism in their investment/consumption behavior. (If the elasticity of intertemporal substitution is less than one, this implies cutting consumption. For values greater than one, it implies cutting investment.)

The reason for this is that Bayesian updating implies that the waiting time to exit the disaster regime displays *negative duration dependence*: the longer the crisis lasts, the longer it is expected to last. As precision declines, this duration sensitivity increases, and the unconditional expected waiting time can become unbounded even holding fixed the mean belief about the probability of exiting. This is depicted graphically in Figure 1 when observers have belief about the switching intensity, λ, that is described by a gamma distribution with mean $E[\lambda] = 1$ and variance $1/a$. The figure illustrates how, for low values of the precision parameter, a, the expected waiting time becomes increasingly determined by "worst-case scenarios", i.e., the possibility that the true value of the exit intensity for the disaster is close to zero. Note that in the plot the explosive behavior of the expected exit time is due *only* to decreases in information precision. For all values of a, the agent has the same mean estimate of the instantaneous probability of an exit ($\mathrm{E}[\lambda] \equiv \hat{\lambda} = 1$ in the figure).

The expected waiting time is not literally an input into the computation of agents' optimal consumption or value function. However, understanding this feature of their beliefs explains what is happening economically within the dynamic program: they behave as if the disaster epoch may be highly persistent. From this observation, we can extrapolate beyond the welfare calculation, and infer the incentives that they would have for policies that we do not incorporate explicitly in the model. Specifically, we expect that their willingness to commit resources for mitigation during a disaster would rise sharply with the imprecision of their information.[2]

Another implication of the model, however, is that the reverse effects apply in advance of a disaster. That is, estimates of the mean arrival rate also exhibit negative duration dependence, which increases when information is imprecise. This, then, can also entail agents acting as if "best-case scenarios" predominate in their beliefs. In particular, uncertainty about the persistence of states can lead to seemingly pessimistic behavior (e.g., about growth rate forecasts) in bad times and optimistic behavior in good times (as illustrated in Figure 2).

We illustrate the implications of this result by repeating our welfare computations in normal times. We show that, in some settings, information about the arrival rate can be welfare-destroying: agents are subjectively better off with imprecise beliefs. We then consider an optimal investment problem where the representative agent can invest in mit-

[2]In a similar vein, Barnett et al. (2021) 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.

Figure 1: Information Precision and Expected Transition Time

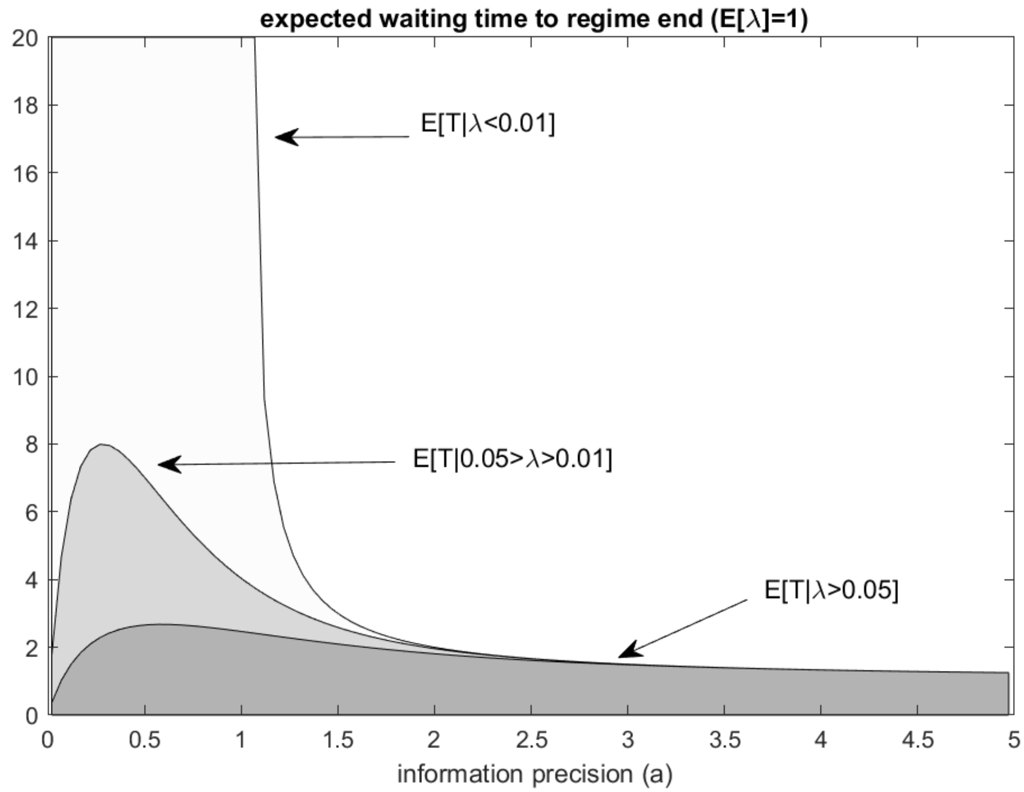

Note: The top line plots the expected waiting time in years for the end of a regime when observers have belief about the intensity per unit time of a switch, λ, that are described by a gamma distribution with mean $E[\lambda] = 1$ and variance $1/a$, where a is the variable on the horizontal axis. The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

igation measures prior the onset of a disaster. We show that, in these settings, agents with imprecise information about disaster frequencies also choose less mitigation than those with full information. As a parallel result, we introduce a disaster insurance contract, and show that information precision can lower the amount agents would pay for it. These results may shed light on "don't look up" behavior of seemingly willful ignorance towards disasters.

Taken together, the model describes a belief dynamic across regimes for which there is empirical support. In fact, a well established branch of behavioral economics takes as given the observation that economic decision makers tend to ignore the risk of rare adverse events in good times and exaggerate them in bad times (see Bordalo et al. (2022) for a recent overview). The theory of *diagnostic expectations* has been formulated precisely to

