Disasters Everywhere: The Costs of Business Cycles Reconsidered*

Òscar Jordà[†] Moritz Schularick[‡] Alan M. Taylor[§]

June 2022

Abstract

This paper introduces a new test to show that business cycles are in general asymmetric, and in particular are characterized by recession events that resemble "mini-disasters." It is well known that stochastic growth is fat-tailed and non-Gaussian, but we present evidence that this is true not only in the widely-studied rare disaster events. Using long-run historical data, we show empirically that this holds for advanced economies since 1870. Focusing on peacetime eras, we develop a tractable local projection framework to estimate consumption processes in normal and financial-crisis recessions. Introducing random coefficient local projections (RCLP) we get an easy and transparent mapping from estimates to a calibrated simulation model of disasters with variable severity. Our simulations show that substantial welfare costs arise not just from the large rare disasters, but also from the smaller but more frequent mini-disasters. On average, even with low risk aversion, households would be willing to pay 12 percent of deterministic consumption to avoid these fat-tailed cyclical fluctuations.

Keywords: macroeconomic fluctuations, asymmetry, skew, fait-tails, hysteresis, local projections, random coefficients, macroprudential policy.

JEL classification codes: E13, E21, E22, E32.

*This work is part of a larger project kindly supported by research grants from the European Research Council (ERC-2017, 772332, SafeHouse) and the Institute for New Economic Thinking (#INO1600004), and we are grateful for this support. Schularick also acknowledges support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1 – 390838866. We are particularly thankful to Johannes Matschke for outstanding research assistance. For helpful comments we thank José Ursúa and seminar participants at UC Berkeley. All errors are our own. The views expressed herein are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco, the Federal Reserve Bank of New York, or the Board of Governors of the Federal Reserve System.

[†]Federal Reserve Bank of San Francisco; and Department of Economics, University of California, Davis (oscar.jorda@sf.frb.org; ojorda@ucdavis.edu).

[‡]Sciences Po Paris; Federal Reserve Bank of New York; Department of Economics, University of Bonn; and CEPR (moritz.schularick@ny.frb.org; moritz.schularick@uni-bonn.de).

[§]Department of Economics and Graduate School of Management, University of California, Davis; NBER; and CEPR (amtaylor@ucdavis.edu).

1. Introduction

A great deal of research in macroeconomics is aimed at understanding the business cycle and thinking about policy responses to dampen cyclical fluctuations. In standard models, agents are assumed to prefer smooth consumption streams and would therefore be willing to insure against fluctuations. But how much they are willing to pay for such insurance depends on how undesirable the otherwise non-smooth consumption streams would be.

Seminal works by Lucas (1987, 2003) showed that assuming transitory i.i.d. Gaussian consumption deviations from the postwar U.S. trend under log utility resulted in very small welfare losses: less than 1/10 of a percent for the representative agent. Even allowing for more volatile consumption, like the pre-war U.S. or as in other countries, or allowing for stochastic growth (Obstfeld, 1994), it is hard to get the costs of fluctuations (and hence potential gains from managing aggregate demand through countercyclical macroeconomic policy) to exceed 1–2 percent. If costs are really this low, there appears to be little upside from stabilization policies.

However, a new take on this welfare conundrum emerges as a spinoff from research on asset pricing anomalies in settings where the endowment process is fat-tailed. In short, if the welfare cost of fluctuations is low, then why is the related risk premium—the compensation for holding risky claims—so puzzlingly high? The "rare disasters" approach resolves this tension via very large but infrequent losses (Rietz, 1988; Barro, 2006; Gabaix, 2008; Wachter, 2013). In a rare disaster setting the welfare costs of fluctuations can be high. Barro (2009) finds welfare costs of 17 percent due to disaster-driven jumps, ten times larger than the 1.6 percent attributable to purely Gaussian growth components.[1] The implied costs can be larger still if utility is recursive, or if disasters have stochastic probability, stochastic size, more persistence, or are permanent rather than transitory (Obstfeld, 1994; Reis, 2009; Gourio, 2012; Nakamura, Steinsson, Barro, and Ursúa, 2013; Barro and Jin, 2011). The key insight is that amplification of both the welfare costs and risk premia derives from the impacts of higher moments of the endowment process (Martin, 2008).[2]

In the purely *rare* disaster perspective, the welfare cost of economic fluctuations mainly stems from such infrequent, but very costly events. But it is possible that the value of stabilization policy is even bigger if disaster dynamics are more pervasive, and in this paper we argue that this could be so and that the rare disaster approach needs amendment: it is not that it is incorrect, but rather that it does not go far enough.

[1]For comparability with the baseline welfare functions we adopt below, this 17% welfare cost is for the CRRA case with $\gamma = \theta = 4$ as reported in Table 3 in Barro (2009).

[2]Of course, in models with heterogeneous agents other sources of risk, notably uninsured idiosyncratic income risk, can also amplify the costs of cycles (Atkeson and Phelan, 1994; Imrohoroğlu, 1989).

In fact, when we use long-run historical data to look at growth at risk in recessions, we find that over the full sample even the typical "normal" business cycle contains a fat-tail that materially amplifies welfare costs. Deviations from the Gaussian benchmark do not just appear in the extreme form of rare disasters—the wars, pandemics, revolutions, financial crises, etc., considered by previous research. Instead, we show that fat tails with strong persistence also appear in peacetime and even in normal recessions that are not coincident with a financial crisis. In short, there are disasters everywhere.[3]

Once we embrace this idea, welfare cost judgments change. In our estimates, even with low risk aversion ($\gamma = 4$) and simple non-recursive power utility, households would be willing to pay 12 percent of deterministic consumption to avoid the peacetime consumption fluctuations seen in advanced economies since 1870. Business cycle volatility then becomes a first-order issue and successful stabilization policy could deliver sizeable welfare gains. Our further examination of the normal versus financial crisis recession dichotomy also speaks to the increased interest in potential gains from macroprudential policies to mitigate crisis risk. Here direct policy actions are being debated and even implemented as we write. Of course, in contrast, some other disastrous recessions like wars or pandemics may be less susceptible to purely economic policy interventions, and not all risks can be mitigated by economic stabilization policy.

Outline The first part of our paper documents the new stylized fact using a comprehensive macro-historical database (Jordà, Schularick, and Taylor, 2017) covering 18 advanced economies since 1870. Using local projections (LP) methods (Jordà, 2005), we present a new *test* for the presence of disasters. Instead of assuming disasters are present, or relying on skewness diagnostics, our test has the virtue of directly mapping into a tractable counterfactual simulation model discussed below. Our results apply to peacetime advanced economies, a sample considered exempt from the more frequent dislocations seen in emerging markets or wartime eras.

Our finding ties into recent research on the importance of skewness for macro and finance puzzles (Colacito, Ghysels, Meng, and Siwasarit, 2016; Dew-Becker, Tahbaz-Salehi, and Vedolin, 2019), and the micro skewness underpinnings at the firm or household level (Busch, Domeij, Guvenen, and Madera, 2018; Salgado, Guvenen, and Bloom, 2019). In this research, skewness is a general phenomenon at all times, not just in disaster episodes.[4] Our analysis also meshes with macro frameworks that embrace asymmetry and hysteresis, from the older Friedman "plucking" model to recent DSGE models with nominal wage rigidity

[3] Deviations from Gaussianity are in line with venerable arguments for asymmetric business cycle dynamics. See, e.g., Keynes (1936); Neftçi (1984); Sichel (1993); Acemoglu and Scott (1994); Morley and Piger (2012).

[4] The large literature on time-varying volatility also points to the importance of higher moments in macro.

(deLong and Summers, 1988; Kim and Nelson, 1999; Dupraz, Nakamura, and Steinsson, 2019; Calvo, Coricelli, and Ottonello, 2013; Schmitt-Grohé and Uribe, 2013). The work of Fatás and Mihov (2013) also echoes this idea, noting that U.S. postwar growth volatility fell in the Great Moderation, but negative skewness increased. But, to our knowledge, we are the first to re-assay the debate over the costs of business cycles in an extended disaster framework with ubiquitous fat tails.

The second part of the paper proposes a new empirical framework for estimating and calibrating a growth process consistent with the above findings. LP estimation has been successfully employed in fixed-coefficient form to document the systematic, large, and persistent differences in paths in normal and financial crisis recessions (Jordà, Schularick, and Taylor, 2013). Here we show how to extend the approach using random-coefficient local projections (RCLP) as a natural way to model variable-severity disasters, which is known to be an essential feature in the data and is a key driver of welfare consequences.[5] We then take the RCLP-estimated growth process and simulate an economy under various parameter configurations to assess counterfactual welfare losses due to peacetime business cycles in the standard way.[6]

As in Barro (2009), we ask how much welfare loss relative to the deterministic baseline is due to Gaussian terms versus disaster terms. Results still depend on the permanent component of the disasters. But since disasters are now everywhere, including in normal recessions, we find that welfare costs are much larger. In a peacetime setting, the Gaussian terms account for only about a 1 percent loss (cf. Obstfeld, 1994); allowing fat tails with hypothetical 100% normal recessions would increase this loss to 9 percent; and with hypothetical 100% financial crisis recessions would increase it to about 20 percent.

Main findings Summing up, we make two main contributions in this paper. First, we present a new empirical methodology for estimation and simulation built around the attractive technique of local projections. It is particularly suited to the problem of measuring disaster losses over multi-period horizons in an parsimonious way without the complexity and fragility of more elaborate methods. In random coefficient form, local projections are well equipped to model disaster gaps with stochastic scaling and persistence in a tractable and flexible fashion. The methods make for an easy and transparent mapping from the LP estimates to the calibrated simulation model.

[5]One could also explore what happens in this setting when the recession-type probability has a conditional mean which depends on covariates; a natural case to consider is when financial crisis probability depends on the history of credit growth (Schularick and Taylor, 2012). We leave this extension for future research.

[6]Economics has little to say about how to stop wars. But other events classified as disasters outside of wars are still very damaging. Here, as is well known, the most damaging type are financial crisis recessions (see, e.g., Muir, 2017). Normal recessions are rarely very disastrous, though probabilistically some will be so.

Second, we offer a new perspective on the important macro question of the cost of business cycles and the importance of stabilization policy. The main new insight is that disasters are possibly everywhere—in the sense that the growth process has fat tails not just in the large Barro-type rare disasters. And although these recessions are smaller in amplitude than the rare big disasters, they are more frequent, and the welfare costs add up to something substantial. The size of the consumption-equivalent welfare loss that we find is 12% over the full historical sample back to 1870. Looking at the crisis-prone post-1985 era in advanced economies, our model says agents would sacrifice about 15 percent of consumption to avoid all business cycles; and about 7 percent just to avoid financial-crisis recessions. This goal was attainable in the 1950s–1960s era when financial crises were absent. This result also speaks to the large potential gains from making our financial systems less crisis-prone. In short, even outside of times of war and pandemics, there is considerable gain to smoothing economic fluctuations, and all the more so in economies subject to financial instability.

2. Disasters everywhere? testing for the presence of disasters

A large literature has documented the evidence against the null assumption of a Gaussian random-walk growth process. In this section we begin with a review of the stylized facts using the most recent data. We then turn to a key question: if the growth process is *not* Gaussian, then which among all the many alternative candidate models should we investigate? At that point we follow the standard literature on disasters and model growth as a jump process with some basic and widely-used features in that jumps are negative and they may exhibit persistence, alhtough in addition they might be quite frequent.

Skewness and kurtosis tests Many researchers in the disasters literature have noted the evidence against the Gaussian growth assumption in long-run cross-country panel data. For sure, the Gaussian assumption can be difficult to reject in a single advanced country, such as in the United States, in so-called New World Offshoots like Australia and Canada, or in neutral countries like Sweden and Switzerland. These fortunate economies largely avoided the huge adverse shocks associated with world wars in the typical post-1870 panel. But the baseline assumption for us, as in the literature, is that the history of any advanced country represents a draw from the underlying distribution of the growth process, so that we must pool all histories to avoid the "peso problem" of drawing only from the urns associated with tranquil development paths.

For example, Barro (2006, Table III) and Ursúa (2011, Table 1.3) report standard tests of skewness and kurtosis tests for log real consumption per capita growth, and we show similarly findings here in Table 1, using the JST Macrohistory annual panel dataset which covers a sample of 17 advanced economies since 1870 (http://www.macrohistory.net/data/). Note that our sample now includes the 2008 crisis and its aftermath, which only serves to make rejection of the Gaussian null even stronger. In the full sample skewness is -0.87 (the null is 0) and kurtosis is 17.25 (the null is 3). Confidence intervals are computed with an 8-year-span panel block bootstrap as in Ursúa (2011, Table 1.3), and the values are significantly different from the null. Rejection is also strong even when wartime periods are dropped, showing that non-Gaussian growth is not driven just by those few major disaster events. Thus, the historical data suggest a non-Gaussian model is needed to match reality.

