

DISCUSSION PAPER SERIES

DP14559

(v. 2)

DISASTERS EVERYWHERE: THE COSTS OF BUSINESS CYCLES RECONSIDERED

Òscar Jordà, Moritz Schularick and Alan M. Taylor

ECONOMIC HISTORY

MONETARY ECONOMICS AND FLUCTUATIONS



DISASTERS EVERYWHERE: THE COSTS OF BUSINESS CYCLES RECONSIDERED

Òscar Jordà, Moritz Schularick and Alan M. Taylor

Discussion Paper DP14559
First Published 04 April 2020
This Revision 07 August 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Economic History
- Monetary Economics and Fluctuations

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Òscar Jordà, Moritz Schularick and Alan M. Taylor

DISASTERS EVERYWHERE: THE COSTS OF BUSINESS CYCLES RECONSIDERED

Abstract

Business cycles are costlier and stabilization policies could be more beneficial than widely thought. This paper introduces a new test to show that all business cycles are asymmetric and resemble “mini-disasters.” By this we mean that growth is pervasively fat-tailed and non-Gaussian. Using long-run historical data, we show empirically that this is true for advanced economies since 1870. Focusing on peacetime eras, we develop a tractable local projection framework to estimate consumption growth paths for normal and financial-crisis recessions. Introducing random coefficient local projections (RCLP) we get an easy and transparent mapping from the estimates to a calibrated simulation model with disasters of variable severity. Simulations show that substantial welfare costs arise not just from the large rare disasters, but also from the smaller but more frequent mini-disasters in every cycle. On average, and in post-WW2 data, even with low risk aversion, households would sacrifice about 15 percent of consumption to avoid such cyclical fluctuations.

JEL Classification: E13, E21, E22, E32

Keywords: Fluctuations, Asymmetry, local projections, Random coefficients, macroprudential policy

Òscar Jordà - oscar.jorda@sf.frb.org

Federal Reserve Bank of San Francisco, University of California, Davis

Moritz Schularick - moritz.schularick@uni-bonn.de

University of Bonn and CEPR

Alan M. Taylor - amtaylor@ucdavis.edu

University of California, Davis and CEPR

Acknowledgements

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, and we are grateful for this support. Schularick 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. 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.

Disasters Everywhere: The Costs of Business Cycles Reconsidered^{*}

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

August 2020

Abstract

Business cycles are costlier and stabilization policies could be more beneficial than widely thought. This paper introduces a new test to show that all business cycles are asymmetric and resemble “mini-disasters.” By this we mean that growth is pervasively fat-tailed and non-Gaussian. Using long-run historical data, we show empirically that this is true for advanced economies since 1870. Focusing on peacetime eras, we develop a tractable local projection framework to estimate consumption growth paths for normal and financial-crisis recessions. Introducing random coefficient local projections (RCLP) we get an easy and transparent mapping from the estimates to a calibrated simulation model with disasters of variable severity. Simulations show that substantial welfare costs arise not just from the large rare disasters, but also from the smaller but more frequent mini-disasters in every cycle. On average, and in post-WW2 data, even with low risk aversion, households would sacrifice about 15 percent of consumption to avoid such cyclical fluctuations.

Keywords: fluctuations, asymmetry, 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, and we are grateful for this support. Schularick 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. 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).

[‡]Federal Reserve Bank of New York; and 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 designed to dampen cyclical fluctuations. In our models, agents are assumed to prefer smooth consumption streams and would therefore be willing to insure against fluctuations in consumption. But how much they are willing to pay for such insurance depends on how undesirable the otherwise non-smooth consumption streams would be. 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. Consumers' willingness to forgo consumption to avoid cyclical fluctuations would therefore be strikingly small. Even allowing for more volatile consumption, like the pre-war U.S. or like other countries; or allowing expected utility and stochastic growth (Obstfeld, 1994), it is hard to get the costs of business cycles (and hence potential gains from managing aggregate demand through economic policy) to exceed 1–2 percent. If these costs are really this low, there would appear to be little scope for countercyclical macroeconomic policy.

