

Disasters with Unobservable Duration and Frequency: A Quantitative Equilibrium Analysis*

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

December 2025

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. Specifically, in a disaster, uncertainty about duration makes disasters subjectively much worse and can make the welfare value of information extremely high. Resolving uncertainty about the disaster can be as or more valuable as ending the disaster itself.

JEL Codes: D6, D8, E21, E32, G10

Keywords : parameter uncertainty, welfare costs, disasters, value of information

*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.

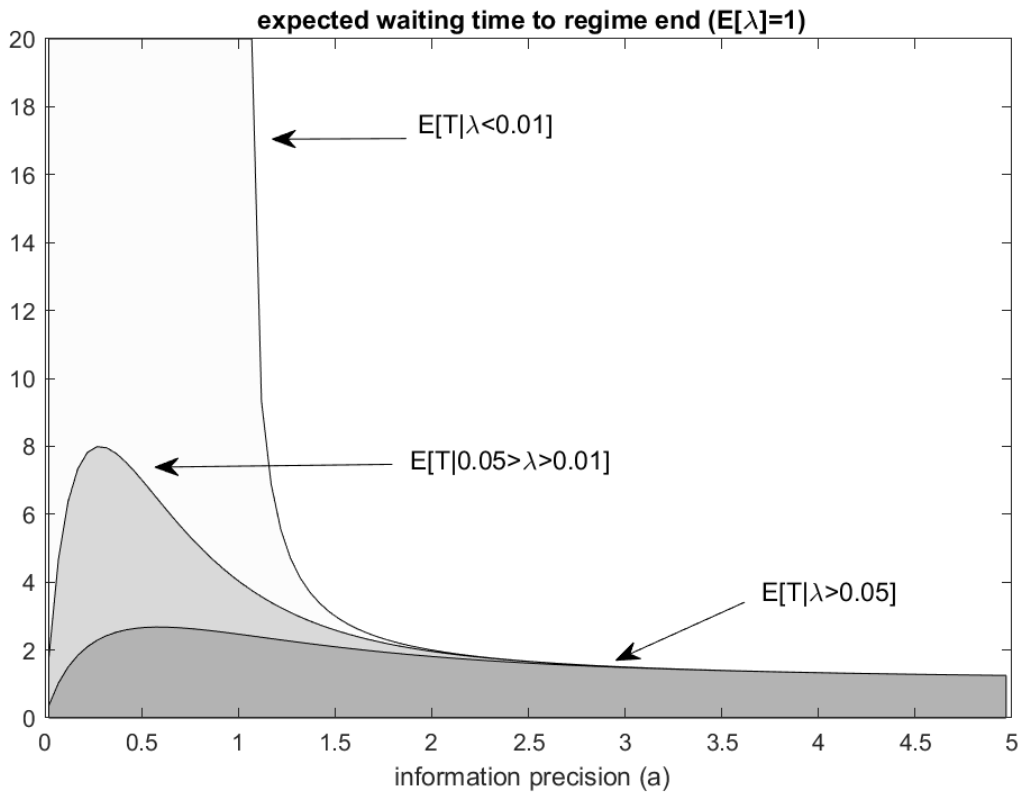
Our model depicts disasters as regimes in which the stock of wealth (potentially including human wealth) is subject to exogenous destruction. The economy transitions stochastically between these episodes and “normal times.” Agents update their beliefs by observing the frequencies of transitions, and 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