Variance ratio tests Evidence against the null also emerges from earlier research on the persistence of shocks using long-run variance-ratio tests. Under the null of an i.i.d. random walk, the ratio VR of $1/k$ times the k-period variance of growth to the 1-year variance should be exactly 1. For example, with this test Cochrane (1988) argued for stable long-run trend-reversion dynamics in the United States for 1869 to 1986. However, for a wider OECD sample Cogley (1990) found that not to be the case for other advanced economies. Expanding to the broadest sample yet, with emerging and developing economies up to 2009, Ursúa (2011, Table 1.43) showed that the U.S. pattern of trend-reversion was the exception not the rule.

Evidence of the same sort for our sample using the JST Macrohistory dataset is given in Table 2. Again, for some of the fortunate countries mentioned above, the null is not rejected, as in Cochrane (1988). However, using the standard Lo-Mackinlay test, we reject the null even for the U.S. using spans of 20 and 30 years. The null fails to be rejected for Canada, for 3 out of 4 Scandinavian countries, and for Switzerland. In all other cases, the null is rejected at the 30 year span, and often at 10 and 20 also. Thus, the historical data suggest a non-random-walk model is needed to match reality.

The next question is what alternative model we might consider. To persuade the reader that our choice of a disaster framework is reasonable, and that this jump-based model can be justifiably extended to a wider range of more frequent events, we first need to develop convincing tools for inference in a nonstandard setting in order to make our case.

Modeling and testing for disasters In the wide range of standard linear models, like the workhorse neoclassical stochastic growth model, where perturbations are symmetrically distributed, bad outcomes are temporary. A quick reversion to the mean is a built-in feature, and although shocks may be persistent they are not skewed. In contrast, disaster models eschew symmetry—but then, one might ask: is there any asymmetry? and, if so, of what form?

Table 1: *Skewness and kurtosis tests for log real consumption per capita growth*

Skewness and kurtosis computed for log real consumption per capita growth based on the JST Macrohistory dataset which covers a sample of 17 advanced economies since 1870 (http://www.macrohistory.net/data/). The non-war sample excludes 1914–21, 1939–49, and the years 1934–38 for Portugal and 1935–36 for Spain. Confidence intervals based on 8-year-span panel block bootstrap.

Sample	Skewness	90% CI	Kurtosis	90% CI
Full sample	-0.87	$[-1.50, -0.10]$	17.25	$[12.69, 21.55]$
Excluding wars	-0.31	$[-0.82, -0.11]$	7.88	$[6.77, 10.01]$

Of course, it might be an obvious conclusion when studying a small number of large rare disaster events with heavily skewed outcomes (e.g., Barro, 2006) that the draws have to be non-Gaussian at these moments. Peak to trough declines ranging from 20% to 60%, as seen in wars or the Great Depression, would be 5σ to 10σ events using back-of-the-envelope math, so a formal test is almost superfluous, and we know what it will say. The world is surely not Gaussian at those moments and that debate is largely settled.

However, we propose to broaden the definition of a "disaster" event. If we seek to extend this way of modeling skewness to a wider range of episodes—which is the goal of this paper—then the rejection of the Gaussian null in this larger class of putative mini-disasters isn't so obvious, and having a formal hypothesis testing structure is an essential piece of the argument to convince the reader. Below we show that even if one excludes well documented disasters from the sample, our tests indicate that economic fluctuations have all the hallmarks of traditional disasters, albeit at a smaller scale but occurring at a higher frequency.

There are two main difficulties in formally testing for disasters. The first difficulty is that definitions of a disaster are self-referential and lead to sample selection bias. Recessions, for example, refer to periods of negative growth after a peak. As we will show, even if the null model is correct, the distribution under the null is shifted by an amount equivalent to the conditional mean given that growth is negative in disaster states, which is different to the unconditional mean under the null. Alternative definitions of disasters (sometimes involving multiple periods) truncate the space of possible outcomes in other ways. The null distribution in these more general cases cannot be easily derived analytically and requires simulation methods—the bootstrap, in particular—as we shall see.

The second difficulty is the specification of the alternative model that describes the disaster asymmetry. What type of dynamic pattern characterizes a financial crisis, for example? Specifying a model, as is usually done in this literature, simplifies the analysis to a few parameters. However, it is less clear how well such models describe the dynamics

Table 2: *Variance ratio tests for log real consumption per capita growth*

Variance ratio tests computed for log real consumption per capita growth based on the JST Macrohistory dataset which covers a sample of 17 advanced economies since 1870 (`http://www.macrohistory.net/data/`). Full sample is used. The Lo-Mackinlay test is reported for spans of 10, 20, and 30 years. VR is variance ratio, R_s is the standardized test statistics which in normally distributed, and p is the significance level.

	span	N	VR	R_s	$p > \|z\|$
AUS	10	117	1.268	0.7021	0.4826
	20	117	2.080	1.8683	0.0617
	30	117	4.062	4.2549	0.0000
BEL	10	74	2.809	3.0690	0.0021
	20	74	5.933	7.3986	0.0000
	30	74	10.733	13.8024	0.0000
CAN	10	116	1.051	0.1109	0.9117
	20	116	1.119	0.1820	0.8556
	30	116	1.770	0.9774	0.3284
CHE	10	117	0.910	−0.2475	0.8045
	20	117	1.210	0.3951	0.6927
	30	117	1.550	0.8461	0.3975
DEU	10	117	2.598	3.2692	0.0011
	20	117	3.417	3.6254	0.0003
	30	117	5.040	5.2321	0.0000
DNK	10	117	0.495	−1.0537	0.2920
	20	117	0.488	−0.8228	0.4106
	30	117	0.634	−0.5089	0.6108
ESP	10	117	1.176	0.4375	0.6618
	20	117	1.747	1.2746	0.2025
	30	117	3.030	2.8724	0.0041
FIN	10	117	1.046	0.0972	0.9226
	20	117	1.488	0.8186	0.4130
	30	117	2.469	2.1061	0.0352
FRA	10	117	1.917	1.3881	0.1651
	20	117	2.349	1.6758	0.0938
	30	117	3.700	3.0682	0.0022
GBR	10	117	1.717	1.6991	0.0893
	20	117	2.866	3.5267	0.0004
	30	117	5.329	7.1691	0.0000
ITA	10	117	3.297	4.4783	0.0000
	20	117	6.580	8.5118	0.0000
	30	117	12.087	14.7778	0.0000
JPN	10	113	2.967	2.7223	0.0065
	20	113	5.098	4.7852	0.0000
	30	113	7.524	7.0151	0.0000
NLD	10	117	0.959	−0.0715	0.9430
	20	117	1.052	0.0767	0.9388
	30	117	1.628	0.8537	0.3932
NOR	10	117	0.647	−0.5670	0.5707
	20	117	0.980	−0.0261	0.9792
	30	117	2.126	1.3776	0.1683
PRT	10	77	3.007	5.1099	0.0000
	20	77	6.030	9.2170	0.0000
	30	77	11.585	16.2933	0.0000
SWE	10	117	0.596	−0.8923	0.3722
	20	117	0.659	−0.5694	0.5691
	30	117	0.932	−0.0976	0.9223
USA	10	117	1.331	0.7821	0.4342
	20	117	2.247	2.0770	0.0378
	30	117	4.072	4.2258	0.0000

seen in the data. By necessity, they must be quite parsimonious to have tractable likelihood functions. But we know reality is often more complicated. Instead, we take a semiparametric flexible approach that preserves tractability while adapting to the dictates of the data. The local projections (Jordà, 2005) toolkit turns out to work quite well for this purpose.

In our application we will examine measures of economic activity generally used in studies of rare disasters and drawn from very long run historical annual panel data (Barro, 2006; Barro and Ursúa, 2008b). We generically label the variable as X, which may refer to either real GDP per capita, Y, or real consumption per capita C. Let $x = \log(X)$, hence, lowercase variables denote logs. We will write the growth rate as $\Delta x_t = x_t - x_{t-1}$. Our focus for now in the main text will be on consumption, as this is the key measure in the rare disaster literature, and will matter for the second part of this paper as we consider welfare implications.

Gaussian null model What does it mean to choose a null Gaussian model? As noted, one canonical model is a deterministic trend with Gaussian shocks (Lucas, 1987, 2003); today, the widely-used baseline is a random walk with drift (Obstfeld, 1994). Working in the rare disasters tradition (Barro, 2006; Barro and Ursúa, 2008b), we follow the latter, where the null characterizes growth as a random walk with drift, with a data generating process

$$\mathcal{M}_0: \quad \Delta x_{t+1} = \mu + \epsilon_{t+1}\,; \qquad \epsilon_{t+1} \sim \mathcal{N}(0, \sigma^2)\,. \tag{1}$$

If adopting this model as the data generating process, the researcher has to find a way to specify only two parameters $\{\mu, \sigma\}$. The notation $\mathcal{M}_0$ denotes "null model." The economic implications of this model are simple. Absent shocks, the economy will grow at a constant μ rate. Shocks will make the economy grow faster or slower but because they are i.i.d. draws, they have no lasting impact on the growth trajectory (though, of course, have a permanent effect on the level). Intuitively, there will likely be few gains from stabilization policy unless the variance of the residuals is quite large.

The mapping of disasters What is a disaster event? There is no single agreed definition. Generally speaking, we have a rule which specifies how the observer maps historical data into a set of events. An example of such a rule is the (Barro, 2006) definition of a rare disaster which we specify here as the year of a local peak, followed by a peak to trough consumption decline exceeding 15% in log terms (i.e., a cumulative change of -0.15 log points). Other observers may choose different rules: for example, the widest definition would be when the events are all consumption recessions, i.e., any consumption peak year.

In general, suppose the data seen by the observer are the historical realizations of the process $\Delta \boldsymbol{x} = \{\Delta x_1, \ldots, \Delta x_T\}$. Based on these realizations, the observer generates binary indicators that capture disasters based on a *mapping*, denoted $m_t(\Delta \boldsymbol{x})$, such that: $d_t = I(m_t(\Delta \boldsymbol{x}))$, where $I(.)$ is the indicator function that takes the value of 1, when at time t there is a disaster according to $m_t(\Delta \boldsymbol{x})$, and is 0 otherwise. Thus $\{d_t\}_{t=1}^{T}$ is a *point process*.

What makes inference complicated? Importantly, we have to note that the mapping is: (1) determined by the observer; and (2) depends on the sample observed. As a result, the mapping under the null results in a truncation of the distribution of the data, generating sample selection bias. Sample selection can lead the analyst astray, as we now show with a simple example, using "negative growth" as a naïve event definition purely for illustration.

Consider the Gaussian null model in Equation 1 and the implied "negative growth" mapping $d_t = 1$ iff $\Delta x_t < 0$, the definition of a recession event. Then note that:

$$E(\Delta x_{t+1}) \;=\; E(\Delta x_{t+1} \mid d_t = 1)\, p \;+\; E(\Delta x_{t+1} \mid d_t = 0)\,(1 - p)\,,$$

with $p = P(d_t = 1) = \Phi(\eta)$ where, $\Phi(.)$ refers to the normal cumulative probability function and hence $\eta = -\mu/\sigma$ due to standardization. From the properties of the truncated Gaussian, we have

$$E(\Delta x_{t+1} \mid d_t = 1) \;=\; \mu \;-\; \sigma\,\frac{\phi(\eta)}{\Phi(\eta)}\,, \qquad E(\Delta x_{t+1} \mid d_t = 0) \;=\; \mu \;+\; \sigma\,\frac{\phi(\eta)}{1 - \Phi(\eta)}\,.$$

For simplicity and noting that $\lambda(\eta) = \sigma\phi(\eta)/(1 - \Phi(\eta))$ is the inverse Mills ratio scaled by σ, we have that

$$\begin{aligned} \mathcal{M}_0: \quad \Delta x_{t+1} \;&=\; \mu \;-\; \lambda(\eta)\frac{1-p}{p} d_t \;+\; \lambda(\eta)(1 - d_t) \;+\; \epsilon_{t+1} \\ &=\; \mu \;+\; \beta_{11}\, d_t \;+\; \beta_{01}\,(1 - d_t) \;+\; \epsilon_{t+1}\,. \end{aligned} \tag{2}$$

The notation $\beta_{11} < 0$ serves to indicate that the "downshift" parameter refers to when $d_t = 1$ and that the effect is felt in period 1. Similarly, the notation $\beta_{01} > 0$ indicates the "upshift" that happens when $d_t = 0$ and the effect is felt in period 1. Note that $\beta_{11} = \beta_{01}(1 - p)/p$ so we can focus just on the "downshifts" in applications and testing, as we do below.