However, if the cost of business cycle fluctuations is low, then compensation for bearing this risk should be low as well. Yet the risk premium—the compensation investors receive for holding risky claims—is puzzlingly high. Why then are risk premiums so high if consumption volatility has only minimal welfare costs? One possible explanation is “rare disasters”, that is, large but low probability output losses, occurring in wars, pandemics, or revolutions. This is the key insight of models that were developed to address the risk premium puzzle, such as Rietz (1988), Barro (2006), Gabaix (2008), and Wachter (2013). But if growth has such fat tails, being negatively skewed by infrequent economic catastrophes, then the welfare costs of consumption fluctuations could also be much higher than in Lucas's original calculations. Bringing these insights back to the welfare calculation using historical panel data, Barro (2009) arrives at a welfare cost estimate of 17 percent due to disaster-driven jumps, ten times larger than the 1.6 percent attributable to Gaussian growth disturbances.¹ The implied costs can be made larger still if utility is recursive, or if disasters have stochastic probability, stochastic size, more persistence, or are permanent rather than transitory (Barro and Jin, 2011; Gourio, 2012; Nakamura, Steinsson, Barro, and Ursúa, 2013; Obstfeld, 1994; Reis, 2009). The key insight is that fat tails matter, with effectively all of the amplification of the welfare costs of business cycles being driven by higher moments (Martin, 2008).²

From the “rare disaster” perspective, the true welfare cost of economic fluctuations stems from such infrequent, but very costly events. But this view might be too optimistic still, and the value of stabilization policy consequently even bigger. In this paper we argue that the “rare disasters” method of measuring the costs of business cycles needs amendment: it is not that it is incorrect, but rather that it does not go far enough.

¹For comparability, this is for the CRRA case with $\gamma = \theta = 4$ reported in Table 3 in Barro (2009).

²Research also shows that other sources of risk, notably uninsured idiosyncratic income risk, can also amplify these costs (Atkeson and Phelan, 1994; Imrohoroglu, 1989).

In fact, when we use long-run historical data to look at growth at risk, we find that many business cycles contain the skew that amplifies welfare costs. Deviations from the Gaussian benchmark do not just appear in the extreme disasters—the wars, pandemics, revolutions, etc., considered by previous research. Instead, we show that fat tails and persistence also appear in “normal recessions”—in short, there are disasters everywhere.³ Once we embrace this idea, welfare cost judgments change profoundly. In our estimates, even with low risk aversion, on average households would sacrifice about 15 percent of consumption to avoid consumption fluctuations. Far from being a side-show for macroeconomics, business cycle volatility is a first-order issue at the heart of the discipline.

The first part of our paper documents this new stylized fact using a comprehensive macro-historical database (Jordà, Schularick, and Taylor, 2017). Using local projections (LP) methods (Jordà, 2005), we present a new direct *test* for the presence of disasters. Instead of assuming disasters are present, or relying on general skewness tests, our test has the virtue of directly mapping into a tractable empirical LP model of disasters 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) as well as 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 present at all times, not just in disaster episodes.⁴ Our analysis also meshes with inherently asymmetric macro frameworks, like deLong and Summers (1988) or the Friedman “plucking” model which has periodically attracted attention (Dupraz, Nakamura, and Steinsson, 2019; Kim and Nelson, 1999). Fatás and Mihov (2013) echo this idea when noting that even in the U.S., as postwar growth volatility fell in the Great Moderation, nonetheless negative skewness increased. But, to our knowledge, we are the first to re-assay the debate over the costs of business cycles in this novel framework of ubiquitous fat tails with disasters everywhere.

The second part of the paper proposes a new empirical framework for accurately estimating and calibrating a growth process with these properties. At each date a recession occurs according to a Bernoulli coin-flip process, but this can be a less-disastrous normal recession or a more-disastrous financial crisis recession. We restrict to these cases and exclude wars to focus on peacetime welfare costs. The recession type is governed by another coin-flip. If no recession occurs, consumption growth is drawn from a Gaussian i.i.d. process. This is a baseline: in general we can allow for many recession types and non-i.i.d. growth. If a recession occurs, then over some horizon the path of consumption is subject to a stochastic penalty, the scale of which varies by horizon, and by recession type. This formulation turns out to be convenient as it maps directly into estimation via local projections (Jordà, 2005), and at larger horizons LP estimation may be preferred to VAR methods when estimation is constrained to finite-lag specifications (Jordà, Singh, and Taylor,

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

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

2020; Montiel Olea and Plagborg-Møller, 2020). 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, an essential feature in the data and a key driver of welfare consequences.⁵

The third part of the paper takes the estimated consumption growth process and simulates an economy under various parameter configurations to assess actual and counterfactual welfare losses due to peacetime business cycles.⁶ Our focus on the normal versus financial crisis recession dichotomy speaks to the increased interest in the potential gains from macroprudential policies to the extent that they can mitigate crisis risk: here direct policy actions are being debated and even implemented as we write; in contrast other disastrous recessions, such as wars or revolutions, may be less susceptible to purely economic policy interventions.