Figure 2: Growth Rate Forecasts: Optimism and Pessimism

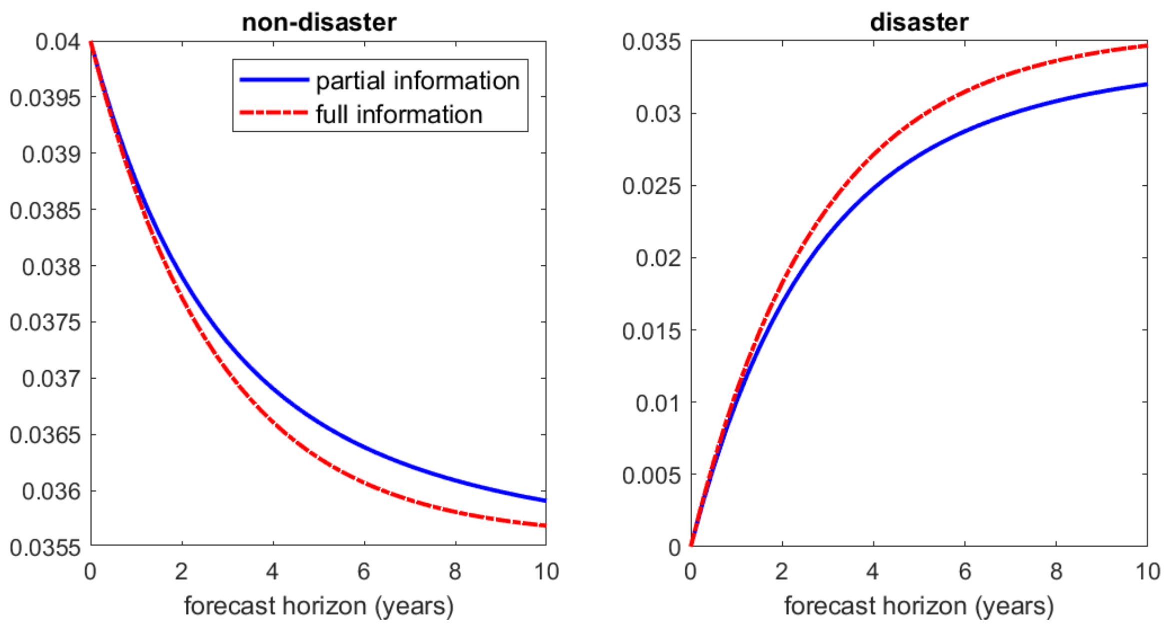

The figure plots subjective expectations for the growth of wealth to different horizons, T. The left panel shows agents' forecasts when in normal times. The right panel shows forecasts during a disaster. In each panel the full-information forecasts are plotted as dotted lines and the partial information ones as solid lines. The plots take the agent's posterior expected switching intensities for the two states to be $(0.05, 0.20)$ with respective posterior standard deviations of $(0.05, 0.10)$.

account for the evidence of this pattern. Moreover, in common with the implication in our model, that theory stresses that agents overreact *more* to recent news when it is more *salient*, which could be viewed as equivalent to settings in which there is less precision of prior information. While agents are not overreacting in our model, they would appear to be doing so to an observer with full information. Their consumption/investment behavior would appear increasingly optimistic prior to a crisis, and then increasingly pessimistic during one. In business cycle terms, their forecasts (for future output, say) would be highest at peaks and lowest at troughs. Survey evidence for such patterns verifies a testable prediction of our framework.

2 Related Literature

There are a number of papers that study learning problems in the context of models with disasters. It may be helpful to highlight distinguishing features of our setting and the focus of our contribution. A feature common to many models is an exogenous shock process (hitting consumption, or output, or the capital stock) whose intensity is unobservable and possibly time-varying. (Such models include Benzoni et al. (2011), Wachter and Zhu (2019), and Hong et al. (2022).) 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 parameter of these shocks. And, for the sake of clarifying terminology, we do not use the word "disaster" to refer to the realization of these shocks, but rather to the epoch in which they occur. In our case, it is the frequency and duration of these epochs about which agents lack full information.

Collin-Dufresne et al. (2016) study a 2-regime rare disaster economy in discrete time 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. In contrast, while our setting is similar, the real effects we document are driven by the drift of the parameter estimates (the duration dependence), not their revisions.[3]

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. Hence agents' beliefs about the persistence of current conditions is formed from a mixture over exponential distributions with posterior weights evolving with experience. In our model, agents do know whether or not they are in a disaster regime; but, in contrast, they do not know the switching intensities conditional on the regime. 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: increasing news about persistence is positive in good times and negative in bad. In Bianchi et al. (2022), agents are uncertain about the duration of monetary policy regimes.

Most of the above papers focus on the implications of their specifications for asset pricing. An exception is Hong et al. (2022) 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, in particular, the interaction between

[3]In addition, many of our findings are *larger* in magnitude when the elasticity of intertemporal substitution is less than the inverse of the coefficient of risk aversion.

unobservable persistence and the current state of the economy in determining the value of information and investment incentives.

3 A Model of Repeated Disasters

In this section, we introduce a regime-switching model of disasters in order to compute welfare in terms of the economy's primitive objects. Our fundamental view of a disaster is as a process that destroys household wealth, as in Gourio (2012), with consumption responding endogenously. For this reason, we work with a production-based framework rather than an endowment economy.

3.1 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 $s \in \{0,1,\}$. We assume that the economy switches between these states with the following transition probabilities:

$$Pr(s_{t+dt} = 1 | s_t = 0) = \eta dt \tag{1}$$

$$Pr(s_{t+dt} = 0 | s_t = 0) = 1 - \eta dt \tag{2}$$

$$Pr(s_{t+dt} = 0 | s_t = 1) = \lambda dt \tag{3}$$

$$Pr(s_{t+dt} = 1 | s_t = 1) = 1 - \lambda dt. \tag{4}$$

That is, η and λ are the probabilities per unit time of switching from one state to the other. These are the parameters that we will assume to be unobservable in later sections.

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, the latter reflecting health as well as intangible 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

$$dq = \mu(s) q dt - C dt + \sigma(s) q dB_t - \chi(s) q dJ_t \tag{5}$$

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 and normalize $\zeta(1) = 1$ for simplicity. 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), this is a matter of semantics. 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.[4]

An assumption worth highlighting concerns the long-run effects of the disaster. Our specification is pessimistic in the sense that loss of wealth due to the J shocks is permanent. Productive capital q does not get restored when the disaster ends. On the other hand, the model is optimistic in the sense that the productive *process*, dq, does fully revert to pre-disaster dynamics. After the disaster, the world looks stochastically the same as it did before. In particular, there are no long-run scarring effects, e.g., on the economy's growth rate, μ. Both assumptions are important for tractability. In Section 5.2 we will consider augmenting the economy to include a real option to invest in mitigation measures prior to a disaster that have the effect of lowering χ.

3.2 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; Duffie and Skiadas, 1994) based on consumption flow rate C, given as

$$\mathbb{J}_t = \mathbb{E}_t \left[\int_t^\infty f(C_{t'}, \mathbb{J}_{t'}) dt' \right] \tag{6}$$

and aggregator

$$f(C, \mathbb{J}) = \frac{\rho}{1-\psi^{-1}} \left[\frac{C^{1-\psi^{-1}} - [(1-\gamma)\mathbb{J}]^{\frac{1}{\theta}}}{[(1-\gamma)\mathbb{J}]^{\frac{1}{\theta}-1}} \right] \tag{7}$$

where $0 < \rho < 1$ is the discount factor, $\gamma \geq 0$ is the coefficient of relative risk aversion (RRA), $\psi \geq 0$ is the elasticity of intertemporal substitution (EIS), and

$$\theta \equiv \frac{1-\gamma}{1-\psi^{-1}} \tag{8}$$

The use of recursive preferences is standard in macro-finance models because of their ability to match financial moments. We recognize the limitations of using a utility speci-

[4]Besides Gourio (2012), important contributions to the literature on time-varying disaster risk include Gabaix (2012), and Tsai and Wachter (2015).

fication driven by consumption goods, particularly within a crisis when other considerations (e.g., health, social interaction, the safety of others) so strongly affect well-being. However, using a familiar formulation ensures that our findings are not driven by non-standard assumptions about utility.

The representative agent's problem is, in each state s, to choose optimal consumption $C(s)$ that maximizes the objective function $\mathbb{J}(s)$.