Equation 2 thus serves to illustrate an important point. Even under the null model in Equation 1, estimation by partitioning the sample according to the point process $\{d_t\}$ will, in general, create bias. Estimation allowing for a potentially different mean when $d_t = 1$ than when $d_t = 0$, generates sample selection bias. A test of the null that the "downshift" (or the "upshift") are zero could be rejected, not because the null model is rejected, but simply because the conditional means differ as a function of d_t. And that sample selection bias, $\lambda(\eta)(1-p)/p$ when $d_t = 1$ and $\lambda(\eta)$ when $d_t = 0$, happens to be a function of the inverse Mills ratio, just as the Heckman (1974) correction would prescribe.

Just to put some numbers on this example, suppose $\mu = 2$ and $\sigma^2 = 4$, for simplicity (close to the actual values for the U.S. in the postwar annual data). That is, in log terms, consumption grows at 2% on average and between 6% and -2% about 95% of the time. Under these parameters, we would find that the downshift under the null is $\beta_{11} = -1.5$ with a t-ratio of approximately 3 in a sample with 100 observation where the share of consumption recession peaks is 16% (one every six years, roughly as in the data). An unsuspecting researcher might therefore conclude that the regression provides ample evidence against the random walk with drift null when in fact there is none.

Of course, outside of this simple mapping, it is usually difficult or impossible to determine the null downshift of the mean and hence the null model itself. However, the downshift can be easily determined by simulation using the bootstrap under the null.

A general alternative model In what follows, we entertain various definitions of disaster events that go well beyond (but include) the simple example from the previous section based on recession peaks. Instead of specifying a particular data generating process (DGP) that can be specified with a few additional parameters (see, e.g. Barro and Ursúa, 2008a; Barro and Jin, 2011; Nakamura et al., 2013), we follow a more general approach. The idea is to be flexible and put as few constrains as possible on the alternative model, while still being able to test the null hypothesis efficiently. We accomplish this with a semi-parametric approach.

In general, the mapping $m(\Delta \boldsymbol{x})$ generates a point process $\{d_t\}_{t=1}^{T}$ that can induce potentially persistent downshifts on mean growth when an event takes place. Given the null model in Equation 1, each mapping will generate a truncation of the null distribution, which can have effects that are observed over several periods.

Consider then, a natural extension of the null model in Equation 2, namely

$$\mathcal{M}_0 : \quad \Delta x_{t+h} \;=\; \mu \;+\; \beta_{1h}\, d_t \;+\; \beta_{0h}\,(1-d_t) \;+\; \epsilon_{t+h}\,; \quad h = 1, \ldots, H\,, \tag{3}$$

which is the usual specification of a local projection. Thus, β_{1h} captures the downshift over time of a disaster that happened h periods ago, whereas β_{0h} captures the upshift. Both are constructed under the null random-walk-with-drift model in Equation 1 in parallel fashion to how we derived the null model when the mapping refers to recession peaks in Equation 2.

The β_{0h} and β_{1h} are, in principle, infinite dimensional moving average terms. Since the growth rate Δx_{t+1} is assumed to be stationary around μ, it is natural to assume absolute summability, that is $\sum_{h=1}^{\infty} |\beta_{jh}| < \infty$ for $j = 0, 1$. Of course, in finite samples, we cannot estimate infinite parameters and hence we have to truncate the series at some horizon, say H. If we are willing to assume that H is chosen such that $H^2/T \to 0$ as $H, T \to \infty$, and furthermore $H^{1/2} \sum_{h=H+1}^{\infty} |\beta_{jh}| \to 0$ as $H, T \to \infty$, for $j = 0, 1$, then truncation will deliver consistent estimates up to horizon H (see, e.g. Kuersteiner, 2005, for a similar condition applied to justify the consistency of truncated infinite vector autoregressions).

Consider now an alternative model, denoted $\mathcal{M}_A$, with a similar structure as $\mathcal{M}_0$ in Equation 3 but with potentially different upshift and/or downshift parameters to admit the possibility of disaster-type deviations from the null, namely

$$\mathcal{M}_A: \quad \Delta x_{t+h} \;=\; \mu \;+\; b_{1h}\, d_t \;+\; b_{0h}\,(1 - d_t) \;+\; \epsilon_{t+h}; \quad h = 1, \ldots, H \tag{4}$$

where we once again assume absolute summability to preserve stationarity, and that the truncation horizon H is chosen under similar conditions as for the null model to ensure consistency.

Given this alternative, the typical hypothesis that one would be interested in testing is

$$H_0: \quad E(\Delta x_{t+h} \mid d_t = 1\,;\, \mathcal{M}_A) = E(\Delta x_{t+h} \mid d_t = 1\,;\, \mathcal{M}_0) \quad \Longrightarrow \quad H_0: \quad \beta_{1h} = b_{1h} \tag{5}$$

Local projection estimation and inference Equation 3 and Equation 4, when truncated at horizon H, fit naturally into the local projections framework. The null can be easily estimated from Equation 1 from which bootstrap samples of the specification in Equation 3 can be easily constructed to obtain the empirical distribution of β_{1h}. Similarly, one can directly estimate Equation 4.

In addition, we can formulate the null and alternative models in cumulative terms to gain more clarity. This can be easily done with the set of local projections,

$$\begin{aligned} \mathcal{M}_0: \quad x_{t+h} - x_t \;&=\; \mu_h \;+\; \sum_{j=1}^{h} \beta_{1j}\, d_t \;+\; \sum_{j=1}^{h} \beta_{0j}\,(1 - d_t) \;+\; u_{t+h} \\ &=\; \mu_h \;+\; \alpha_{1h}\, d_t \;+\; \alpha_{0h}\,(1 - d_t) \;+\; u_{t+h}\,; \quad h = 1, \ldots, H\,, \end{aligned} \tag{6}$$

with $\mu_h = \mu\, h$, and where $\alpha_1 = \beta_1$, $\alpha_h = \beta_1 + \ldots + \beta_h$. Similarly, the alternative model becomes

$$\mathcal{M}_1: \quad x_{t+h} - x_t \;=\; g_h \;+\; a_{1h}\, d_t \;+\; a_{0h}\,(1 - d_t) \;+\; v_{t+h}\,; \quad h = 1, \ldots, H\,, \tag{7}$$

and where $a_1 = b_1$, $a_h = b_1 + \ldots + b_h$.

The hypothesis of interest now becomes $H_0 : \alpha_h = a_h$ for $h = 1, \ldots, H$. Again, we note that estimation of the null model using Equation 1 and using the bootstrap to obtain the empirical distribution of the α_h is easy, while the a_h can be directly estimated by local projections.

The key aspect of inference is that $a_h - \alpha_h < 0$ gives evidence against the Gaussian null, indicating downshifts at one or more horizons that exceed what should be seen under the null. With that framing in mind, we turn to the results of these tests in long-run panel data.

3. Results from applying the tests

Using the local projection tools just developed, we can now present—given any mapping—our estimates of the parameters under the null β_h and α_h, and likewise the data-based estimates of b_h and a_h to evaluate the hypothesis of "disaster" type deviations from the Gaussian null.

The data used for our estimates are from the JST Macrohistory annual panel dataset which covers a sample of 17 advanced economies since 1870 as originally developed by Schularick and Taylor (2012); Jordà, Schularick, and Taylor (2013), and subsequently updated (`http://www.macrohistory.net/data/`). The consumption data here are an extended version of the data in Barro (2006).

We should further note here that the JST dataset excludes emerging markets, a subset which accounts for many of the peacetime rare disasters in Barro (2006) and where growth is typically judged to be more fat-tailed in general. Our sample selection is just the advanced economies, and that will tend to work against our main hypothesis, likely making it harder to reject the Gaussian null, and making this a sterner test of our our proposed alternative "disasters everywhere" hypothesis.

To recap, we now estimate for various mappings the null and alternative local projections in long-run panel data, over countries $i = 1, \ldots, N$ and years $t = 1, \ldots, T$, for the models

$$\begin{aligned}
\mathcal{M}_0 : \quad \Delta x_{i,t+h} &= \mu_i + \beta_{1h} d_{it} + \beta_{0h} (1 - d_{it}) + \epsilon_{i,t+h} \; ; \quad h = 1, \ldots, H \,, \\
\mathcal{M}_A : \quad \Delta x_{i,t+h} &= \mu_i + b_{1h} d_{it} + b_{0h} (1 - d_{it}) + \epsilon_{i,t+h} \; ; \quad h = 1, \ldots, H \,.
\end{aligned}$$

Again, the former uses the bootstrap assuming the null; the latter is estimated from the data. We present the "downshift" parameters (β_{1h}, b_{1h}) and cumulatives (α_{1h}, a_{1h}), as well as bootstrap confidence bands for the null. Note that for the panel we include country fixed effects to control for different long-run growth trends, and we estimate the local projections using OLS with inverse propensity score weights based on the country-level disaster-event probabilities $\frac{1}{p_i}, \frac{1}{1-p_i}$, to address sample selection bias at the panel level.

To facilitate a comparison of the alternative views of disasters we will examine three different mappings of consumption recession peaks $d_t = I(m_t(\Delta \boldsymbol{x}))$, as follows,

$$\begin{aligned} All: \quad d_{it} = 1 \iff (i,t) \in Peaks_{All} &= \{(i,t) \mid \Delta x_{i,t} \geq 0, \Delta x_{i,t+1} < 0\}, \\ BRD: \quad d_{it} = 1 \iff (i,t) \in Peaks_{BRD} &= \{(i,t) \mid \Delta x_{i,t} \geq 0, \Delta x_{i,t+1} < 0, PT(\Delta \boldsymbol{x})_{it} < \theta\}, \\ xBRD: \quad d_{it} = 1 \iff (i,t) \in Peaks_{xBRD} &= \{(i,t) \mid \Delta x_{i,t} \geq 0, \Delta x_{i,t+1} < 0, PT(\Delta \boldsymbol{x})_{it} \geq \theta\}, \end{aligned}$$

where $PT(\Delta \boldsymbol{x})_{it}$ denotes the percentile (in the sample) of the peak-to-trough decline after the peak denoted by $d_{it} = 1$, and where the θ cutoff is chosen to match the frequency of Barro rare disasters (BRD) in the actual data (i.e., those that incur peak-to-trough declines larger than -0.15 log units).

Note that by using a percentile cutoff we focus on the left tail. In the end, we will thus be asking whether the worst fraction of Gaussian events can match the same worst fraction of observed events, using a cutoff that captures the BRD definition.

Here, the notation "All" denotes the simple set of all consumption recession peaks, "BRD" denotes peaks that are followed by Barro rare disasters (under the percentile definition), and "xBRD" denotes the complement, the peaks that are not followed by Barro rare disasters (likewise).

To get an idea of the fairly typical parameters under the hypothetical null random walk with drift, and the frequency of standard recession peaks and BRD peaks, we can consult Table 3. Average annual drift is about 2% and volatility is about 4%, in line with expectations. A typical country has about 20 recessions in 120 peacetime years, or one every six years. The consumption rare disasters are quite infrequent outside of wars, and number only 15.

Main result: all recessions have non-Gaussian permanent losses Figure 1 presents our main result, with baseline estimates on the full sample, 1870–2016, all 17 countries. The mapping *All* is used here, so the definition of a disaster event is any recession peak. The sample is all peacetime years, with the exclusion of WW1 and WW2 (and subsequent recovery windows), and the Spanish Civil war periods that effected the economies of both Spain and Portugal. These match the war-related disasters in Barro (2006).

In the figure, the left panel displays b_h coefficient estimates from the data, with point estimates shown by a navy square. The right panel shows the corresponding cumulative coefficient estimates of a_h from the data as a navy line. Both panels also show with purple circles the corresponding coefficients under the Gaussian null, β_h and α_h, with 90%

Table 3: *Summary data: Gaussian null drift and volatility, recessions and rare disasters, full sample (ex. war)*

The table shows the country-specific drift μ_i and volatility σ in a model of random walk with drift,

$$\Delta x_{t+1} = \mu + \epsilon_{t+1}\,; \qquad \epsilon_{t+1} \sim \mathcal{N}(0, \sigma^2)$$

The variable x is real consumption per capita and the sample used is the baseline full sample, excluding wars. The table also shows the number of country-specific recession peaks and the number of peaks which satisfy the *BRD* criterion of a peak-trough log $\times 100$ change less than -15.