As in Barro (2009), we explore how much welfare loss is due to the Gaussian terms versus the disaster terms. Results still depend on assumptions about the permanent component of the disasters. But since disasters are now everywhere, including in normal recessions, we find that the welfare costs of business cycles are much larger. Under simulation we reject the Gaussian null in favor of our RCLP-estimated dynamic moments. 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 10 percent; and allowing fat tails with hypothetical 100% financial crisis recessions would increase it to about 25 percent.

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 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 way. The methods make for an easy and transparent mapping from the LP estimates to the calibrated simulation model.

Second, we offer a new perspective on the perennial macro question of the cost of business cycles and the importance of stabilization policy. The main new insight is that *disasters are everywhere*—in the sense that the growth process has fat tails everywhere, not just in the large Barro-type rare disasters. And though these recessions are smaller in amplitude than the rare big disasters, they are much more frequent, and we show that the welfare costs really add up to something substantial. Looking at the post-1985 era in advanced economies, our model says agents would sacrifice about 20 percent of consumption to avoid all business cycles; and about 10 percent just to avoid financial-crisis

⁵One 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.

⁶Economics 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.

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

As noted, one canonical model of per capita consumption growth used in Lucas (1987, 2003) is a trend with Gaussian shocks. Subsequently a widely-used baseline for a large literature in macroeconomics and in finance has been a random walk with drift. Both processes are symmetric.

As a first approximation, this model has served the profession well. However, as remarked in the introduction, evidence that the business cycle is asymmetric abounds. This section documents some basic moments in the data on per capita consumption (and per capita output) growth for 17 advanced economies since 1870 drawn from the macrohistorical dataset of of Jordà, Schularick, and Taylor (2017). The statistical features that we derive form the basis of the model specification that we adopt in the next section, where we calculate the welfare costs of all “disasters.”

The analysis shows non-Gaussian or fat-tailed consumption growth dynamics, or what one could term mini-disasters. This is true even outside the established rare disaster events in the benchmark chronologies of Barro (2006), Barro and Ursúa (2008), and related studies. Basically, we show that the entire universe of peacetime recessions in advanced economies since 1870 displays the same qualitative fat-tailed behavior as the Barro disasters.

We will examine two alternative measures of economic activity that we generally label as X and which refer to either real GDP per capita, Y , or real consumption per capita C . Let $x = \log(X)$, hence, lowercase variables denote logs. We use a null model that characterizes the growth rate of x (in logs) as a random walk with drift, such as

$$\Delta x_t = x_t - x_{t-1} = \gamma + \epsilon_t; \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2). \quad (1)$$

Although our data are organized as a panel, we omit cross-section index to avoid cluttered notation.

Our next step is to expand on this null model with a set of alternatives using a formal treatment of event timing definitions. First, we define the key concept of a mapping m , assumed to be chosen by the observer, which produces the classification of some time periods as *events* or “disasters.” The index of events given the observer’s choice of mapping is denoted $r_m \in \mathbb{Z}$ for $r_m = 1, \dots, R_m(T)$, that is, given a sample of size T and based on the mapping, there are at most R_m such events. When clear from context, we will omit the subindex m to simplify the presentation.

We will use the notation $\tau(r)$ to refer to the calendar time period where the r^{th} event happens. Hence $H(r) \equiv \tau(r+1) - \tau(r)$ refers to the number of calendar time periods that elapse between two consecutive events, $r-1$ and r . This is a common notation in studies where the arrival of events follows a stochastic process (see, e.g., Hamilton and Jordà, 2002).

As an illustration, consider one of the mappings we use below. In a sample where the focus is on periods of recession regardless of their origin, a convenient way to date events is with the algorithm of [Bry and Boschan \(1971\)](#). With annual data such as ours, $\tau(r_{All}) = t$ if $x_{t-1} < x_t > x_{t+1}$ with r_{All} denoting the index of all recession peaks resulting from the mapping $m = All$. Of course, [Barro \(2006\)](#), [Barro and Ursúa \(2008\)](#), and [Nakamura, Steinsson, Barro, and Ursúa \(2013\)](#) are canonical examples of an alternative rare-disaster event mapping that we also explore in more detail shortly.