3.3 Full information solution

We now characterize the solution to the optimization problem when agents know the switching intensities η and λ. For ease of notation, define the following combination of preference parameters:

$$e_0 \equiv \frac{\theta}{\psi}\,\rho^{\psi} \quad \text{and} \quad e_1 \equiv -\frac{\psi}{\theta}. \tag{9}$$

Proposition 1. *Denote*

$$g(s) \equiv \theta\,\rho - (1-\gamma)\left(\mu(s) - \frac{1}{2}\gamma\sigma(s)^2\right) - \zeta(s)\,\left(\left[1-\chi(s)\right]^{1-\gamma} - 1\right) \tag{10}$$

for $s \in \{0,1\}$. Let $H(s)$'s denote the solution to the following system of recursive equations:

$$g_0 \equiv g(0) = e_0\,(H(0))^{e_1} + \eta\left[\frac{H(1)}{H(0)} - 1\right] \tag{11}$$

$$g_1 \equiv g(1) = e_0\,(H(1))^{e_1} + \lambda\left[\frac{H(0)}{H(1)} - 1\right] \tag{12}$$

Assuming the solutions are positive, optimal consumption in state s is

$$C(s) = \rho^{\psi}\;(H(s))^{e_1}\,q, \tag{13}$$

and the value function of the representative agent is

$$\mathbb{J}(s) \equiv \frac{H(s)q^{1-\gamma}}{1-\gamma}. \tag{14}$$

Note: All proofs appear in the appendix.

The recursive system is straightforward to solve numerically. The unknown functions $H(s)$ are necessarily bounded by the limiting solutions in which the economy is never in a disaster, H_0^{min}, or is always in a disaster, H_1^{max}.[5] The former corresponds to $\eta = 0$ and the latter to $\lambda = 0$. It is straightforward to show that these constants are given by

$$H_0^{min} = \left(\frac{g_0}{e_0}\right)^{1/e_1} \qquad \text{and} \qquad H_1^{max} = \left(\frac{g_1}{e_0}\right)^{1/e_1}.$$

These quantities are real and positive if g_0, g_1, and e_0 all have the same sign. Given this, it can be shown that a necessary and sufficient condition for existence of a unique solution is that $g_1 < g_0$. Henceforth we implicitly assume the parameters are such that these regularity conditions are satisfied. It is worth noting that the parameters defining the exogenous wealth process only affect the system (and hence its solution) through the constants g_0 and g_1. So, while the model may seem to involve a large number of parameters, its effective dimensionality is low.

4 Solution under Parameter Uncertainty

4.1 Information Structure

Recall that in our model η is the intensity of switching from state 0 ("good") to state 1 ("bad") and λ is the intensity of switching from 1 to 0. In this section, we assume that agents have imperfect information about these intensities. As discussed in the introduction, 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.

Let us stipulate that at time zero the agent has beliefs about the two intensity parameters that are described by *gamma* distributions, which are independent of each other. Each gamma 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

$$\mathbb{E}[\eta] = \frac{a^\eta}{b^\eta}, \qquad \text{Std}[\eta] = \frac{\sqrt{a^\eta}}{b^\eta}, \tag{15}$$

and likewise for λ. The *relative precision* about η, defined as its mean divided by its standard deviation, is $\sqrt{a^\eta}$. The choice of the gamma family is motivated by tractability, as it

[5]Note that with $\gamma > 1$, the value function is negative. Hence smaller (positive) values of H correspond to better states.

induces a finite-dimensional representation of agents' beliefs. However, the exact functional form does not drive our results. A property that is important is that distribution is not bounded away from zero: agents are not able to rule out *a priori* that the switching intensities could be arbitrarily small.

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 $N^\eta = N^\lambda = 0$.

We thus paste 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.)

This type of gamma-exponential conjugate system is well studied in stochastic process theory (e.g., see Harris and Singpurwalla (1968) and Rubin (1972)), and has some important properties. Most notably, under the observer's subjective probability measure, the system exhibits negative duration dependence, meaning the longer the *experienced* waiting time for a switch, the longer the *expected* remaining waiting time (since the estimated switching intensity grows weaker). As a result, unlike a pure exponential system (driven by a Poisson process with known intensity), the subjective expected waiting time until the next switch is not the inverse of the (expected) intensity, but is larger. In fact, the explicit measure for the switching time is described by a Lomax distribution (Lomax (1954)), whose expectation (in the $s = 0$ regime) is $1/\hat{\eta}$ times $(a^\eta/(a^\eta - 1)$. Note that this can be infinite when the relative precision of knowledge of η is low (as illustrated in Figure 1). Similarly the variance of the waiting time explodes for low precision. As we will see, these features have important consequences for the agents' welfare and optimal behavior in our economy.

4.2 Partial-Information Solution

Under the above information structure, 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 actually be reduced to three when solving for the agent's value function.

Since the switches between states alternate, let us define an integer index M_t to be the total number of switches $N_t^\eta + N_t^\lambda$ and then (assuming we are in state 0 at time 0) $N_t^\eta = M_t/2$ 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. Thus, solutions to the model can be described as a sequence of functions $H_M(\hat{\eta}, \hat{\lambda})$ for the agent's value function at step M.

Compared to the full-information model in Section 3, within each regime the only new changes to the state 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

$$\begin{aligned} d\hat{\eta}_t = d\frac{a_t^\eta}{b_t^\eta} &= a_t^\eta d\frac{1}{b_t^\eta} \\ &= -\frac{a_t^\eta}{(b_t^\eta)^2}\,dt \\ &= -\frac{(\hat{\eta}_t)^2}{a_t^\eta}\,dt. \end{aligned} \tag{16}$$

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 drives the main results.[6]

Under partial information, we proceed as in Section 3.3 to write-out the HJB equation with the state variables following the dynamics determined by the representative agent's

[6]The ODE in (16) 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.

information set. As before, we can conjecture a form of the value function

$$\mathbb{J} = \frac{q^{1-\gamma}}{1-\gamma}\, H(\hat{\eta}, \hat{\lambda}, M; C). \tag{17}$$

And, as before the first order condition for consumption yields $C = q\ (\rho^{\psi})\ H^{e_1}$ (where e_1 is defined in (9)). This follows because consumption does not enter into any of the new terms involving the information variables.

Using these results, it can be shown that (See the Appendix for a derivation of (18)-(19)):

Proposition 2. *The HJB system characterizing the value function in the presence of parameter uncertainty can be written as a set of infinite-dimensional linked PDEs, where M runs over the even integers, with the constants g_0 and g_1 are as defined in Section 3.3:*

$$g_0 = e_0 H_M^{e_1} + \hat{\eta}\left(\frac{H_{M+1}}{H_M} - 1\right) - \frac{(\hat{\eta})^2}{a^{\eta} H_M}\,\frac{\partial H_M}{\partial \hat{\eta}} \tag{18}$$

$$g_1 = e_0 H_{M+1}^{e_1} + \hat{\lambda}\left(\frac{H_{M+2}}{H_{M+1}} - 1\right) - \frac{(\hat{\lambda})^2}{a^{\lambda} H_{M+1}}\,\frac{\partial H_{M+1}}{\partial \hat{\lambda}}. \tag{19}$$

For large M, the estimation errors for both η and λ, expressed as a fraction of the posterior estimates, go to zero:

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

The system always converges to the full-information solution, providing one boundary condition, which, together with the single-regime solutions on the edges of the $(\hat{\eta}, \hat{\lambda})$ plane, enables computation of all individual H functions. Knowing the solution for higher M enables direct evaluation of the jump-terms in (18)-(19). Knowing the solution on the inner edges enables explicit approximation of the first partial derivatives.