	μ_i	σ	No. of peaks	No. of BRD peaks	List of BRD peaks (peak-trough change, expressed as $100 \times \log C$)							
Australia	1.31	3.76	24	3	1891	(-39.3)	1896	(-17.8)	1930	(-23.7)		
Belgium	1.80	3.76	10	0								
Canada	1.88	3.76	16	2	1873	(-16.5)	1929	(-26.1)				
Denmark	1.53	3.76	20	0								
Finland	2.39	3.76	18	2	1928	(-22.2)	1989	(-15.1)				
France	1.61	3.76	24	0								
Germany	2.27	3.76	17	0								
Italy	1.56	3.76	18	0								
Japan	2.54	3.76	21	0								
Netherlands	1.50	3.76	22	0								
Norway	1.96	3.76	17	0								
Portugal	3.05	3.76	12	0								
Spain	1.92	3.76	22	3	1894	(-23.0)	1901	(-15.6)	2007	(-16.3)		
Sweden	1.99	3.76	23	0								
Switzerland	1.35	3.76	22	4	1876	(-25.5)	1881	(-15.3)	1885	(-15.2)	1887	(-17.1)
UK	1.54	3.76	20	0								
USA	1.89	3.76	20	1	1929	(-23.4)						

confidence intervals from the bootstrap (500 replications).

In the left panel Figure 1a, the growth penalty relative to the null $b_h - \beta_h = 0$ cannot be rejected at $h = 1$, meaning that the first year of a recession is not typically different from the Gaussian null. However, we can see that evidence of drag accumulates in later years, with the null rejected in most years for $h \geq 2$. This has corresponding implications for the cumulative path of deviations shown in the right panel Figure 1b. Again, the path aligns with the null at at $h = 1$, but then diverges down and away from the null thereafter. By $h = 10$ the cumulative level penalty relative to the null $a_h - \alpha_h$ becomes substantial, amounting to a highly statistically significant 4% loss in levels terms at year 10. After a recession event, growth is skewed negatively relative to the Gaussian null.

We briefly remark on the null path here. It is, of course, a dogleg down in levels at $h = 1$ and then flat. A recession under the Gaussian null must have negative growth draw at $h = 1$, by construction, so the mean is a truncated Normal; after that, with an iid process, any draw can obtain and the unconditional growth mean is zero. Using non-panel notation to avoid clutter, we have $E(\Delta x_{t+1} \mid d_t = 1) = \mu - \sigma\phi(\eta)/\Phi\eta$ as before , but $E(\Delta x_{t+h} \mid d_t = 1) = 0$ for $h > 1$, given iid draws, and the event probability would be

Figure 1: *Tests for disasters: all recessions, advanced economies 1870–2017 (ex. war)*

The figure shows the b_h and cumulative a_h coefficients, with $a_h = \sum_{j=0}^{h} b_j$ for $h = 1, \ldots, 10$, with 90% confidence intervals from the bootstrap in brackets (500 replications). The event indicator denotes all recessions in the full sample. The coefficients are also estimated separately for normal and financial crisis recessions, following the classification approach in Jordà, Schularick, and Taylor (2013). See text. Units are in log times 100.

(a) *Growth coefficients, null β_h and data b_h, all recessions, full sample*

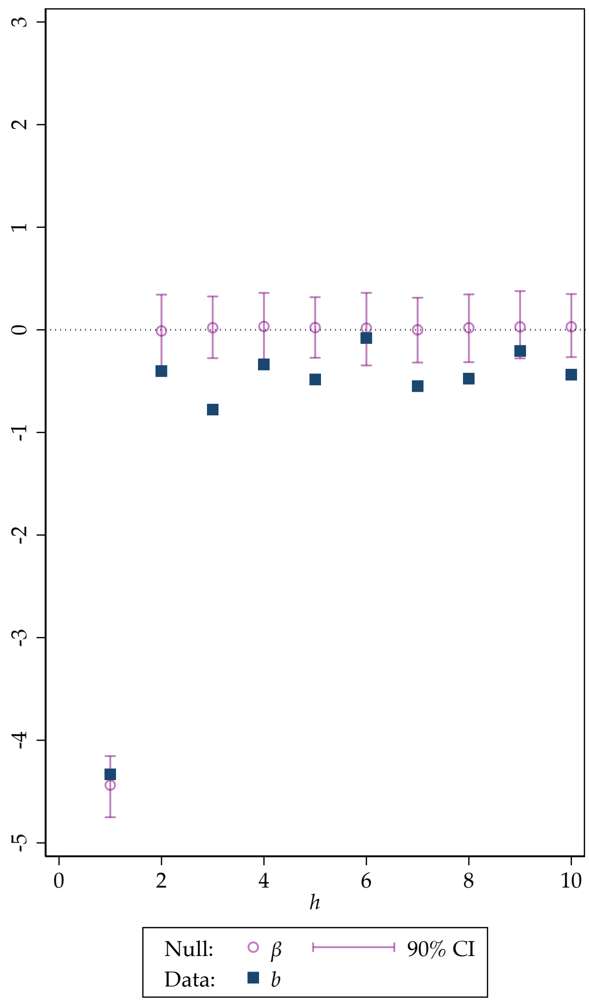

(b) *Level coefficients, null α_h and data a_h, all recessions, full sample*

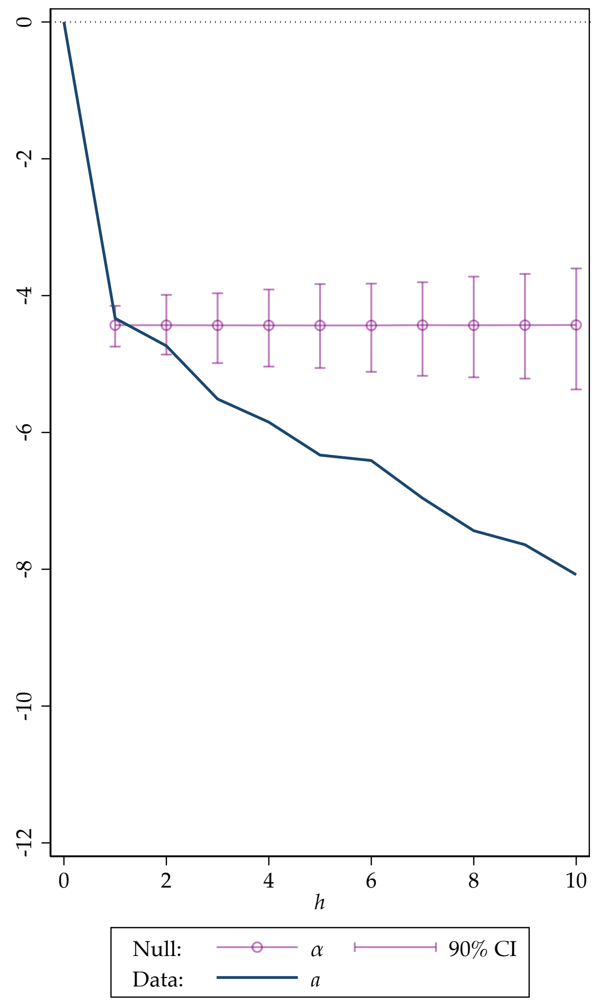

$p = P(d_t = 1) = \Phi(\eta)(1 - \Phi(\eta))$ since a peak event requires a positive growth draw followed by a negative growth draw.[7]

This is our main result. The cumulative penalty relative to the Gaussian null is found to be substantial and statistically significant for the sample of all recessions, what we term "disasters everywhere." We now probe the robustness of this result.

Robustness: the main result is not driven by just the Barro rare disaster events One objection to our main result might be that such negative skew is not surprising in a sample that includes quite a few Barro (2006)-type rare disasters. There is a straightforward way to address this concern: repeat the analysis using a subset of events that includes only the Barro (2006) rare disasters, and another subset that excludes them. This is put into effect by using the mappings *BRD* and *xBRD* defined above, where the cutoff for a rare disaster event is a peak-trough decline larger than 15 in log $\times 100$ units.

And the main result continues to hold, as reported in Figure 2. Both panels shows the cumulative coefficient estimates of a_h from the data as a navy line, and the corresponding coefficients under the Gaussian null, α_h, as purple circles with 90% confidence intervals. Note that the scales are very different in the two panels, with of course the trajectories for the BRD case deviating negatively to a much larger degree than in the xBRD case.

The left panel Figure 2a reports results where only *BRD* peaks are included, both in the null and alternative. The null is just rejected here, but rejection would be much stronger if we used the full set of disasters covering wartime events as well (it is in wars that most of the largest peak-trough declines of 40%–60% are observed). The right panel Figure 2b reports results with the *BRD* peaks excluded, both in the null and alternative. The null is strongly rejected here, even more convincingly than in Figure 2a. Even outside of rare disasters growth trajectories deviate sharply from the Gaussian null after a recession. By year 10 the shift is a highly statistically significant 2% loss in levels terms.

We briefly remark on the null paths here. Whilst above we could obtain simple closed-form solutions for the null path, here we cannot and we must turn to simulation. In Figure 2a, the rare disasters encompass very large losses but they can be spread over $1, 2, 3, \ldots$ or many years. Hence, when the null picks the worst θ fraction of events, the path declines gradually, reflecting a weighted average of all those possible disaster realizations that qualify. Conversely, in the right panel Figure 2b we have the average path that obtains under the null in all recessions minus the weighted contribution of those rare disaster paths. So this path doglegs down at first but then bounces back up, as we have excluded the slow moving disasters where the worst downside materializes.

[7]This setup would then have introduced a bivariate Mills ratio into the corresponding version of Equation 2,

Figure 2: *Tests for disasters: Barro rare disasters or not, advanced economies 1870–2017 (ex. war)*

The figure shows the cumulative a_h coefficients, with $a_h = \sum_{j=0}^{h} b_j$ for $h = 1, \ldots, 10$, with 90% confidence intervals from the bootstrap in brackets (500 replications). The event indicator denotes all BRD recessions in the full sample in the left panel, all $xBRD$ recessions in the full sample in the right panel. See text. Units are in log times 100.