We now allow for the possibility that after an event r , the economy takes a hit beyond what the Gaussian null model with random drift in [Equation 1](#) would prescribe. Hence, [Equation 1](#) can be recast by first noting that

$$\Delta x_t = \gamma + \psi_h + \epsilon_t, \quad (2)$$

for $t = \tau(r) + h$ and $h \geq 1$. Thus, ψ_h denotes the amount by which growth gets knocked from trend growth γ , h periods after an event r . It may seem natural to assume that $\psi_h \rightarrow 0$ as h grows, but in a rare disasters model such losses can be permanent, so we do not take that for granted. In sum, $\psi = (\psi_1 \dots \psi_H)$ forms a sort of impulse response. It captures, on average, the negative effects on the growth of x from experiencing a *disaster*, relative to the random walk with drift null path.

In order to characterize the cumulative effects of the disaster, we define $\Delta_h x_t \equiv x_t - x_{\tau(r)}$, where $t = \tau(r) + h$ as earlier. Note that $\Delta_h x_t = \Delta x_t + \dots + \Delta x_{\tau(r)+1}$. Thus, using [Equation 2](#) to characterize each term in this summation, it is easy to see that

$$\Delta_h x_t = \gamma h + \sum_{j=1}^h \psi_j + \sum_{j=1}^h \epsilon_{\tau(r)+j} = \gamma h + \alpha_h + u_t, \quad (3)$$

where it is clear that we define $(\psi_1 + \dots + \psi_h) = \alpha_h$; and $(\epsilon_{\tau(r)+1} + \dots + \epsilon_t) = u_t$. It is straightforward to test whether the data support departures from the null model in [Equation 1](#). Based on [Equation 3](#), one would test the null $H_0 : \alpha_1 = \dots = \alpha_H = 0$.

However, the null hypothesis can be sharpened by realizing that a given mapping m , under the null, will generate, by construction, a recession path even under that null. In other words, the mere definition provided by the mapping restricts the probability space of possible disaster paths. We illustrate precisely how each of the mappings that we entertain below generate such paths and how to properly test for the absence of event-specific distortions to the null model of [Equation 1](#).

In a sample of data, a good way to illustrate what the cumulative path of x looks like after a disaster relative to its null growth path is to examine the cumulative changes of x following disasters. In practice, for a choice of finite H , these can be estimated using [Jordà \(2005\)](#) local projections with

$$\Delta_h x_t = \gamma h + (\alpha_1 D_t^1 + \dots + \alpha_H D_t^H) + u_t, \quad (4)$$

where we define the indicator variable $D_t^h = 1$ for $h = 1, \dots, H$, and again $t = \tau(r) + h$, and where $D_t^h = 0$ otherwise. Now, *depending on the mapping m* , one obtains a sequence of values *under the*

null model of the cumulative impulse response, denoted $\alpha_h^{0,m}$, with which to compare the sample estimates from Equation 4, $\hat{\alpha}_h$, which arise from the data. Thus, when testing the null that disasters are simply rare events consistent with the null model, we will be interested in testing the hypothesis $H_0 : \alpha_h = \alpha_h^{0,m}$.

To reiterate, different mappings m truncate the calendar-time data-generating-process in Equation 1 differently. Events, or disasters, defined by a given m , refer to particular regions of the distribution of x rather than the entire sample space. For this reason, except for some mappings where the null can be constructed analytically, we rely on numerical approximations to the null that can be made as precise as needed. The next section provides a more detailed look at the dynamics of disasters in our sample using these methods.

2.1. Testing for the role of different disasters

Our statistical framework guides the testing strategy that we now follow. We begin with a baseline scenario where the theoretical null can be easily computed to illustrate how the definition of the mapping—the rule chosen to define peaks or subsets, like disaster peaks—changes the LP path under the Gaussian null. Thus, in Test 1 we consider all peaks of economic activity ($m \equiv All$) in our sample (which excludes World Wars and the Spanish Civil War), where an event is defined as: $t = \tau(r_{All})$ if $x_{t-1} < x_t > x_{t+1}$, that is, a peak in activity defined as in Bry and Boschan (1971).

Now it is quite possible that a rejection of the null in Test 1 could be solely because some of the events defined by our criterion will include Barro-Ursúa disasters. Test 2 hence begins by assessing the null against these disasters, with the mapping denoted $m \equiv BU$. We will show this is indeed the case, thus begging the question: Would we still find a rejection of the null if we excluded Barro-Ursúa disasters from our sample?

We probe this question in Test 3, with the notation $m \equiv All - BU$. We again reject the null despite excluding Barro-Ursúa disasters. But as Jordà, Schularick, and Taylor (2013); Muir (2017) have shown, recessions associated with financial crises tend to be deep and protracted. Hence one might ask if the null would be rejected if we focused only on typical recessions not associated with a Barro-Ursúa disaster nor on any remaining financial crises. This is what we do in Test 4 using the notation $m \equiv N$ for normal recessions and $m \equiv F$ for recessions associated to a financial crisis.