5 Results

We now turn to numerical analysis to illustrate the effect of information precision on the economy's properties. We will highlight, in particular, differences in these effects depending on the current state.

To undertake the analysis, we specify a benchmark set of parameters described by the values shown in Table 1. These can be partitioned into preference parameters and those specifying the (pre-consumption) dynamics of wealth, q. The baseline preference parameters are broadly consistent with the macro-finance literature under stochastic differential utility. For the q dynamics, we choose to fix the growth rate $\mu(s)$ and Gaussian volatility $\sigma(s)$ across regimes in order to focus attention on the role of the disaster shock size χ. (The former values are chosen to approximately match the growth rate and volatility of aggregate dividends in non-disaster times.) We fix the disaster shock intensity to be 1.0 in order to interpret χ as the expected loss of wealth per year.

5.1 Information Precision and Welfare in a Disaster

To start, we examine the welfare consequences of parameter uncertainty within 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 such risks. 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 precision about the persistence parameters? We also explore which features of the model make imprecision more costly.

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 willingness-to-pay definition is standard in the literature.

5.1.1 Welfare Gain from Ending a Disaster

To quantify the severity of disasters under our parameterization, Table 2 reports the welfare gain for 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 panel shows the result when there is no uncertainty about the parameters. Here the upper left cell shows that, in this benchmark case, agents would be willing to pay between roughly 5% and 20% 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 duration* of the disaster $(1/\lambda)$ times the expected loss of wealth per year, $\chi = 0.04$. When disasters are expected to reoccur more frequently ($\eta = 0.05$) the welfare gain is smaller. Reading across the top panel, changing the preference parameters does not have large effects on the the full-information values.

The bottom panel shows the same computation when agents' current uncertainty about the timing parameters (their posterior standard deviation) is equal to their mean belief about each of them, or their relative precisions are 1.0 for both. This is our baseline case of partial information.[7] Compared to the top panel of the table, it is clear that the partial information situation is subjectively much worse. Adding parameter uncertainty greatly increases the resources that the economy would be willing to expend to find a cure or otherwise curtail the damage.

Among the preference parameters, we note that the information effects are more extreme with lower EIS (ψ) and lower time discount factor (ρ). Also, perhaps surprisingly, the relative effect is strongest when the estimated switching intensity to return to the normal state $\hat{\lambda}$ is higher. This is counterintuitive because one might think information about timing would matter less when disasters are shorter-lived. We return to these effects below.

5.1.2 Welfare Gain from Resolving Parameter Uncertainty

The results above immediately raise the question of how much agents would be willing to pay to resolve parameter uncertainty, even staying within the disaster state. Table 3 answers this question. In fact, for each of the preference configurations considered and for almost all values of $\hat{\eta}$ and $\hat{\lambda}$, the value of resolving the parameter uncertainty is as large or larger than the value of resolving the present disaster under full information.

It is perhaps not surprising that risk averse agents would be willing to pay something to resolve parameter uncertainty. However, as we will see below, this need not always be the case. Moreover, here, it is the *magnitude* of the value that is surprising. The numbers are much larger than typically found in analogous calculations in the literature for other types of risk. As one example, in the current model, when we compute the welfare gain to eliminating the stochastic nature of disaster shocks[8] we find welfare gains on the order

[7] In this case the gamma prior is just an exponential distribution. There is also nothing critical about the value 1.0: the policies and value functions are all continuous in the precision parameter. Hence the numerical results are similar for higher initial precision.

[8] This entails dropping the Poisson jump term in the dq specification and lowering the disaster drift, $\mu(1)$, by the expected jump magnitude, $\zeta\chi$.

of 1 percent or smaller, much less than those in Table 3. We are not the first to make this observation. Collin-Dufresne et al. (2016) show that, measured by the one-period utility loss compared to an adaptive expectations 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.

From a policy perspective, the implication of the extreme value of information in a disaster is that, while working to curtail the disaster itself is enormously valuable, equally and perhaps even more valuable is investment and effort to uncover fundamental knowledge about the underlying mechanisms, including biological, physical, and sociological, governing the duration of these episodes. Note too that the welfare gain that the model identifies does not even consider the direct effect that such knowledge has on our ability to prevent or mitigate disasters. The model contains no channel by which *knowing more about* λ and η allows agents to change them.

5.1.3 Comparative Statics

Returning to Table 3, comparing the results across preference specifications, the information value – the value of resolving parameter uncertainty – increases with higher risk-aversion (γ), and is lower with a lower time discount factor (ρ). The γ effect is intuitive: parameter risk increases the subjective volatility of wealth, which agents dislike. Less apparent is the effect of ρ. Why should agents with a longer time horizon (lower subjective discount rate) care so strongly about information? Recall that when agents do not know the true value of λ, their expected time until the end of the disaster is governed by a Lomax distribution whose expectation explodes as the precision of their information declines. Put differently, with parameter uncertainty, *worst case scenarios* come into play. When there is little experience from which to learn, the current disaster may look effectively permanent. Hence, its impact on welfare can be enormous when discount rates are low.

This effect also explains why the results in the table show surprisingly little sensitivity to $\hat{\lambda}$. As noted above, this seems counterintuitive. Here the key observation is the seemingly paradoxical fact that the explosion of the expected waiting time for the end of the disaster happens *for any level* of belief about the intensity of switches. Again, this can be viewed as due to increased weight being placed on worst case scenarios: agents cannot rule out that $\lambda \sim 0$, i.e., that the disaster will effectively last for their entire lifetime.

The largest effects in Table 3 come from lowering 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 precisely the opposite.[9]

To understand this, we observe 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 investment risk than consumption smoothing. Agents with a high EIS therefore decrease investment in a disaster, since investment is exposed to greater risk. So, the propensity to consume out of wealth increases for these agents during a disaster. These responses are plotted in the top panel of Figure 3.

However, the differing consumption responses do not make disasters worse *per se* for agents with a low EIS: the top panel of Table 2 shows little effect of the EIS under full information. Instead, as the plot in the lower panel of Figure 3 shows, 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, the withdrawal of consumption starts to look effectively permanent.[10]

To summarize, the crucial effect of parameter uncertainty within a crisis is that, with less precise information, agents place more weight on increasingly bad possibilities, or low values of λ. This leads to very large welfare losses and consumption/investment distortions, which are stronger when agents have longer time horizons (low ρ) or less intertemporal elasticity of substitution (low ψ). It is worthwhile to remark that the effects are mostly driven by agents' beliefs about the *current* disaster: the worst-case scenario is that it lasts forever.

5.2 Parameter Uncertainty Prior to a Disaster

The analysis above immediately suggest an unexpected corollary: all of the conclusions may be *reversed* prior to a disaster. Low precision of information about the disaster intensity in normal times could lead to agents acting as if they overweight *best* case scenarios,

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

[10] 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.