(a) *Level coefficients, null α_h and data a_h, BRD recessions, full sample*

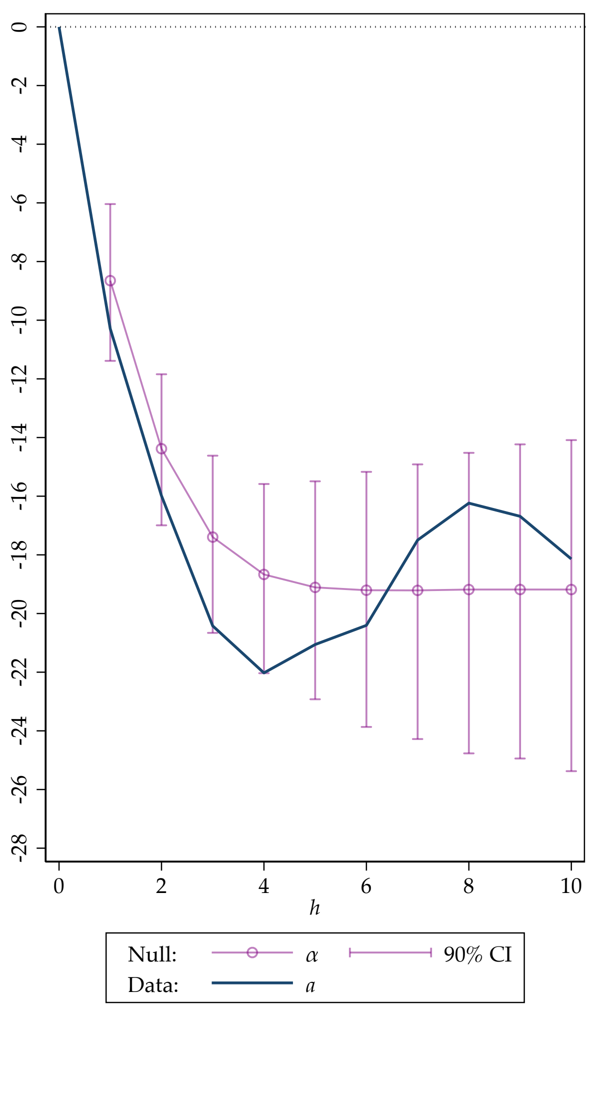

(b) *Level coefficients, null α_h and data a_h, xBRD recessions, full sample*

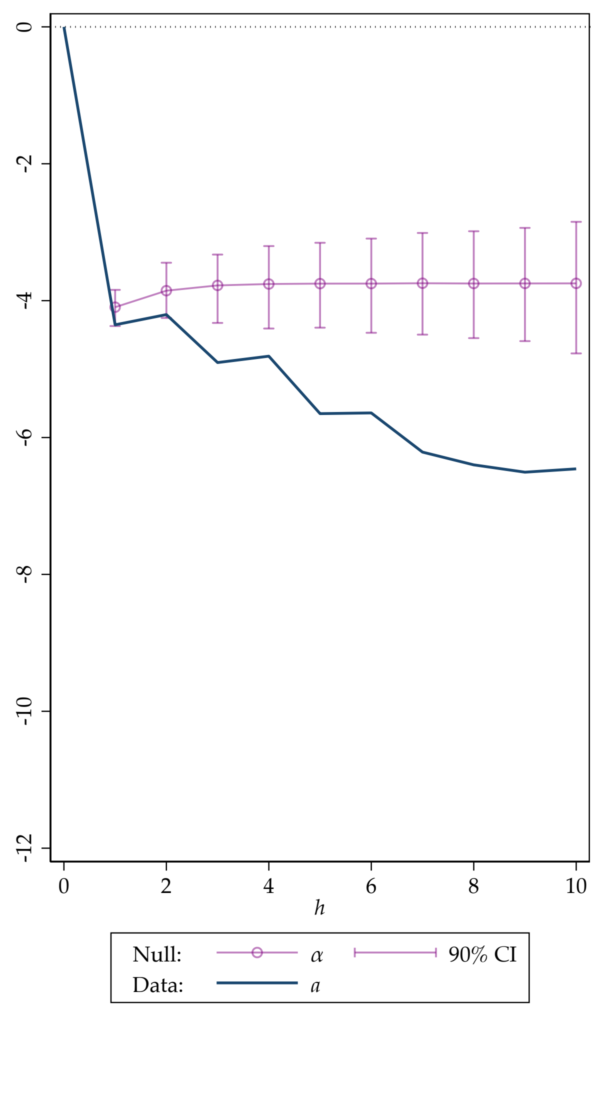

Robustness: the main result holds even in the post-WW2 sample Another concern might be that our results are driven by unusually volatile and skewed outcomes in the pre-WW2 period, and especially in the interwar period which includes the Great Depression. Even though some of those skewed outcomes might be filtered out above using the BRD sample exclusion, this is a more stringent test. To address that we repeat the analysis by fully excluding the entire pre-WW2 period from the sample.

And the main result still continues to hold, as reported in Figure 3a. Even in the Post-WW2 period, seen as being dominated by the era of so-called Great Moderation, and when there is only one BRD event (Spain, 2007), growth trajectories still deviate sharply downwards after a recession, and much more so than one would expect under the Gaussian null. By year 10 the deviation is a highly statistically significant 3% loss in levels terms.

This is a strong finding, since the rare disaster literature, and also the equity premium puzzle, literature has gone to great lengths to point out how we might have a biased view of the frequency of bad events if we only look at the nice and smooth growth outcomes seen from the 1950s to the 1990s in the advanced economies. So this sample restriction is, in a sense, precisely what one should *not* do if one wants to respect the idea that only the widest sample of events should be used so that we avoid cherry-picking good times. And yet, even when we try to cherry pick here, our test still throws up a rejection of the Gaussian null.

Robustness: the main result holds for normal and financial crisis recessions In a final robustness check, we confront another question. Is the non-Gaussianity we find just a manifestation of only a few moderately rare events that align with serious downturns, namely the phenomena of occasional financial crises? If so, our result may simply expand the scope of disaster dynamics by a little.

But the main result still continues to hold whether we look at financial crisis recessions or normal recessions, as reported in Figure 3b. The coefficients here are estimated separately for the two types of recessions, following the classification approach in Jordà, Schularick, and Taylor (2013). For sure, the deviation from the Gaussian null is much more adverse in financial crisis recessions, but it is still not trivial in a normal recession. In the latter case, average growth trajectories still deviate downwards after a recession, with a highly statistically significant 2% loss in levels terms by year 10. In the former case, the loss accumulates to about 7% at year 10.

which is why we used a simpler example at that point.

Figure 3: *Tests for disasters: postwar sample and normal-financial, advanced economies (ex. war)*

The figure shows the b_h and cumulative a_h coefficients, with $a_h = \sum_{j=0}^{h} b_j$ for $h = 1, \ldots, 10$, with 90% confidence intervals from the bootstrap in brackets (500 replications). The event indicator denotes all recessions in the full sample. The coefficients are also estimated separately for normal and financial crisis recessions, following the classification approach in Jordà, Schularick, and Taylor (2013). See text. Units are in log times 100.

(a) *Level coefficients, null α_h and data a_h, all recessions, post-WW2 sample*

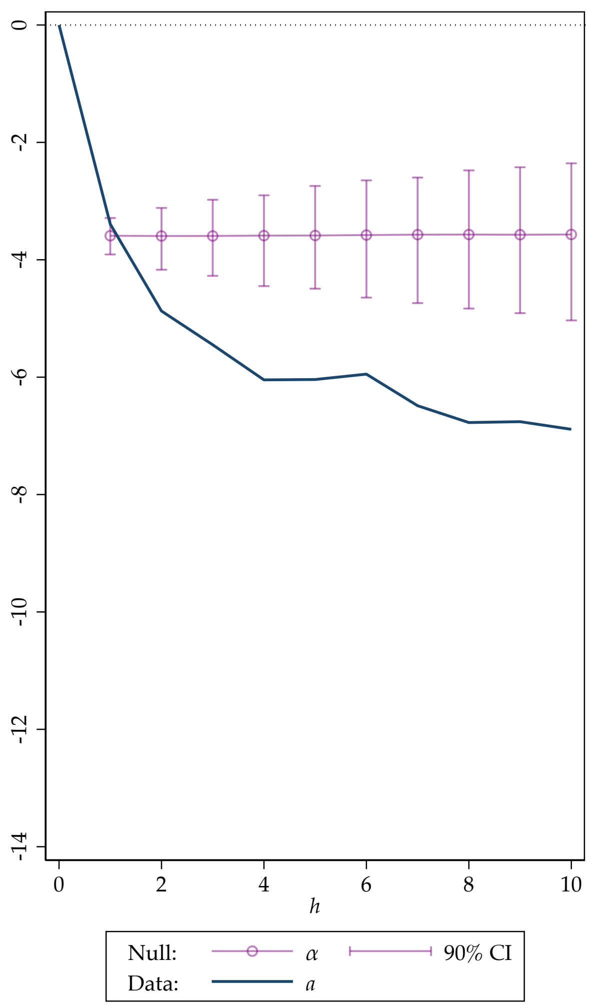

(b) *Level coefficients, null α_h and data a_h, normal and financial crisis recessions, full sample*

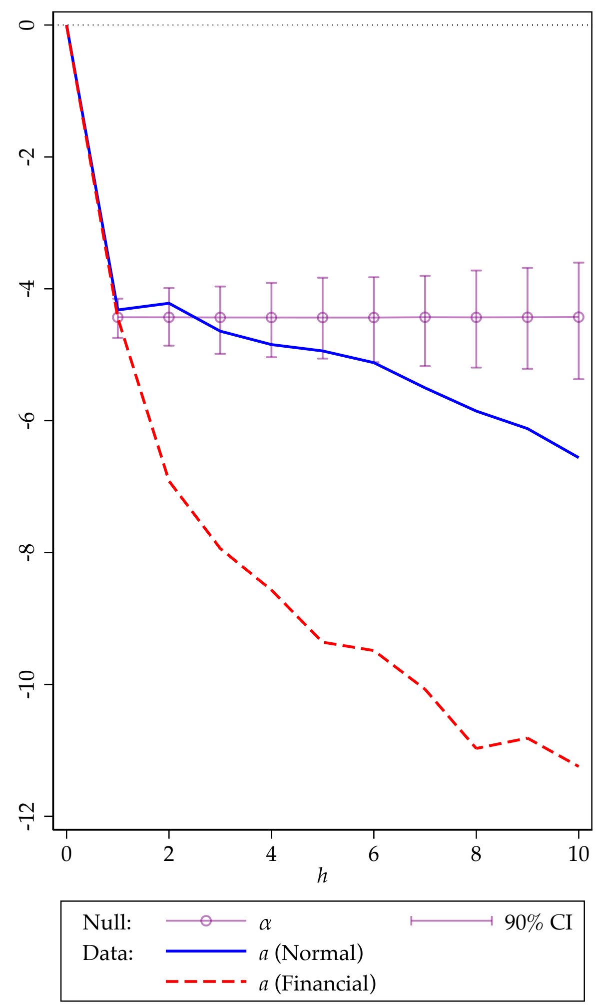

Table 4: *Level penalties, all and normal-financial recessions, advanced economies 1870–2017 (ex. war)*

The table shows the cumulated level penalty coefficients, defined as the deviation of the estimated coefficients from the Gaussian null, $\psi_h = a_h - \alpha_h$ for $h = 1, \ldots, 10$, estimated as in the text with 95% confidence intervals from the bootstrap in brackets ($B = 500$ replications). In column 1 the mapping (*All*) denotes all recessions in the full sample. Additionally in columns 2 and 3 the coefficients are estimated separately for normal and financial crisis recessions, following the classification approach in Jordà, Schularick, and Taylor (2013). Sample size is 2,149 observation on real per capita consumption growth Δx_{it}, with some losses to leads and lags at each h. See text. Units are in log times 100.

	All recessions	Normal recessions	Financial recessions
ψ_1	0.11	0.12	-0.01
	[-0.18, 0.42]	[-0.17, 0.43]	[-0.29, 0.31]
ψ_2	-0.29	0.23	-2.47
	[-0.73, 0.14]	[-0.22, 0.66]	[-2.91, -2.04]
ψ_3	-1.08	-0.22	-3.51
	[-1.55, -0.53]	[-0.68, 0.34]	[-3.98, -2.96]
ψ_4	-1.45	-0.45	-4.17
	[-1.98, -0.85]	[-0.98, 0.15]	[-4.70, -3.57]
ψ_5	-1.96	-0.57	-4.99
	[-2.56, -1.34]	[-1.17, 0.05]	[-5.59, -4.36]
ψ_6	-2.05	-0.76	-5.13
	[-2.66, -1.38]	[-1.37, -0.09]	[-5.74, -4.45]
ψ_7	-2.60	-1.15	-5.72
	[-3.23, -1.86]	[-1.77, -0.40]	[-6.34, -4.97]
ψ_8	-3.10	-1.52	-6.63
	[-3.81, -2.34]	[-2.23, -0.76]	[-7.34, -5.87]
ψ_9	-3.33	-1.81	-6.51
	[-4.08, -2.55]	[-2.56, -1.03]	[-7.25, -5.72]
ψ_{10}	-3.80	-2.28	-6.97
	[-4.63, -2.86]	[-3.11, -1.34]	[-7.79, -6.02]
N	2149	2149	2149

Summary In Table 4 we recap our main finding. After a recession event the *level penalties*, defined as the deviation of the estimated level coefficients from the Gaussian null model $\psi_h \equiv a_h - \alpha_h$ grow to become negative and highly significant out to the 10 year horizon. The full-sample ψ_h estimates show losses, along with 95% confidence intervals, and the loss amounts to -3.8% after 10 years in all recessions. Deviations are more substantial for financial crisis recessions, as expected, at around -6.9% after 10 years, but they are nontrivial even in normal recessions, reaching -2.3% after 10 years. These losses outstrip what would be expected from the Gaussian null model.

4. Alternative model: disaster paths with randomly varying severity

The previous section establishes that, on average, even plain vanilla recessions are like mini-disasters. The Gaussian null is rejected, and not just on account of a few very large, rare disasters. Instead, fat tails are pervasive, and growth skews negative in recessions.

But while we know that we can reject the null, that does not tell us how to set up an appropriate alternative model. The goal now is to propose a simple and tractable alternative model that builds on the well-established rare disaster framework, allows for variable severity disasters, and permits a direct calibration based on the local projection estimation toolkit. We can then use the model to conduct counterfactual welfare calculations.