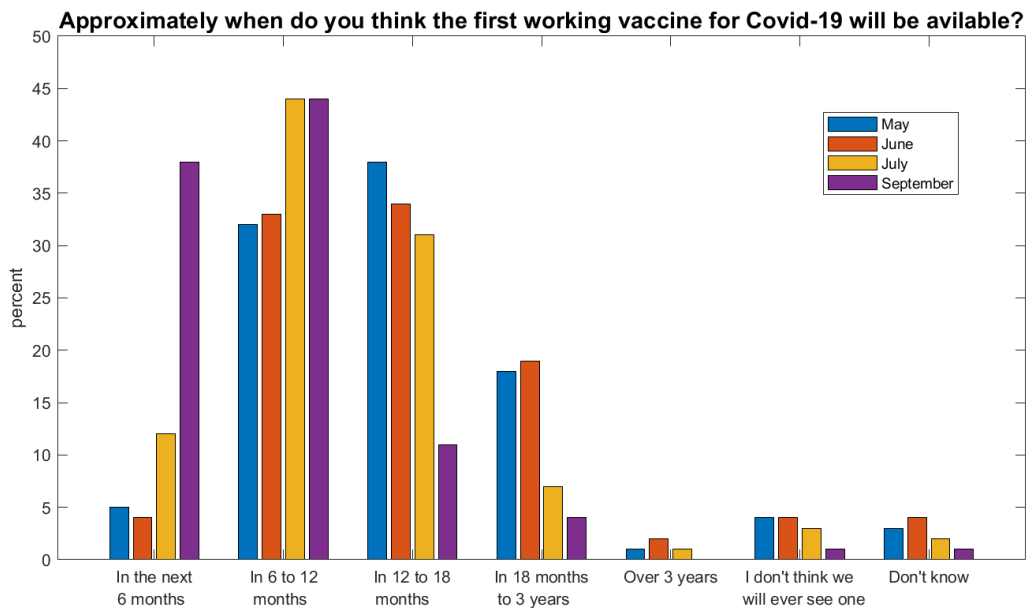
The top line plots the expected waiting time in years for the end of a regime when observers have belief about the intensity per unit time of a switch, λ , that are described by a gamma distribution with mean $E[\lambda] = 1$ and variance $1/a$, where a is the variable on the horizontal axis. The lower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

¹These 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

Next, 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.

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

²In 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.

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.

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

$$dq = \mu(s)qdt - Cdt + \sigma(s)qdB_t - \chi(s)q dJ_t \quad (1)$$

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

³Also, 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.

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

$$\mathbb{E}[\eta] = \frac{a^\eta}{b^\eta}, \quad \text{Std}[\eta] = \frac{\sqrt{a^\eta}}{b^\eta}, \quad (2)$$

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

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

and aggregator

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

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

⁴As 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.

⁵We 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:

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

⁶The 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*

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

and the value function of the representative agent is

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

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

3 Results: Information Precision in a Disaster

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.

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 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 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

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

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.

3.2 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 without curtailing the current disaster. 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, 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.

Comparing the results in 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,

⁸The 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.

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

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

3.3 Incentives to Mitigate and Acquire Information

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 Appendix (A.3.1), 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.

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. To show this, we can again consider augmenting the economy to have a one-time opportunity to purchase a signal about λ , (e.g., via a “laboratory experiment”) in the form of a realization of N transition times of the underlying process for a cost $c(N)$. Note that the outcome will also convey information about the level of λ , which may entail good or bad news. In particular, the worst-case scenario of a very low λ could be revealed,

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

¹¹In 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.

making the purchase of the signal subject to the same risks that are reflected in the partial information value function. It is thus not *a priori* clear that this information-production opportunity will have positive expected value.

In the Appendix (A.3.2), we show that, under a simple linear information cost structure, agents do optimally choose to spend significant fractions of their endowment to buy signals. The expected value from the signal increases rapidly for small increases in precision. 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.

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. From a policy perspective, reducing uncertainty about the evolution of disasters may be an important mechanism for alleviating their perceived harm.