Figure 3: Consumption and the Elasticity of Intertemporal Substitution

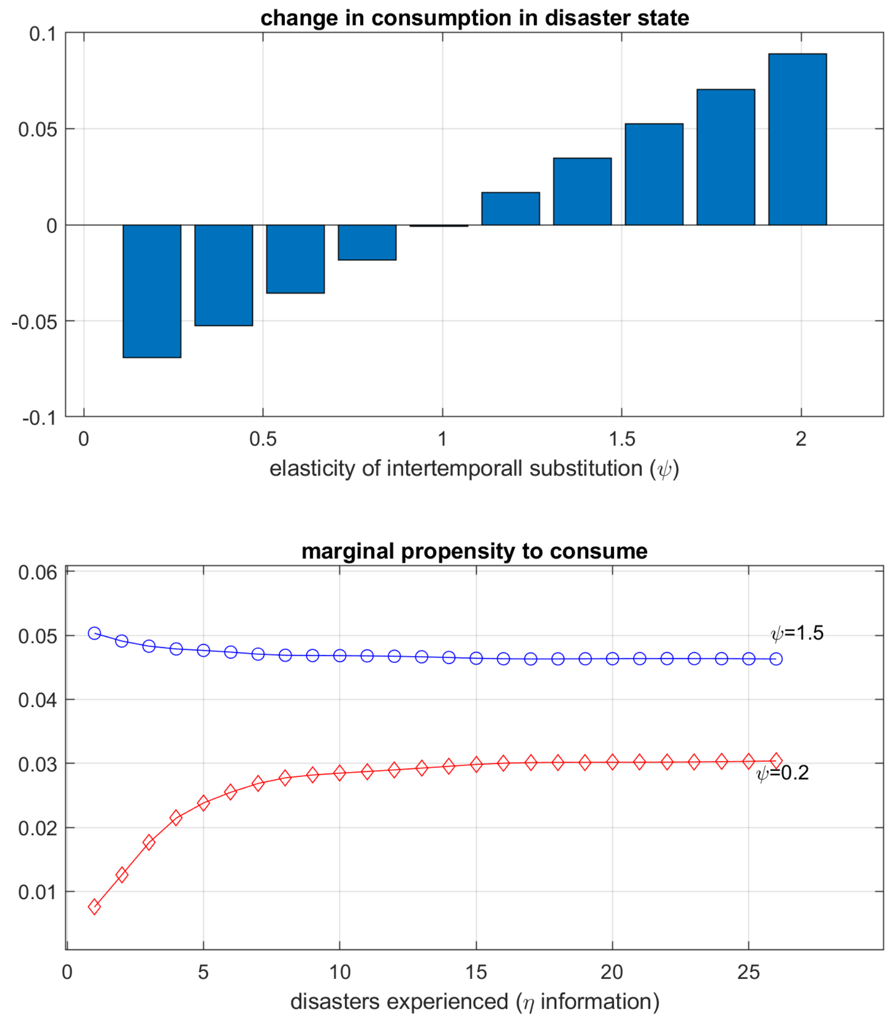

The top panel plots the ratio, $c(1)/c(0)$ of the marginal propensity to consume in states 1 (disaster) and 0 (non-disaster) states under full information. The bottom panel plots the respective propensities in successive disaster states, starting from the prior distributions given by $a^{\eta} = a^{\lambda} = 1$. Both panels use the benchmark parameters from Table 1. In both panels, the current beliefs are $\hat{\eta} = 0.03, \hat{\lambda} = 0.5$.

namely, that a disaster will never materialize. We now show that, indeed, this can be the case. Moreover, we will see that *both* types of effects – pessimistic in a disaster and optimistic beforehand – may co-exist.

5.2.1 Welfare Gain from Resolving Parameter Uncertainty

To set the stage, we start by examining the welfare effect of uncertainty about η when $s = 0$. We isolate this effect by setting the prior precision for λ to be very high, so that, effectively agents know its value. Table 4 shows the value of information under these conditions. In the baseline case, the value of information about η indeed can be negative, although the magnitude is not large economically. Working against the effect of longer subjective waiting time until a disaster is the effect of risk aversion: the value function is concave in $\hat{\eta}$, so higher posterior variance lowers welfare through this channel. The third panel in the table exhibits that with $\gamma = 2$ the effect can be economically significant: when the point estimate $\hat{\eta}$ is large the representative agent would be willing to give up to 2.1% of wealth to *not* learn the true disaster frequency.

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 from the observation that the percentage drift in the means (which drives the effect) scales with the ratio of the mean to the precision parameter. Thus, when $\hat{\eta}/a_0^{\eta}$ and $\hat{\lambda}/a_0^{\lambda}$ are close in size, we obtain similar belief dynamics in the two states. Figure 2 in the introduction illustrates this co-existence of pessimism and optimism in terms of growth rate expectations. Using the parameters in that figure together with $\gamma = 1$, the welfare cost of parameter uncertainty is 3.2% of wealth in the disaster and -3.5% before it. Hence, the incentives to acquire information alternate sign in the two states.

5.2.2 Disaster Mitigation Incentives

We observed above that, when information about disaster duration was imprecise, agents had stronger incentive to end the disaster (c.f. Table 2). But that logic would now also be expected to flip. When agents place more weight on best-case scenarios, their incentives to invest in mitigation are weaker. To make this explicit, consider endowing the economy with a one-time real option to expend resources, I, out of the stock of wealth, q, to invest in a mitigation technology that can lower the severity, χ, of future disasters via $\chi = g(i)$, where $i = I/q$, for some positive smooth function g with $g' < 0$. Parameterizing mitigation technology is beyond the scope of this work. However, in this setting, the strength of the welfare effects of mitigation map directly to the degree of optimal investment for any technology. Specifically, if the (log) value function coefficient H is less sensitive to χ under partial information than under full information, then it is straightforward to show that optimal investment is lower as well.[11]

Figure 4 plots the log value function as a function of the disaster severity under full and partial information. The figure uses the same parameters as in Figure 2 with γ close to one.[12] When $s = 0$ in the left panel, we verify the conjecture that the slope is less positive, and the relation is reversed when $s = 1$ in the right panel. We can conclude that, for any reasonable investment technology, lower precision of information will result in underinvestment in mitigation or preparedness relative to full-information in advance of a disaster.

5.2.3 Pricing of Disaster Insurance

Another way of capturing preparedness incentives is via willingness to pay for insurance against a disaster. So 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. Then,

Proposition 3. *The price, P, in the non-disaster state of the claim which pays 1 upon the arrival of the next disaster, satisfies the equation*

$$-\frac{(\hat{\eta})^2}{a^{\eta}}\,\frac{\partial P}{\partial \hat{\eta}} \;+\; \hat{\eta}\,\frac{H(\hat{\eta},\hat{\lambda},M+1)}{H(\hat{\eta},\hat{\lambda},M)}\,(1-P) \;-\; r_0\,P = 0 \tag{20}$$

where r_0 is the riskless rate.[13]

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

[11] The assertion is that, for two otherwise equal economies E1 and E2, if the sensitivity of the value function, H, to χ is weaker in E1 than in E2, then, if a solution to the real-options problem exists in E2, a solution also exists in E1 with smaller optimal investment. To see this, view H as a function of χ, and the problem is to choose i to maximize the $H(g(i))(1-i)^{1-\gamma}/(1-\gamma)$ with first order condition $-g'(i)\,\partial \log H(g(i))/\partial\chi = (\gamma-1)/(1-i)$. Assume $\gamma > 1$. Then the right side (the marginal cost) is an unbounded increasing function of i on $[0,1)$ which is the same for both economies. Call it $RHS(i)$. On the left side (the marginal benefit), the first term is the same for both economies. The hypothesis is that $\partial \log H(\chi)/\partial\chi$ is smaller in E1 than in E2 for all χ implying that the second term is smaller. Hence $LHS1(i) < LHS2(i)$ for all i. Assume $LHS2$ is continuous and declining. Then, if an interior solution, i_2^*, exists, it follows that on $[i_2^*, 1)$ we have $LHS1 < LHS2 < RHS$, meaning that there cannot be a solution for E1 in this region. Hence, either there is a solution $i_1^* < i_2^*$ or no interior optimum exists and $i_1^* = 0$ in E1.