Model: Random Coefficients Local Projections Following the rare disaster paradigm, we begin by conceiving of the economy as evolving under baseline stochastic growth path with a recurring risk of a disaster-type event defined as a recession of any form, as before. When such an event strikes, the economy is shunted onto a new forward-looking path whose dynamics we aim to characterize using a parsimonious set of stochastic shifts over a finite horizon H. These shifts apply repeatedly to future growth outcomes whenever a disaster-type event happens. For the moment, we restrict attention to a single disaster type for clarity. Also note that we present the model by omitting the panel dimension of our application to avoid notational clutter.

A simple way to pursue estimation is to build out from our earlier local projections specification in as follows,

$$\Delta x_t \;=\; \mu \;+\; \sum_{h=1}^{10} s_h d_{t-h} \zeta_{t-h} \;+\; \epsilon_t \; ; \quad \log \zeta_t \sim \mathcal{N}(-v^2/2, v^2) \; ; \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2) \; ; \qquad (8)$$

where μ is baseline drift, ζ and ϵ are independent from each other, and the $\{d_t\}_{t=1}^{T}$ are event indicators for a recession event, with $d_t = \mathbb{1}(\Delta x_t < 0 \;\&\; \Delta x_{t-1} \geq 0)$, i.e., an event is now defined as one year after a cyclical peak.

In Equation 8 the simplifying assumption is that if a disaster strikes at time $t-h$, then for all h, all impulse response coefficients s_h at horizon h after that event are shifted up or down by the same random-draw from a log-normal, but these will be different (random) amounts when looking across disaster events. That is, we implement variable disasters in our LP framework. Note that we will set the mean of the log-normal to $\mathbb{E}(\log \zeta) = -v^2/2$ so that $\mathbb{E}(\zeta) = 1$, using the properties of the log-normal, and hence these are pure scaling factors relative to the mean shift s_h. We also allow for no contemporaneous shifter, only with a lag, given that the Gaussian null was not rejected earlier in the first year after a peak.

We refer to this model the *random coefficient local projections* or RCLP model. To get a sense of what the RCLP specification buys us, notice that $\zeta \in (0, \infty)$ since $\log \zeta \in (-\infty, \infty)$. So, as $\log \zeta \to -\infty$, the disaster approaches the Gaussian null model and shown in Equation 1. As $\log \zeta \to +\infty$, the disaster becomes increasingly severe. Below, we estimate $\sigma_\zeta \approx 0.39$. This means that the penalty coefficients, s_h are scaled up or down as follows. Take a centered 95% probability range for ζ. At the low end, s_h is scaled down by a factor of approximately 0.47. At the other end, it gets scaled up by a factor of approximately 2.1. In other words, in this range, a harsh recession is about 4.6 times worse than a mild one when comparing the 2.5% quantile to the 97.5% quantile of the ζ draw.

We set a maximal horizon for estimation $H = 10$ years as a practical matter since few expansions last more than 10 years, so we have insufficient observations at large horizons to get plausibly accurate parameter estimates. Beyond that horizon we implicitly assume in the growth rate regression that the process reverts to the random walk with drift for $h > H$. That is, the growth losses captured by the s_h parameters will cease at that point.

Estimation of Equation 8 is mostly standard, even when extended to a panel setting, with allowance for country fixed effects to capture differences in drift. The only complication is to obtain an estimate of the variance of $\log \zeta$, denoted v^2, where we have

$$\begin{aligned}
V(\Delta x_t) &= V\left(\sum_{h=1}^{H} s_{t-h}\zeta_{t-h}\right) + V(u_t) \\
\implies V(\Delta x_t) - V(u_t) &= \sum_{h=1}^{H} s_h^2[q(1-q)V(\zeta) + q(1-q) + V(\zeta)q^2] \\
&= \sum_{h=1}^{H} s_h^2 q(1-q)V(\zeta) + \sum_{h=1}^{H} s_h^2 q(1-q) + \sum_{h=1}^{H} s_h^2 q^2 V(\zeta) \\
\implies V(\zeta) &= \frac{V(\Delta x_t) - V(u_t) - \sum_{h=1}^{H} s_h^2 q(1-q)}{q\sum_{h=1}^{H} \psi_h^2} \\
\implies \quad v &= \log\left[V(\zeta) + 1\right]^{1/2}
\end{aligned}$$

Results Using the methods above, we now estimate a baseline RCLP model on the data that we subsequently simulate for use in welfare calculations of the cost of business cycles. Hence, we focus only on real consumption per capita in what follows. Thus, let x_{it} denote the log of annual real consumption per capita across years $t = 1, ..., T$ and countries $i = 1, ..., I$. The relevant sample will be drawn from the same peacetime dataset as above, with recession events sorted into two types, recessions associated with financial crises (F) and normal recessions (N) as before.

Figure 4: *Recession paths for normal and financial crisis recessions: Random Coefficient Local Projections*

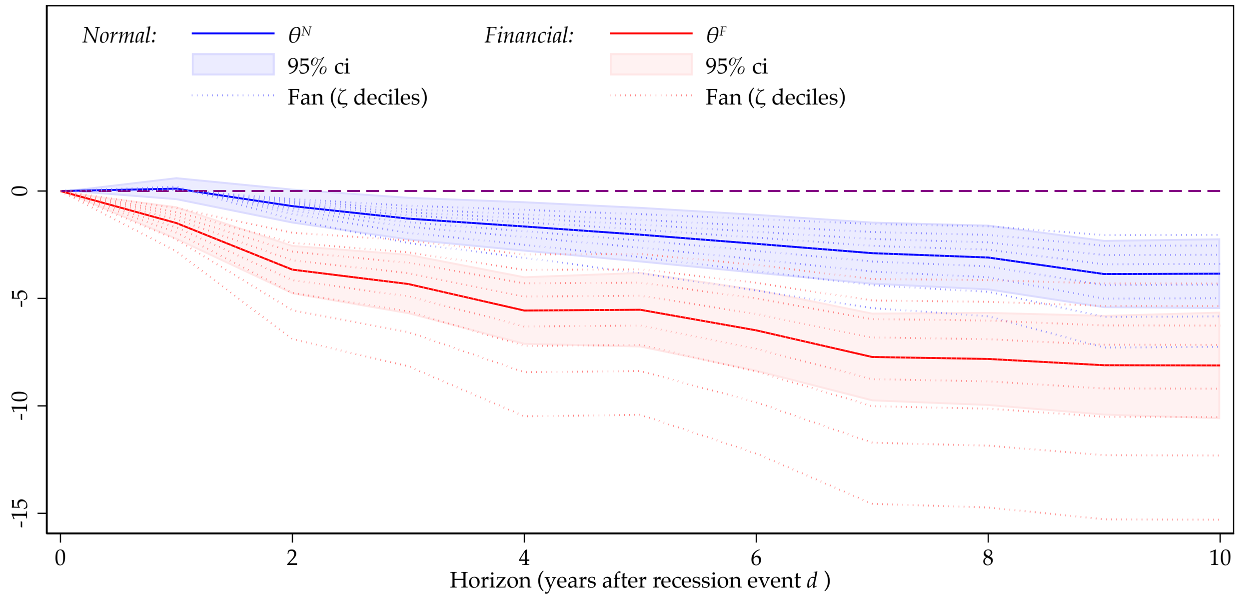

Notes: The figure shows Normal versus Financial Crisis recession paths based on Random Coefficient Local Projections. The event at year $h = 0$ is the first year of a recession, so the peak is a time $h = -1$. There is no effect at $h = 0$. The coefficients of the growth versions are reported in cumulative form by adding up the coefficient estimates appropriately. The estimates are scaled by 100 to show the results in percent (log $\times$100) deviations from the peak of consumption. 95% confidence intervals are provided. See text.

Hence, we estimate the panel RCLP model

$$\Delta x_t = \mu + \sum_{h=1}^{10} (s_h^N d_{t-h}^N + s_h^F d_{t-h}^F)\zeta_{t-h} + u_i + \epsilon_{it} \, ; \, \log \zeta_t \sim \mathcal{N}(m, v^2) \, ; \, \epsilon_t \sim \mathcal{N}(0, \sigma^2) \, ;$$

where the intercept μ is the average of the country fixed effects, with the country-weighted average of the u_i constrained to zero.

Our baseline estimate of the RCLP model is shown in Table 5 and Figure 4. In the table, the growth penalties are substantial, resemble what we saw earlier in the LP tests, and cumulate over 10 years to -3.8% for normal recessions and -8.1% for financial crisis recessions. In the figure, estimated cumulative coefficients $\theta_h^k = \sum_{j=1}^h s_j^k$, with $k = N, F$ are shown, with 95% confidence intervals. Recall that these coefficients are average level effects h years after an event, corresponding to a draw of $\zeta = 1$. To show the range of variable penalties the figure also displays a fan chart of dotted lines corresponding to deciles of the distribution of ζ, where blue denotes normal and red denotes financial crisis recessions.

Summing up, our estimates of the data generating process appear reasonable and closely resemble the type of negatively-skewed deviations from the Gaussian null seen in our earlier LP tests, for example, as in Figure 3b. With these estimates in hand we now turn to simulating these DGPs for use in counterfactual welfare analysis.

Table 5: *Random Coefficient Local Projections: Estimates*

	Δx
s_1^N	0.109
	(0.271)
s_2^N	-0.811**
	(0.275)
s_3^N	-0.582*
	(0.275)
s_4^N	-0.365
	(0.278)
s_5^N	-0.382
	(0.278)
s_6^N	-0.426
	(0.280)
s_7^N	-0.434
	(0.281)
s_8^N	-0.203
	(0.282)
s_9^N	-0.770**
	(0.281)
s_{10}^N	0.0179
	(0.277)
s_1^F	-1.490***
	(0.397)
s_2^F	-2.168***
	(0.402)
s_3^F	-0.669
	(0.405)
s_4^F	-1.235**
	(0.402)
s_5^F	0.0369
	(0.406)
s_6^F	-0.957*
	(0.424)
s_7^F	-1.242**
	(0.442)
s_8^F	-0.0910
	(0.459)
s_9^F	-0.292
	(0.496)
s_{10}^F	-0.00921
	(0.504)
μ	2.643***
	(0.132)
$\sum_{h=1}^{10} s_h^N$	-3.847***
	(0.838)
$\sum_{h=1}^{10} s_h^F$	-8.117***
	(1.277)
σ	3.552
v	0.494
N	1919

Notes: Dependent variable: $\Delta c_{i,t} \times 100$. The table displays regression coefficients for the RCLP model with H=10. Standard errors in parentheses. $^* p < 0.10, ^{**} p < 0.05, ^{***} p < 0.01$. The row v displays the estimate of the variance of the random lognormal draw. The estimated drift is the intercept μ, which is the average of the country fixed effects. See text.

5. Costs of business cycles with disasters everywhere

Recessions are not just bad consumption draws from a random walk with drift. They have patterns that defy this Gaussian null model in ways we have shown above—in this sense, disasters are everywhere, and not only of the rare kind.

If so, how much is a representative consumer willing to pay to avoid such pervasive and weighty left-tail randomness? In this section we take the lessons from our empirical work seriously and apply them in a standard welfare counterfactual exercise.

Model setup and calibration Our approach will be to simulate the RCLP model of consumption growth as the true data generating process using the above estimated parameters.

After a "good" no-disaster event at time t, consumption will not be affected in any future years. But if a "bad" disaster event takes place, with some probability the normal or financial type of disaster is decided, $k \in \{N, F\}$, and then the path of consumption deviates at $t+h$ by a penalty s_h^k as above, adjusting its severity by the log-normal draw of ζ_t. Thus, the growth deviations applied to the Gaussian DGP at time $t+h$ are $s_h^k \zeta_t$ as $t+h$ counts upwards within a given cycle for $h = 1, \ldots, H$, with all such penalties applied additively at any future date.

To proceed with the above model, we simulate data for annual consumption growth. The simulation is for an individual country, so the i subscript is dropped. Just as described above, we will assume consumption follows a process given by a drift with penalties,

$$\Delta c_t \ = \ \hat{\mu} \ + \underbrace{\sum_{h=1}^{10} \left[\hat{s}_h^F d_{t-h}^F \zeta_{t-h} + \hat{s}_h^N d_{t-h}^N \zeta_{t-h} \right]}_{penalties} + \epsilon_t \,, \tag{RCLP}$$

where ϵ_t is drawn from $\mathcal{N}(0, \hat{\sigma}^2)$, and $\log \zeta_t$ is drawn from $\mathcal{N}(-v^2/2, v^2)$. The parameters $\hat{\mu}$, $\hat{s}_h^f$, $\hat{s}_h^N$, $\hat{\sigma}^2$, and v^2 are taken from the RCLP estimates reported in Table 5.