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.
- 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.
- 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.
- 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					Partial Information: $a_0^\lambda = 10,000$				
	$\hat{\lambda}$					$\hat{\lambda}$			
	0.2	0.5	1.0			0.2	0.5	1.0	
$2^*\hat{\eta}$	0.01	0.1789	0.0873	0.0451	$2^*\hat{\eta}$	0.01	0.1791	0.0874	0.0451
	0.05	0.1414	0.0788	0.0431		0.05	0.1434	0.0790	0.0431
Partial Information: $a_0^\lambda = 10$					Partial Information: $a_0^\lambda = 1$				
	$\hat{\lambda}$					$\hat{\lambda}$			
	0.2	0.5	1.0			0.2	0.5	1.0	
$2^*\hat{\eta}$	0.01	0.2194	0.1147	0.0624	$2^*\hat{\eta}$	0.01	0.3449	0.3045	0.2583
	0.05	0.1715	0.1054	0.0640		0.05	0.2655	0.2545	0.2343
(B) Varying preference parameters									
Full Information ($\gamma = 2.0, \psi = 2.0$)					Full Information ($\gamma = 3.8, \psi = 0.20$)				
	$\hat{\lambda}$					$\hat{\lambda}$			
	0.2	0.5	1.0			0.2	0.5	1.0	
$2^*\hat{\eta}$	0.01	0.1575	0.0781	0.0416	$2^*\hat{\eta}$	0.01	0.2003	0.0915	0.0459
	0.05	0.1322	0.0716	0.0398		0.05	0.1599	0.0840	0.0441
Partial Information: $a_0^\lambda = 1$ ($\gamma = 2.0, \psi = 2.0$)					Partial Information: $a_0^\lambda = 1$ ($\gamma = 3.8, \psi = 0.20$)				
	$\hat{\lambda}$					$\hat{\lambda}$			
	0.2	0.5	1.0			0.2	0.5	1.0	
$2^*\hat{\eta}$	0.01	0.3243	0.2624	0.2065	$2^*\hat{\eta}$	0.01	0.8103	0.8363	0.8505
	0.05	0.2673	0.2364	0.2021		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.

Table 2: The Value of Information During a Disaster

$a_0^\lambda = 1, \gamma = 3.8, \psi = 2.0$					$a_0^\lambda = 10, \gamma = 3.8, \psi = 2.0$				
$\hat{\lambda}$					$\hat{\lambda}$				
0.2 0.5 1.0					0.2 0.5 1.0				
$2^*\hat{\eta}$	0.01	0.1716	0.1937	0.1686	$2^*\hat{\eta}$	0.01	0.0384	0.0201	0.0121
	0.05	0.1278	0.1798	0.1776		0.05	0.0244	0.0247	0.0214
$a_0^\lambda = 1, \gamma = 2, \psi = 2.0$					$a_0^\lambda = 1, \gamma = 3.8, \psi = 0.20$				
$\hat{\lambda}$					$\hat{\lambda}$				
0.2 0.5 1.0					0.2 0.5 1.0				
$2^*\hat{\eta}$	0.01	0.1432	0.1411	0.1149	$2^*\hat{\eta}$	0.01	0.9111	0.9043	0.8888
	0.05	0.1125	0.1401	0.1298		0.05	0.8967	0.9010	0.8911

The panel 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. 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.

A Proofs and Derivations

A.1 Full Information

To prove Proposition 1, we first treat the case of full-information in which the only state variables are $s \in \{0,1\}$ and q . For ease of notation, define the following combination of preference parameters:

$$e_0 \equiv \frac{\theta}{\psi} \rho^\psi \quad \text{and} \quad e_1 \equiv -\frac{\psi}{\theta}. \quad (\text{A.1})$$

Also define $\lambda(0) = \eta, \lambda(1) = \lambda$.

Lemma *Denote*

$$g(s) \equiv \theta \rho - (1 - \gamma) \left(\mu(s) - \frac{1}{2} \gamma \sigma(s)^2 \right) - \zeta(s) \left([1 - \chi(s)]^{1-\gamma} - 1 \right) \quad (\text{A.2})$$

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

$$g_0 \equiv g(0) = e_0 (H(0))^{e_1} + \eta \left[\frac{H(1)}{H(0)} - 1 \right] \quad (\text{A.3})$$

$$g_1 \equiv g(1) = e_0 (H(1))^{e_1} + \lambda \left[\frac{H(0)}{H(1)} - 1 \right] \quad (\text{A.4})$$