[12] The plot use take $\gamma = 1.01$ for convenience. Recall the full value function is negative, so higher values of H are worse.

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

Figure 4: Parameter Uncertainty and Mitigation Incentives

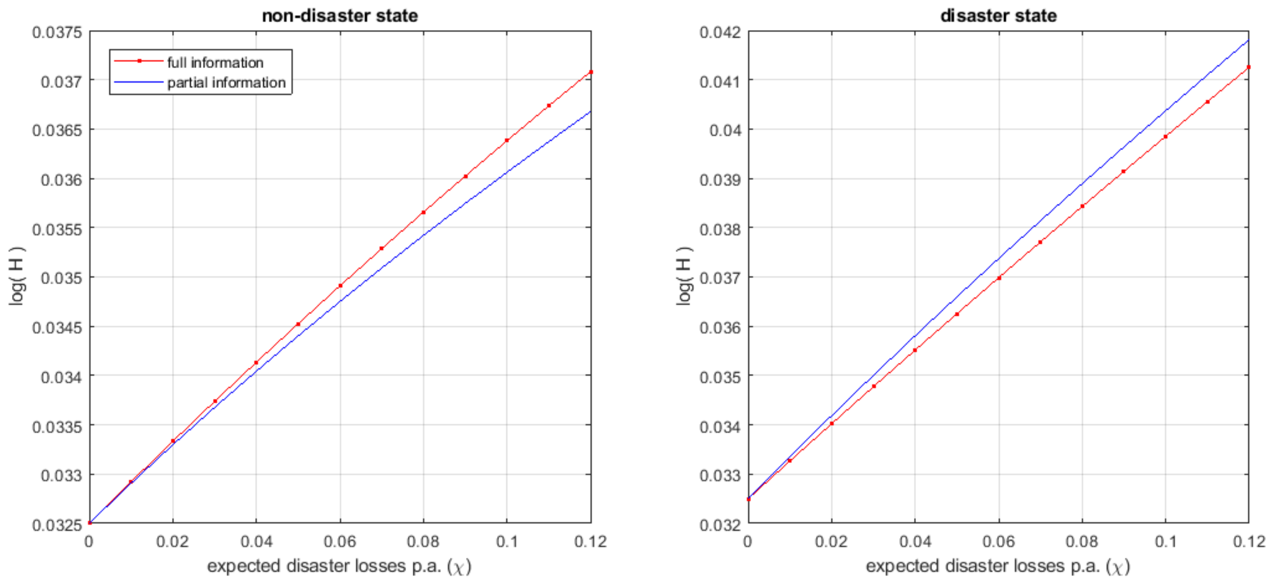

The figure plots log value function multiplier, H, as a function of the disaster severity. χ. In each panel, the full-information values are plotted as dotted lines and the partial information ones as solid lines. The plots use the benchmark parameters in Table 1 with $\gamma = 1.01$. The plots take the agent's posterior expected transition intensities for the two states to be $(0.05, 0.20)$ with respective posterior standard deviations of $(0.05, 0.10)$.

been considering. In line with the intuition that partial information leads to longer expected waiting times, we see that the contract is substantially underpriced relative to its full information value.

This section has shown that information about disaster frequency can be welfare reducing because, with less information, agents rationally believe a disaster may never materialize (the expected waiting time becomes unbounded) even when the mean intensity of disasters is held fixed. The other phenomena that we have illustrated (optimistic forecasts, underinvestment in mitigation, undervaluing insurance) are all manifestations of the same belief dynamic. The negative value of information may shed light on failure to prepare adequate for disasters and on "don't look up" behavior of seemingly willful ignorance towards their threat.[14]

[14]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).

Figure 5: Disaster Insurance Pricing

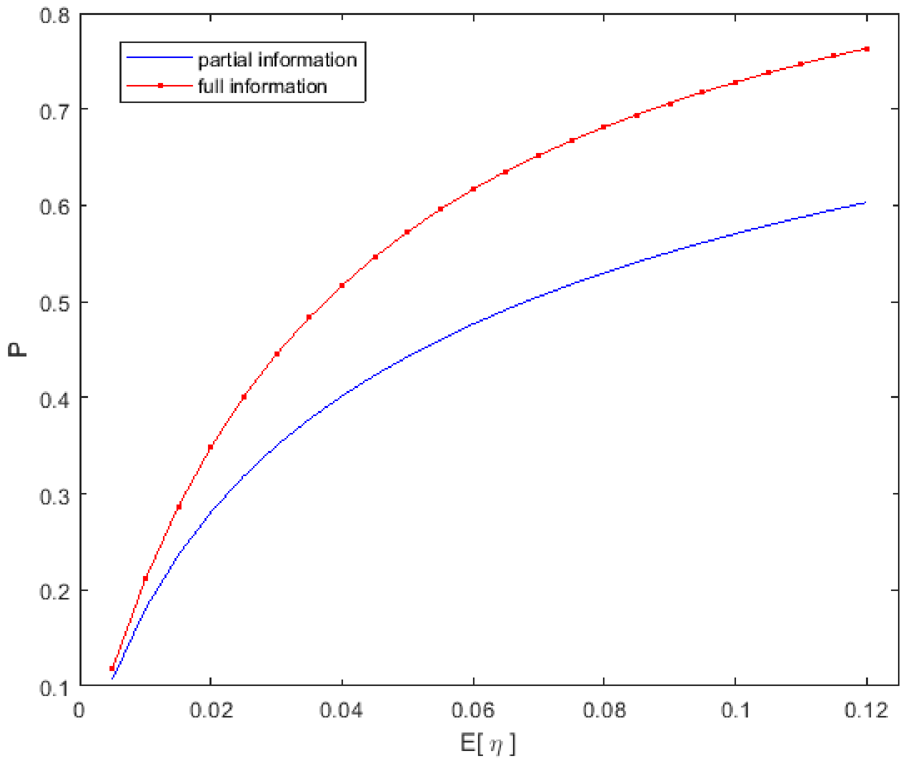

The figure plots the price of a contract paying 1 upon the arrival of the next disaster as a function of the mean arrival intensity, $\hat{\eta}$. Other parameter values are the same as in Figure 4.

6 Conclusion and Discussion

Motivated by the ongoing COVID-19 pandemic, we augment a standard regime-switching model to include uncertainty about the duration of both the disaster and the non-disaster states. We find a perhaps surprising dichotomy: the welfare benefit of information about persistence of the state can be extreme within a disaster and small or even negative prior to one. The mechanism behind the welfare effects is that imprecision heightens the possibility that the true transition probability is small. Hence, agents in the economy act, in effect, as if the current state may never end, even holding fixed their estimate of the instantaneous probability of it ending. This can result in seemingly exaggerated pessimistic behavior within a disaster (e.g., extreme reductions in consumption) at the same time as seemingly exaggerated optimistic behavior (e.g. reduced investment in mitigation or insurance) in non-disaster times.