Here, a disaster event is endogenously given by $d_t = \mathbb{1}(\Delta x_t < 0 \ \& \ \Delta x_{t-1} \geq 0)$, i.e., an event is now defined as one year after a cyclical peak. We then further define two types of events $d_t^N = d_t N_t$ and $d_t^F = d_t F_t$, wher a normal recession indicator is $N_t = 1$ and a financial recession indicator is $F_t = 1$.

To fully specify the simulation, it remains to define the process governing the evolution of the disaster type, for which we define event-type indicators indicating when there is a recession (normal or financial), namely $R_t = \{N_t, F_t\}$. This requires that we make a choice of an independent draw probability (q_n, q_f) for whether the type is normal or financial, with $q_n + q_f = 1$. So N_t, F_t are Bernoulli i.i.d. draws, with $P(N_t = 1) = q_N$ and $P(F_t = 1) = q_F$.

Candidates for these probabilities are chosen as follows, based on empirical frequency in the data in the full sample of the JST dataset, excluding wars, and these will help us illustrate how varying the sub-draw probabilities (q_n, q_f) affects welfare outcomes:

- **Zero financial crisis risk** $q_f = 0, q_n = 1$, empirical frequencies for the "quiet" 1950s–60s era.
- **Medium financial crisis risk** $q_f = \bar{q}_f = 0.25, q_n = \bar{q}_n = 0.75$, approximate empirical frequencies for the full sample.
- **High financial crisis risk** $q_f = 0.50, q_n = 0.50$, approximate empirical frequencies of the post-1985 era.
- **Variable financial crisis risk** any q_f, q_n combination.

To afford welfare comparisons with other benchmarks from the literature, we complement our simulated models above with three additional simulated models of consumption growth: a deterministic trend, to be used as a welfare baseline; a "Deterministic plus Gaussian" path, which adds i.i.d. Gaussian shocks to get a random walk with drift; and a simulated Barro RD rare disaster process, with Bernoulli probability π and draws b_t from an empirical distribution of the 65 output disasters in Barro (2006), and applied additively to the Deterministic + Gaussian drift model above. Formally, we can write these as follows,

$$\Delta c_t = \hat{\mu}\,, \qquad \text{(Deterministic)}$$
$$\Delta c_t = \hat{\mu} + \epsilon_t\,, \qquad \text{(Deterministic + Gaussian)}$$
$$\Delta c_t = \mu^b + \pi b_t + \epsilon_t\,. \qquad \text{(Deterministic + Gaussian + Barro RD)}$$

Here, $\hat{\mu}$ is the conditional drift in our RCLP model estimation above, and $\hat{\sigma}_\epsilon^2$ is residual variance. Note that $\hat{\mu} + E(penalties) = E(\Delta c)$ For comparability, the RD drift is chosen so its DGP has the same mean, with $\mu^b + E(\pi b) = E(\Delta c)$.

By construction, the full model matches the observed mean drift. But what is being assumed in the reference Deterministic counterfactual? As is standard practice in the disasters literature, we assume that all permanent losses and temporary skew arising from the disaster or penalty terms are completely switched off in the counterfactuals, as well as the Gaussian residual (see, e.g., Barro, 2006).[8] In the case of paths with permanent losses (i.e., our recession paths or the Barro RD paths) this means that the counterfactual process

[8] An implicit assumption is that a move to the counterfactual will not damage mean drift, but only tamp down higher moments of growth. Policies may or may not achieve that objective and specifics matter. This remains an open debate in theory and empirics (see, e.g., Rancière, Tornell, and Westermann, 2008).

will inherit a higher mean growth, as well as being stripped of negative-skewed terms. To be clear, these changes will contribute to higher welfare for the representative consumer from increases in the mean as well as the elimination of all higher moments.

For numerical implementation, we simulate a large history of all series over $T = 10{,}000$ years for the benchmark models indexed by $M \in \{\text{Deterministic}, \text{Deterministic} + \text{Gaussian}\}$ and for the simulated RCLP models indexed by $M = q_f \in [0, 1]$ using a grid of size 0.05. We obtain a welfare estimate for each model by averaging the discounted lifetime utility (truncating over 500 years forwards) at all dates. We assume CRRA preferences with a baseline low risk aversion parameter $\gamma = 4$ (as in Barro, 2006, 2009) and a discount factor of $\beta = 0.96$ for annual data.[9] The welfare level for model M is then a function of simulated real consumption per capita $C_{M,t}$ in that model,

$$\text{Welfare}_M = W(C_{M,t}) \equiv \frac{\sum_{s=1}^{T} \sum_{k=1}^{K} \beta^k \frac{C_{M,s+k}^{1-\gamma}}{1-\gamma}}{S},$$

where s refers to the start period of the sample, and $K = 500$ is the truncation point. We will then compute the consumption-equivalent welfare ratio of any model versus the Deterministic baseline CE_M implied by the formula $W(C_{M,t}) = W(CE_M \times C_{Deterministic,t})$.

Comparing welfare under actual and counterfactual histories Our main result is shown in Figure 5. Using the deterministic model as a baseline, the vertical axis shows consumption-equivalent welfare ratio relative to the baseline. The horizontal axis shows the relative probability of a financial crisis for the subdraw q_f which varies between zero and one. The reference level for the deterministic model is the thin dotted blue line at $CE = 1$.

Consider first the Deterministic + Gaussian model, shown by the thin dashed blue line. The welfare loss is fixed, about 1%, as expected and consistent with (Obstfeld, 1994). In contrast, the red line shows the welfare losses for our baseline RCLP disaster model, which are roughly and order of magnitude larger, ranging between about 9% and 20%.

The red line shows how financial crisis disaster frequency q_f impacts the losses relative to the deterministic path. In a simulated RCLP world of high financial crisis risk the welfare loss of 15% is not so different from the loss of 18% generated in a standard Barro RD calibration, shown by the thin solid blue line, under the same CRRA welfare function.[10]

[9] As for our functional form, Lucas (1987, 2003) and Barro (2006) preferences are CRRA, but Barro (2009) employs Epstein-Zin-Weil preferences. The latter tend to magnify welfare losses, all else equal. However, even in our CRRA environment with a conservative choice of γ, we can generate considerable welfare losses: in our setup, disaster-style fat tails are a feature of all recessions, and weigh heavily in the full welfare cost.

[10] This 18% figure for the Barro RD case is close to the 17% in footnote 1. The small difference is due to a slightly different calibration of drift and residual variance for consistency with our baseline.

Figure 5: *Welfare simulations: Main results*

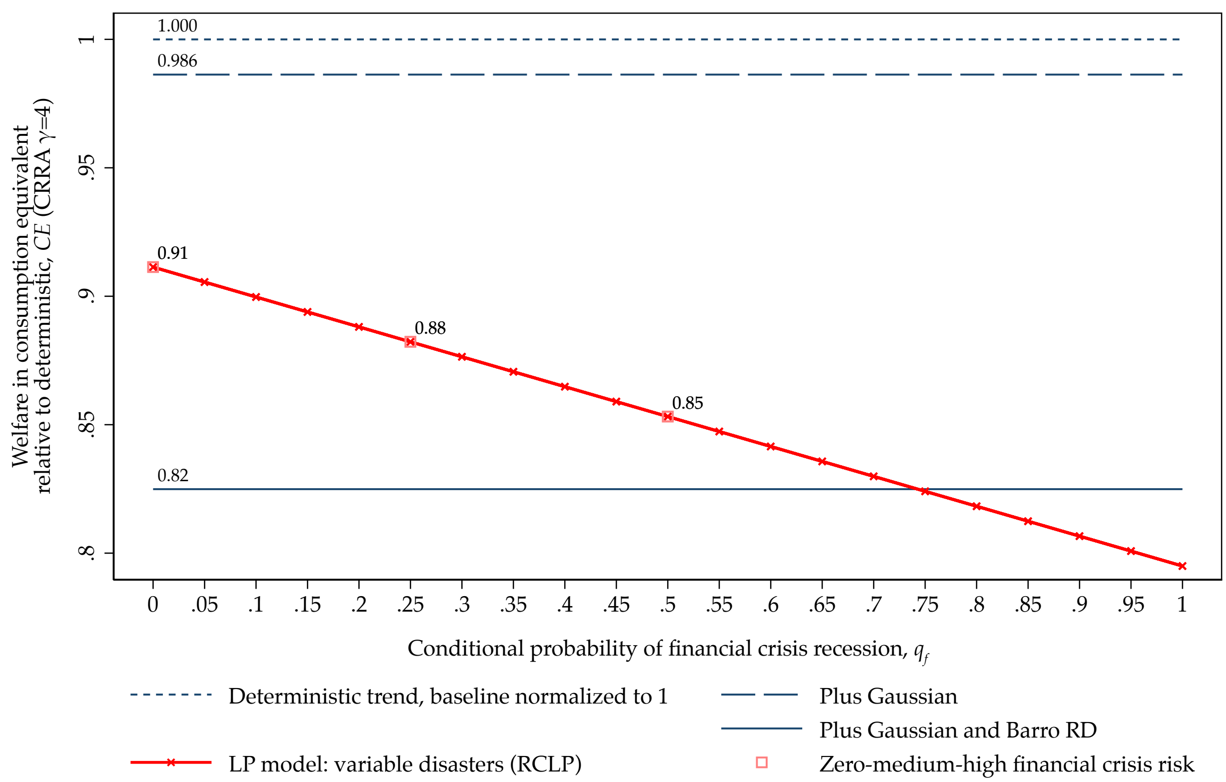

Notes: Consumption-equivalent welfare ratios associated with our baseline model and for the benchmarks, relative to the Deterministic model. The squares correspond to the zero, medium, and high frequency crisis risk levels. See text.

To focus on a few reference levels, three points on the line are picked out with squares corresponding to zero, medium, and high frequency crisis risk levels. A representative agent in a world of high financial crisis risk, like the recent decades, ($q_f = 0.50$, post-1985) suffers a welfare loss of about 15% compared to the deterministic path. In a world of medium financial crisis risk ($q_f = \bar{q}_f = 0.25$, the historical average) they lose about 12%. But even with *zero* financial crisis risk, and all recessions constrained to be of the normal type ($q_f = 0.00$, 1950s–1960s), they still lose 9% compared to the deterministic path.

These are large welfare costs, and they arise without wartime rare disasters present at all in the analysis, and with no rare disasters of any kind in the simulation ($q_f = 0$) restricted to purely normal peacetime recessions. These results speak not only to the massive welfare losses associated with frequent financial crises, but also to the hitherto ignored but nontrivial welfare losses felt even in normal recessions. The latter effects have not been captured in traditional models Lucas (1987); Obstfeld (1994); Lucas (2003); Barro (2006) which treat normal recessions as deriving from Gaussian processes with no fat tail disaster attributes. Instead, as our LP tests have shown, even normal recessions are non-

Figure 6: *Welfare simulations: Shutting down growth penalties in normal recessions*