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

$$C(s) = \rho^\psi (H(s))^{e_1} q, \quad (\text{A.5})$$

and the value function of the representative agent is

$$\mathbb{J}(s) \equiv \frac{H(s)q^{1-\gamma}}{1-\gamma}. \quad (\text{A.6})$$

Proof. Using the evolution of capital stock for the representative agent (1) the Hamilton-

Jacobi-Bellman (HJB) equation for each state s can be written:

$$\begin{aligned}
0 = \max_C & \left[f(C, \mathbb{J}(s)) + \mathbb{J}_q(s)(q\mu(s) - C) \right. \\
& + \frac{1}{2} \mathbb{J}_{qq}(s) q^2 \sigma(s)^2 + \zeta(s) [\mathbb{J}(s, q(1 - \chi(s))) - \mathbb{J}(s, q)] \\
& \left. + \lambda(s) [\mathbb{J}(s') - \mathbb{J}(s)] \right]
\end{aligned} \tag{A.7}$$

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

Taking the first-order condition with respect to $C(s)$ in (A.7), we obtain

$$f_c(C, \mathbb{J}(s)) - \mathbb{J}_q(s) = 0. \tag{A.8}$$

Using $f(C, \mathbb{J})$ from (4) and taking the derivative with respect to C , we obtain

$$f_c = \frac{\rho C^{-\psi-1}}{[(1 - \gamma)\mathbb{J}(s)]^{\frac{1}{\theta}-1}}. \tag{A.9}$$

Substituting the conjecture $\mathbb{J}(s)$ in equation (7) yields

$$f_c = \frac{\rho C^{-\psi-1}}{H(s)^{\frac{\gamma-\psi-1}{1-\gamma}} q^{\gamma-\psi-1}}. \tag{A.10}$$

Then, for state s , we obtain by substituting $\mathbb{J}_q(s) = H(s)q^{-\gamma}$ in (A.8), and simplifying:

$$C(s) = \frac{H(s)^{-\psi/\theta} q}{\rho^{-\psi}} \tag{A.11}$$

which agrees with (6) using the definitions of the constants in (A.1).

To verify the conjectured form of the value function, we plug it in to the HJB equation (A.7) 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.7);
2. cancel the terms in q which have the same exponent; and
3. group constant terms not involving Hs and define them to be $g(0)$ for state 0 and $g(1)$ for state 1.

The third step yields the system of recursive equations A.3, A.4. □

Regularity Conditions

The 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} . It is straightforward to show that these constants are given by

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

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 $g_1 < g_0$.

A.2 Proposition 1: 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-disaster epochs and odd numbered states are the disasters. Also, in the non-disaster states, $\hat{\lambda}$ is constant, while $\hat{\eta}$ is constant in disasters. 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 (5) 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 HJB equations under partial information are the same as (A.7) 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:

$$-\frac{(\hat{\eta})^2}{a^\eta} \frac{\partial \mathbb{J}(0)}{\partial \hat{\eta}} \tag{A.12}$$

for $s = 0$, and

$$-\frac{(\hat{\lambda})^2}{a^\lambda} \frac{\partial \mathbb{J}(1)}{\partial \hat{\lambda}} \tag{A.13}$$

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. This follows because consumption does not enter into any of the new terms involving the information variables. 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 showing up on the right hand side, in a system that is otherwise identical to the full-information system (A.3) and (A.4).

$$g_0 = e_0 H_M^{e_1} + \hat{\eta} \left(\frac{H_{M+1}}{H_M} - 1 \right) - \frac{(\hat{\eta})^2}{a^\eta H_M} \frac{\partial H_M}{\partial \hat{\eta}} \quad (\text{A.14})$$

$$g_1 = e_0 H_{M+1}^{e_1} + \hat{\lambda} \left(\frac{H_{M+2}}{H_{M+1}} - 1 \right) - \frac{(\hat{\lambda})^2}{a^\lambda H_{M+1}} \frac{\partial H_{M+1}}{\partial \hat{\lambda}} \quad (\text{A.15})$$

where the constants g_0 and g_1 are as defined in Lemma 1 above. □

A.2.1 Solution Algorithm

In the full information case, solution of the algebraic system over a grid in the $(\hat{\eta}, \hat{\lambda})$ plane is straightforward. The unknown constants $H(s)$ are 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} . The former corresponds to $\eta = 0$ and the latter to $\lambda = 0$.