In the introduction, we cited the behavioral economics literature that documents the tendency of managers and investors (as well as professional forecasters) to overweight recent experience in forming their beliefs. We believe evidence for this may also be interpreted as consistent with the specific type of parameter uncertainty that we study here. In particular, the literature advancing diagnostic expectations stresses the pattern of decision-makers ignoring "tail events" in good times and exaggerating their effects when they occur. In those models, unlike ours, investors are incorrectly overweighting current data, for which there is evidence in survey data.[15] Rather than rebut this alternative to our story, however, we would suggest that our conclusions overall are likely to go through – or even be strengthened – under it. An agent in a crisis who places *even more* weight on worst case scenarios than does a rational Bayesian will be willing to pay even more than the magnitudes we find in order to rule them out. Likewise, the subjective well-being of such an agent in good times is similarly likely to be even less improved by ruling out best case scenarios or in expending resources on mitigation.

[15]Some of the patterns in the survey data can also be explained by models of learning about low frequency features of the data generating process. See Farmer et al. (2021).

References

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. Working Paper 27289, National Bureau of Economic Research, April 2021.

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.

Francesco Bianchi, Martin Lettau, and Sydney C Ludvigson. Monetary policy and asset valuation. *The Journal of Finance*, 77(2):967–1017, 2022.

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.

Darrell Duffie and Costis Skiadas. Continuous-time security pricing: A utility gradient approach. *Journal of Mathematical Economics*, 23:107–132, 1994.

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.

Leland Farmer, Emi Nakamura, and Jón Steinsson. Learning about the long run. Working Paper 29495, National Bureau of Economic Research, December 2021.

Xavier Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics*, 127(2):645–700, 2012.

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.

François Gourio. Disaster risk and business cycles. *American Economic Review*, 102(6): 2734–66, May 2012.

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. Technical report, National Bureau of Economic Research, February 2022.

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. Disasters and the welfare cost of uncertainty. *American Economic Review*, 98(2):74–78, 2008.

Ian WR Martin and Robert S Pindyck. Averting catastrophes: the strange economics of scylla and charybdis. *American Economic Review*, 105(10):2947–85, 2015.

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.

Jerry Tsai and Jessica A. Wachter. Disaster risk and its implications for asset pricing. *Annual Review of Financial Economics*, 7(1):219–252, 2015.

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: Parameter Values

Parameter	Symbol	Value
Coefficient of relative risk aversion	γ	4.0
Elasticity of intertemporal substitution	ψ	1.5
Rate of time preference	ρ	0.04
Expected growth of wealth	$\mu(0) = \mu(1)$	0.04
Volatility of wealth	$\sigma(0) = \sigma(1)$	0.05
Destruction intensity in disasters	ζ	1.0
Disasters shock magnitude	χ	0.04

The table shows the parameter values for preferences and the capital stock dynamics that form the baseline case in the numerical analysis.

Table 2: Welfare Gain to Ending Disaster

Low Uncertainty

Benchmark

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.188	0.090	0.046
	0.05	0.147	0.081	0.044

$\psi = 0.20$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.205	0.093	0.046
	0.05	0.162	0.085	0.045

$\gamma = 2$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.160	0.079	0.042
	0.05	0.135	0.072	0.04

$\rho = 0.02$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.225	0.100	0.048
	0.05	0.166	0.091	0.047

High Uncertainty

Benchmark

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.294	0.212	0.14
	0.05	0.212	0.169	0.128

$\psi = 0.20$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.715	0.716	0.704
	0.05	0.559	0.585	0.610

$\gamma = 2$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.242	0.154	0.095
	0.05	0.195	0.137	0.094

$\rho = 0.02$

		$\hat{\lambda}$ 0.2	0.5	1.0
$\hat{\eta}$	0.01	0.466	0.436	0.392
	0.05	0.297	0.292	0.287

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 the top four panels, agents in the economy know the parameters λ and η. In the bottom four panels they have posterior standard deviation equal to their point estimates of these quantities. The benchmark parameters are given in Table 1.

Table 3: The Value of Information

	Benchmark				$\psi = 0.20$			
		$\hat{\lambda}$				$\hat{\lambda}$		
		0.2	0.5	1.0		0.2	0.5	1.0
$\hat{\eta}$	0.01	0.170	0.161	0.114	0.01	0.865	0.828	0.767
	0.05	0.129	0.162	0.139	0.05	0.851	0.832	0.786
	$\gamma = 2$				$\rho = 0.02$			
		$\hat{\lambda}$				$\hat{\lambda}$		
		0.2	0.5	1.0		0.2	0.5	1.0
$\hat{\eta}$	0.01	0.114	0.094	0.064	0.01	0.564	0.571	0.488
	0.05	0.095	0.106	0.086	0.05	0.404	0.514	0.504

The table shows the fraction of wealth the agent would be willing to surrender for a transition from partial information to full information (as defined in Table 2) while remaining in the disaster state. The benchmark parameters are given in Table 1.

Table 4: The Value of Information about Disaster Frequency

	Benchmark				$\psi = 0.20$			
		$\hat{\lambda}$				$\hat{\lambda}$		
		0.2	0.5	1.0		0.2	0.5	1.0
$\hat{\eta}$	0.01	0.001	0.001	0.000	0.01	0.012	0.002	0.001
	0.05	-0.008	-0.000	0.001	0.05	0.079	0.026	0.007
	$\gamma = 2$				$\rho = 0.02$			
		$\hat{\lambda}$				$\hat{\lambda}$		
		0.2	0.5	1.0		0.2	0.5	1.0
$\hat{\eta}$	0.01	-0.001	-0.000	-0.000	0.01	0.050	0.015	0.004
	0.05	-0.021	-0.007	-0.002	0.05	0.044	0.047	0.023

The table shows the fraction of wealth the agent would be willing to surrender for a transition from partial information to full information about the disaster intensity η while remaining in the non-disaster state. The agent is assumed to have full information about λ. The benchmark parameters are given in Table 1.

Appendix

A Proofs and Derivations

A.1 Proof of Proposition 1

Proof. For ease of notation, define $\lambda(0) = \eta, \lambda(1) = \lambda$. Using the evolution of capital stock for the representative agent (5) the Hamilton-Jacobi-Bellman (HJB) equation for each state s can be written:

$$\begin{aligned} 0 = \max_{C} \Big[& f(C, \mathbb{J}(s)) - \rho \mathbb{J}(s) + \mathbb{J}_q(s)(q\mu(s) - C) \\ & + \frac{1}{2}\mathbb{J}_{qq}(s) q^2 \sigma(s)^2 + \zeta(s) \left[\mathbb{J}(s)\left(q(1-\chi(s))\right) - \mathbb{J}(s)(q)\right] \\ & + \lambda(s) \left[\mathbb{J}(s')(q) - \mathbb{J}(s)(q)\right]\Big] \end{aligned} \tag{A.1}$$

for $s = \{0,1\}$ and $s' = \{1,0\}$.