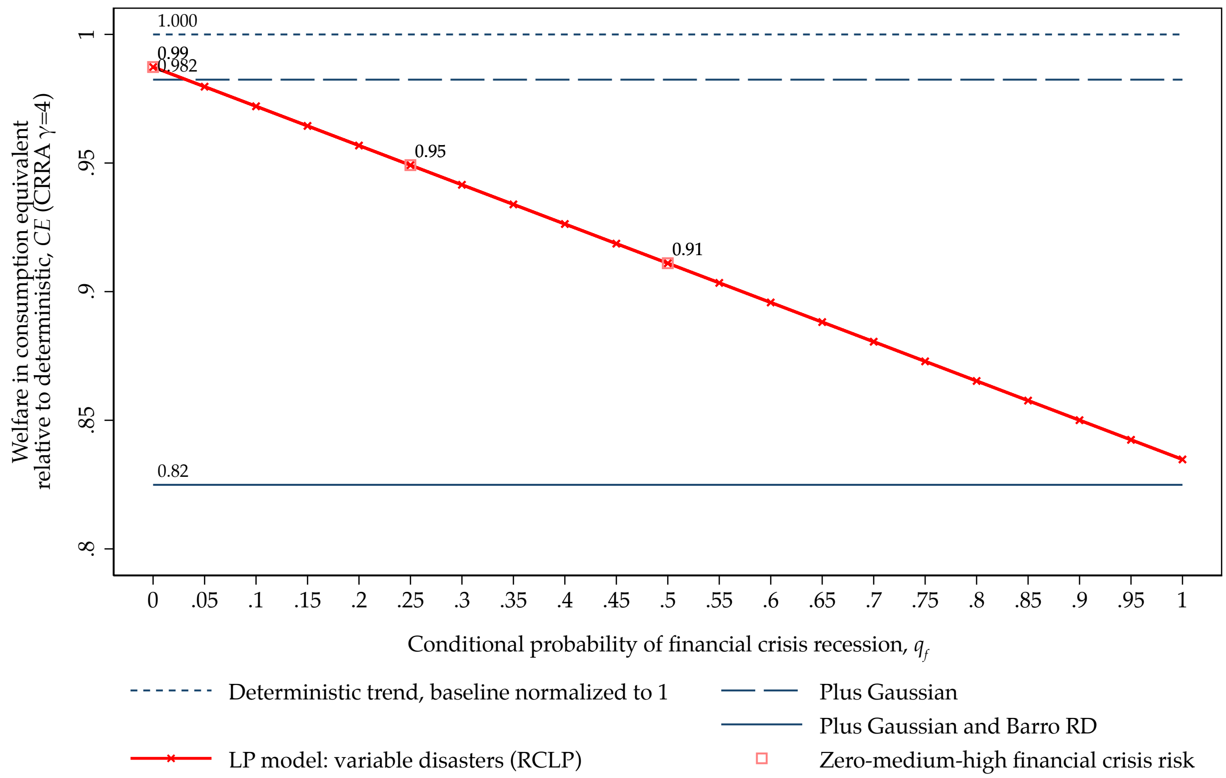

Notes: Consumption-equivalent welfare ratios associated with our baseline model with the restriction $s_h^N =$ and for the benchmarks, relative to the Deterministic model. The squares correspond to the zero, medium, and high frequency crisis risk levels. See text.

Gaussian, with significant fat tails in the consumption path relative to the null. In welfare terms, this really matters. If the rare disaster model has reawakened concerns about the welfare costs of non-deterministic fluctuations, then our findings about the broader costs of "disasters everywhere" with pervasive fat-tailed dynamics are of comparable quantitative import.

Robustness: Shutting down growth penalties in normal recessions A robustness check shown in Figure 6 examines shutting down the fat-tail for normal recessions. The RCLP is re-estimated with all normal recession penalty coefficients S_h^N constrained to zero, and the new parameters are used to construct a new set of simulated consumption paths as before.

This change should tilt the dependence of welfare costs on crisis risk substantially. In the case $q_f = 0$, when there are no crisis financial recessions, the DGP (RCLP) will be simply the Deterministic + Gaussian model, and welfare losses would then be tied down close to the 1% level. However, when $q_f = 1$ when there are *only* crisis financial recessions, the welfare losses would then inherit the full set of variable disaster penalties as before.

As Figure 6 shows, this intuition is correct. However, we see that even moderate levels of crisis risk can drive up welfare losses quite quickly. In a world of medium financial crisis risk ($q_f = 0.25$, the historical average) the consumption-equivalent welfare loss is 5%. But in a world of high financial crisis risk, ($q_f = 0.50$, post-1985) the loss rises to 9%.

We caution that our RCLP estimates clearly reject this model, with skew evident even in normal recessions. However, some might have a strong a priori that a normal recession should leave no permanent mark. But even then, the non-trivial risk of financial crises and the strong evidence of persistent drag from these events (Cerra and Saxena, 2008; Jordà, Schularick, and Taylor, 2013; Reinhart and Rogoff, 2014) should focus special concern on this case, with large potential gains from careful macroprudential policies which can mitigate these damaging left-tail scenarios.

6. Conclusions

At the time of writing this paper, the world's economies were recovering from one of the largest and most sudden declines in output due to the COVID-19 pandemic. This is one example of the rare disasters that Barro (2006, 2009) saw as a cause of the large equity premium and large welfare losses.

Based on a historical frequency of less than 2%, a person with a life expectancy of 80 years will see on average 1.5 rare disasters in their lifetime—maybe a great war, a depression, or even a pandemic. But these are just the very worst events, and beyond this we would infer that they will have to endure on average 12 other recessions, of which 3 on average will be financial crises. Those events may not be as catastrophic, but given the pattern of persistent downside skew we have documented, the associated welfare losses from such asymmetric shocks look far from trivial and can still mount up.

Our argument in this paper is that the rare disasters in the outer-tail are not the only events that matter. We find in the long-run historical data that business cycles tend to be asymmetric and the inner-tail events resemble "mini-disasters". Consumers also experience considerable welfare losses from less extreme but more frequent peacetime recessions. The depth and duration of such recessions are variable, and because they cause skewed deviations from trend growth that can last for extended periods, households would pay a nontrivial cost to insure against them. Disasters, in this sense, are everywhere.

Our paper re-calculates the costs of business cycles in this setting of frequent fat tails in advanced economies. The size of the consumption-equivalent welfare loss that we find is large: 15% of permanent consumption for cycles in the crisis-prone past three decades under very moderate risk aversion, and 12% over the full historical sample back to 1870. These losses are more than two magnitudes above that in a Lucas (1987) trend-stationary model, and more than one magnitude above that in an Obstfeld (1994) random-walk-with-drift model. These losses are in the same order of magnitude as those arising from a pure rare disaster model like Barro (2006).

The welfare costs of business cycles have increased in the recent decades as financial crises have become more frequent. If these results are a good approximation of reality, then substantial gains in welfare could be achieved from well-designed policies to prevent financial crises and mitigate even normal recessions.

References

Acemoglu, Daron, and Andrew Scott. 1994. Asymmetries in the Cyclical Behaviour of UK Labour Markets. *Economic Journal* 104(427): 1303–1323.

Atkeson, Andrew, and Christopher Phelan. 1994. Reconsidering the Costs of Business Cycles with Incomplete Markets. *NBER Macroeconomics Annual* 9: 187–207.

Barro, Robert J. 2006. Rare Disasters and Asset Markets in the Twentieth Century. *Quarterly Journal of Economics* 121(3): 823–866.

Barro, Robert J. 2009. Rare Disasters, Asset Prices, and Welfare Costs. *American Economic Review* 99(1): 243–264.

Barro, Robert J., and Tao Jin. 2011. On the Size Distribution of Macroeconomic Disasters. *Econometrica* 79(5): 1567–1589.

Barro, Robert J., and José F. Ursúa. 2008a. Consumption Disasters in the Twentieth Century. *The American Economic Review* 98(2): 58–63.

Barro, Robert J., and José F. Ursúa. 2008b. Macroeconomic Crises since 1870. *Brookings Papers on Economic Activity* 2008: 255–335.

Busch, Christopher, David Domeij, Fatih Guvenen, and Rocio Madera. 2018. Asymmetric Business-Cycle Risk and Social Insurance. NBER Working Papers 24569.

Calvo, Guillermo, Fabrizio Coricelli, and Pablo Ottonello. 2013. Jobless Recoveries During Financial Crises: Is Inflation the Way Out? Working Paper 19683, National Bureau of Economic Research.

Cerra, Valerie, and Sweta Chaman Saxena. 2008. Growth Dynamics: The Myth of Economic Recovery. *American Economic Review* 98(1): 439–57.

Cochrane, John H. 1988. How Big Is the Random Walk in GNP? *Journal of Political Economy* 96(5): 893–920.

Cogley, Timothy. 1990. International Evidence on the Size of the Random Walk in Output. *Journal of Political Economy* 98(3): 501–518.

Colacito, Riccardo, Eric Ghysels, Jinghan Meng, and Wasin Siwasarit. 2016. Skewness in Expected Macro Fundamentals and the Predictability of Equity Returns: Evidence and Theory. *Review of Financial Studies* 29(8): 2069–2109.

deLong, J. Bradford, and Lawrence H. Summers. 1988. How Does Macroeconomic Policy Affect Output? *Brookings Papers on Economic Activity* 1988(2): 433–494.

Dew-Becker, Ian, Alireza Tahbaz-Salehi, and Andrea Vedolin. 2019. Macro skewness and conditional second moments: evidence and theories. Northwestern University. Unpublished.

Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson. 2019. A Plucking Model of Business Cycles. NBER Working Paper 26351.

Fatás, Antonio, and Ilian Mihov. 2013. Recoveries. CEPR Discussion Paper 9551.

Gabaix, Xavier. 2008. Variable Rare Disasters: A Tractable Theory of Ten Puzzles in Macro-Finance. *The American Economic Review* 98(2): 64–67.

Gourio, François. 2012. Disaster Risk and Business Cycles. *American Economic Review* 102(6): 2734–2766.

Heckman, James. 1974. Shadow prices, market wages, and labor supply. *Econometrica: journal of the econometric society* 679–694.

Imrohoroğlu, Ayşe. 1989. Cost of Business Cycles with Indivisibilities and Liquidity Constraints. *Journal of Political Economy* 97(6): 1364–1383.

Jordà, Òscar. 2005. Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review* 95(1): 161–182.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2013. When Credit Bites Back. *Journal of Money, Credit and Banking* 45: 3–28.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2017. Macrofinancial History and the New Business Cycle Facts. In *NBER Macroeconomics Annual 2016, Volume 31*, 213–263. Chicago: University of Chicago Press.

Keynes, John Maynard. 1936. Notes on the Trade Cycle. In *The General Theory of Employment, Interest, and Money*, 279–295. London: Macmillan.

Kim, Chang-Jin, and Charles R. Nelson. 1999. Friedman's Plucking Model of Business Fluctuations: Tests and Estimates of Permanent and Transitory Components. *Journal of Money, Credit and Banking* 31(3): 317–334.

Kuersteiner, Guido M. 2005. Automatic inference for infinite order vector autoregressions. *Econometric Theory* 85–115.

Lucas, Robert E., Jr. 1987. *Models of business cycles*. Oxford: Basil Blackwell.

Lucas, Robert E., Jr. 2003. Macroeconomic Priorities. *American Economic Review* 93(1): 1–14.

Martin, Ian W. R. 2008. Disasters and the Welfare Cost of Uncertainty. *American Economic Review* 98(2): 74–78.

Morley, James, and Jeremy Piger. 2012. The Asymmetric Business Cycle. *Review of Economics and Statistics* 94(1): 208–221.

Muir, Tyler. 2017. Financial Crises and Risk Premia. *Quarterly Journal of Economics* 132(2): 765–809.

Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa. 2013. Crises and Recoveries in an Empirical Model of Consumption Disasters. *American Economic Journal: Macroeconomics* 5(3): 35–74.

Neftçi, Salih N. 1984. Are Economic Time Series Asymmetric over the Business Cycle? *Journal of Political Economy* 92(2): 307–328.

Obstfeld, Maurice. 1994. Evaluating risky consumption paths: The role of intertemporal substitutability. *European Economic Review* 38(7): 1471–1486.

Rancière, Romain, Aaron Tornell, and Frank Westermann. 2008. Systemic Crises and Growth. *Quarterly Journal of Economics* 123(1): 359–406.

Reinhart, Carmen M., and Kenneth S. Rogoff. 2014. Recovery from Financial Crises: Evidence from 100 Episodes. *American Economic Review* 104(5): 50–55.

Reis, Ricardo. 2009. The Time-Series Properties of Aggregate Consumption: Implications for the Costs of Fluctuations. *Journal of the European Economic Association* 7(4): 722–753.

Rietz, Thomas A. 1988. The equity risk premium a solution. *Journal of Monetary Economics* 22(1): 117–131.

Salgado, Sergio, Fatih Guvenen, and Nicholas Bloom. 2019. Skewed Business Cycles. NBER Working Paper 26565.

Schmitt-Grohé, Stephanie, and Martin Uribe. 2013. Downward Nominal Wage Rigidity and the Case for Temporary Inflation in the Eurozone. *Journal of Economic Perspectives* 27(3): 193–212.

Schularick, Moritz, and Alan M. Taylor. 2012. Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008. *The American Economic Review* 102(2): 1029–1061.

Sichel, Daniel. 1993. Business Cycle Asymmetry: A Deeper Look. *Economic Inquiry* 31(2): 224–36.

Ursúa, José F. 2011. *Essays in Macroeconomics and Asset Pricing: Long-Run Risks, Disasters, and Trends*. Ph.D. thesis, Harvard University.

Wachter, Jessica A. 2013. Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance* 68(3): 987–1035.