For the general case, we pick a large even integer M^{max} and assume that the economy has converged to the full information solution with $s = 0$ at M^{max} and $s = 1$ at $M^{max} - 1$. Given these solutions, the HJB system for $M = M^{max} - 2$ is just a first order ODE, since the jump terms in (A.14)-(A.15) can be explicitly evaluated. For even values of M , the boundary condition at $\hat{\eta} = 0$ is again the full-information solution because the posterior standard deviation $\sqrt{a^\eta \hat{\eta}}$ is also zero. (Note that the value of $\hat{\lambda}$ is immaterial if disasters cannot arise.) Likewise, for odd values of M , the boundary condition at $\hat{\lambda} = 0$ is given by the full-information solution. Hence, the first-order ODEs can be explicitly solved in alternating directions. The procedure is then repeated for all lower values of M .

A.3 Real Options

A.3.1 Mitigation

The text in Section 3 describes endowing the model economy with a one-time real option to invest in a mitigation technology to alter a structural parameter, χ , via $\chi = g(i)$ where

I is a lump-sum investment and $i = I/q$. Since the option is a one-shot decision, the post-investment economy is identical to the original model (without the technology) and hence its value function is as derived in the main propositions.

Then, 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.

A.3.2 Information Production

The top panel of the table below presents the optimal information investment as a fraction of wealth when the economy contains a technology allowing agents to purchase a realization of N transitions of the disaster process, e.g., in a laboratory. The realization increases a^λ by N but also alters the mean $\hat{\lambda}$ depending on the (random) time-length of the realization.¹ The table assumes that the information production function is $N = 200i$, where $i = I/q$ is the lump-sum investment. The option to make this investment is a one-time occurrence at the on-set of a disaster.

The lower panel reports the welfare gain, in units of wealth, of the investment. The difference between the respective panels can be interpreted as the value-added of the technology.

¹The time-length is a virtual output. The experiment is assumed to be atemporal. Agents receive the results immediately.

Table A.1: Information Production

(A) Optimal Investment									
Benchmark					$\psi = 0.20$				
$\hat{\lambda}$					$\hat{\lambda}$				
0.2 0.5 1.0					0.2 0.5 1.0				
$2^*\hat{\eta}$	0.01	0.015	0.015	0.015	$2^*\hat{\eta}$	0.01	0.030	0.035	0.040
	0.05	0.015	0.015	0.015		0.05	0.035	0.030	0.030
$\gamma = 2$					$\rho = 0.02$				
$\hat{\lambda}$					$\hat{\lambda}$				
0.2 0.5 1.0					0.2 0.5 1.0				
$2^*\hat{\eta}$	0.01	0.015	0.015	0.010	$2^*\hat{\eta}$	0.01	0.030	0.035	0.035
	0.05	0.010	0.015	0.015		0.05	0.025	0.030	0.03015

(B) Welfare Gain									
Benchmark					$\psi = 0.20$				
$\hat{\lambda}$					$\hat{\lambda}$				
0.2 0.5 1.0					0.2 0.5 1.0				
$2^*\hat{\eta}$	0.01	0.038	0.046	0.047	$2^*\hat{\eta}$	0.01	0.314	0.346	0.343
	0.05	0.032	0.042	0.042		0.05	0.284	0.320	0.343
$\gamma = 2$					$\rho = 0.02$				
$\hat{\lambda}$					$\hat{\lambda}$				
0.2 0.5 1.0					0.2 0.5 1.0				
$2^*\hat{\eta}$	0.01	0.031	0.032	0.020	$2^*\hat{\eta}$	0.01	0.114	0.164	0.188
	0.05	0.021	0.030	0.026		0.05	0.085	0.130	0.143