Taking the first-order condition with respect to $C(s)$ in (A.1), we obtain

$$f_c(C, \mathbb{J}(s)) - \mathbb{J}_q(s) = 0. \tag{A.2}$$

Using $f(C, \mathbb{J})$ from (7) and taking the derivative with respect to C, we obtain

$$f_c = \frac{\rho C^{-\psi^{-1}}}{[(1-\gamma)\mathbb{J}(s)]^{\frac{1}{\theta}-1}}. \tag{A.3}$$

Substituting the conjecture $\mathbb{J}(s)$ in equation (14) yields

$$f_c = \frac{\rho C^{-\psi^{-1}}}{H(s)^{\frac{\gamma-\psi^{-1}}{1-\gamma}} q^{\gamma-\psi^{-1}}}. \tag{A.4}$$

Then, for state s, we obtain by substituting $\mathbb{J}_q(s) = H(s)q^{-\gamma}$ in (A.2), and simplifying:

$$C(s) = \frac{H(s)^{-\theta\psi^{-1}} q}{\rho^{-\psi}} \tag{A.5}$$

which agrees with (13) using the definitions of the constants in (9).

To verify the conjectured form of the value function, we plug it in to the HJB equation (A.1) and reduce it to the recursive system in the proposition via the following steps:

1. substitute the optimal policy $C(s)$ into the HJB equation (A.1);

2. cancel the terms in q which have the same exponent; and

3. group constant terms not involving Hs into $g(0)$ for state 0 and $g(1)$ for state 1.

After some rearrangement, we reach equations (10) – (12).

□

A.2 Proposition 2: HJB System with Parameter Uncertainty

Proof. As noted in the text, the model can be parameterized in terms of the state variables $M, \hat{\eta}, \hat{\lambda}$, and q, where $M = M_t$ is an integer counter that increases on a state switch such that $M_0 = 0$ and even numbered states are the non-pandemic epochs and odd numbered states are the pandemics. Also, in the non-pandemic states, $\hat{\lambda}$ is constant, while $\hat{\eta}$ is constant in pandemics. As a consequence, compared with the derivation above for the full-information case, there is now only one additional source of variability in each regime. The dynamics of $\hat{\eta}$ are given in (16) with an analogous expression for and $\hat{\lambda}$. And note that, under the agents' information set, the dynamics of the wealth variable q are identical to the full information dynamics.

As a result, the HBJ equations under partial information are the same as (A.1) above (with state 0 and state 1 being replaced by M and $M+1$) with the addition of a single term on the right side:

$$-\frac{(\hat{\eta})^2}{a^{\eta}} \frac{\partial \mathbb{J}(0)}{\partial \hat{\eta}} \tag{A.6}$$

for $s = 0$, and

$$-\frac{(\hat{\lambda})^2}{a^{\lambda}} \frac{\partial \mathbb{J}(1)}{\partial \hat{\lambda}} \tag{A.7}$$

for $s = 1$. Since, under the agent's information set, the state switches are a point-process with instantaneous intensities $\hat{\eta}$ and $\hat{\lambda}$, these quantities also replace their full information counterparts, η and λ, in multiplying the jump terms in the respective equations.

The next steps in the derivation involving the first order condition for optimal consumption are unchanged from the full-information case. Replace $\mathbb{J}$ by the conjecture $\frac{q^{1-\gamma}}{1-\gamma} H(\hat{\eta}, \hat{\lambda}, M)$, then a common power of q term is cancelled, and the whole equation is divided by H. These manipulations lead to the above two terms becoming the right-most terms in (18) and (19), which are otherwise identical to the full-information system (11) and (12).

□

A.3 Pricing Kernel, Riskless Rate and Proposition 3

This section first derives the pricing kernel and riskless rate under partial information. The results are then used to prove Proposition 3 Section 5.2 which describes the pricing equation of insurance against a disaster.

Under stochastic differential utility, the kernel can be represented as

$$\Lambda_t = e^{\int_0^t f_{\mathbf{J}} du} f_C \tag{A.8}$$

where the aggregator function is given in (7). With the form of the value function and the optimal consumption rule from Section 4.2, evaluating the partial derivatives yields (after some rearrangement)

$$\Lambda_t = q^{-\gamma}\ H(\hat{\eta},\hat{\lambda},M)\ e^{\int_0^t [c_u\ (\theta-1)-\rho\theta] du} \tag{A.9}$$

where $c = c(\hat{\eta},\hat{\lambda},M) \equiv C/q$ is the marginal propensity to consume.

The riskless rate is minus the expected rate of change of $d\Lambda_t/\Lambda_t$ under the agents' information set. Applying Itô's lemma, for even values of M, the expected change is

$$c\ (\theta-1) - \rho\theta\ - \gamma(\mu - c) + \frac{1}{2}\gamma(\gamma+1)\sigma^2$$

$$-\frac{(\hat{\eta})^2}{a^{\eta}}\ \frac{1}{H}\frac{\partial H}{\partial \hat{\eta}}\ + \hat{\eta}\left(\frac{H(M+1)}{H(M)} - 1\right).$$

A key simplification is to observe that, by the HJB equation derived above (see (18)), the latter two terms in this expression can be replaced by $g_0 - \frac{\theta}{\psi}c$. This causes all of the terms involving c to exactly cancel. Using the definition of g_0 in (10), the remaining terms are just $-\mu + \gamma\sigma^2$. Hence we have shown

$$r_0 = \mu - \gamma\sigma^2.$$

Repeating the above steps for odd values of M and applying the same trick yields

$$r_1 = \mu - \gamma\sigma^2 - \zeta\chi(1-\chi)^{-\gamma}.$$

Turning to the insurance claim, the asset is assumed to make a terminal payout of 1.0 upon the occurrence of the next disaster. Proposition 3 characterizes its price in normal-times prior to that disaster.

Proof. We conjecture that the price, P, of the insurance is not a function of wealth, q. Moreover, when $s = 0$, the state variables a^η, a^λ, and $\hat{\lambda}$ are all fixed, and $\hat{\eta}$ evolves deterministically according to (16).

By the definition of the pricing kernel, for any claim in the economy, its instantaneous payout per unit time (in this case, zero) times Λ must equal minus the expected change of the product process $P\Lambda$, or

$$\mathcal{L}(\Lambda(q_t, s_t, \hat{\eta}_t)\ P(s_t, \hat{\eta}_t))/\Lambda_t = 0, \tag{A.10}$$

where $\mathcal{L}(X)$ is the drift operator $\mathrm{E}[dX]/dt$ under the agents' information set.

Using Itô's lemma for jumping processes to expand the expected change,

$$-\frac{(\hat{\eta})^2}{a^\eta}\frac{\partial P}{\partial \hat{\eta}} + \mu_\Lambda P + \hat{\eta}\left(\frac{H(M+1)}{H(M)} - P\right) = 0$$

where we have written μ_Λ for the deterministic terms in $d\Lambda_t/\Lambda_t$ and used the fact that $P(M+1) = 1$.

Next, add and subtract $\hat{\eta}(\frac{H(M+1)}{H(M)} - 1)P$ and use the fact that the expected growth rate of the pricing kernel is minus the riskless rate:

$$r_0 = -\mu_\Lambda - \hat{\eta}\left(\frac{H(M+1)}{H(M)} - 1\right)$$

to get (20):

$$-\frac{(\hat{\eta})^2}{a^\eta}\frac{\partial P}{\partial \hat{\eta}} - r_0 P + \hat{\eta}\frac{H(M+1)}{H(M)}(1-P) = 0.$$

□