Summing up the preview of our results: disasters are everywhere. Economic slowdowns are not merely unfortunate draws along trend growth. They are economic events whose economic consequences need to be reevaluated in light of the results of Tests 1–4. This we do in the sections that follow. But first, here is the summary of all four tests.

Test 1: All peacetime recessions This benchmark test applies to the mapping that uses all recessions, denoted $m = All$. Every recession peak date is defined as a local maximum of x , whether $x = y$ (output peaks and output recessions) or $x = c$ (consumption peaks and consumption recessions) following Bry and Boschan (1971).

This particular mapping has a closed form solution for the null. By construction, any year after a peak will have conditional mean growth $\bar{\delta}^- = E(\Delta x | \Delta x < 0)$, whereas all other years will have unconditional mean growth $\bar{\delta} = \gamma$. Given the Gaussian null, the former is $\bar{\delta}^- = \gamma - b$, where $b = \sigma[\phi(-\gamma/\sigma)/\Phi(-\gamma/\sigma)]$ is the closed-form expression for the expected loss at the peak under Gaussianity with ϕ and Φ referring to the normal density and distribution functions respectively. Thus, under the null, the local projection Equation 4 should yield the following predicted path estimates: $\Delta_h \hat{x} = \gamma h - b D_t^1$. That is, the null predicts a dogleg jump down of size b in the level after a peak at $h = 1$, but after that growth reverts to trend as it is drawn afresh every subsequent period from the Gaussian null for all $h > 1$.

To confirm that this is correct, Figure 1 and Table 1 show numerical null paths exactly that exactly match the closed-form solution just discussed. For clarity, here and below, all LP paths are in deviations relative to the mean growth path γh . The sample is all countries ($N = 17$), all peacetime years (1870–2008) from the Jordà, Schularick, and Taylor (2017) dataset, excluding the two World Wars and the Spanish Civil War. For log real GDP per capita, $\gamma = 1.96\%$, $\sigma = 3.62\%$, and the implied jump loss under the null is $b = -4.24\%$. For log real consumption per capita, $\gamma = 1.78\%$, $\sigma = 4.02\%$, and the implied null has $b = -4.42\%$. Both match the simulation LP output.

We now use historical data to test whether reality accords with the null, and it is easy to see that the null is generally too high and easily rejected.⁷ In the figures, the 95% confidence intervals are shown for the LP. For GDP the difference between the LP and the null paths is negative and statistically significant when $h \geq 2$, and for consumption when $h = 1$ and $h \geq 4$. The behavior is as one might have expected. Including all recessions in the sample leads in general to post-recession paths skewed more to the downside than a Gaussian null would predict, and more in line with the disaster model.


Test 2: All peacetime Barro-Ursúa disasters This benchmark test applies to the mapping used by an investigator who employs the Barro and Ursúa (2008) rare disaster criterion, which we denote $m = BU$. In BU, the mapping encompasses disaster events defined as a peak-to-trough drop in annual real GDP per capita of at least 15%, resulting in a set of disaster events with frequency $p = 2.4\%$ in the peacetime sample considered above. We construct the Gaussian null numerically by imposing the same conceptual frequency cutoff.⁸

That is, we take all simulated peak-to-trough declines, and find the worst $X\%$ among them by peak-to-trough ranking (as in BU), and then designate these as the rare disaster dates, choosing X such that their empirical frequency among all years is 2.4%. Specifically, in the simulated Gaussian null data we find that 12% of all observations are recession peaks, and so we choose $X = 20\%$, to generate the desired 2.4% rare disaster frequency.

⁷Rejection in one period is enough to reject the joint null, hence the joint null is not reported even if the size of the implicit test is too small. Evidence from the data is overwhelming.

⁸That is, we do not use a calibration based on the drop cutoff, as it would yield hardly any disasters under the Gaussian null, which is another way of stating the problem with the Gaussian null, as will become apparent.

Figure 1: LP test for non-Gaussianity, all JST recession peaks, all countries, peacetime sample



Notes: See text. The null is a Gaussian process. The LP path and 95% confidence interval are from the data and disaster mapping used.

Table 1: LP test for non-Gaussianity, all JST recession peaks, all countries, peacetime sample

(a) Real GDP per capita, log y ($\times 100$)										
Horizon h	1	2	3	4	5	6	7	8	9	10
$\alpha_h^{0,All}$	-4.24	-4.25	-4.25	-4.25	-4.25	-4.25	-4.25	-4.25	-4.25	-4.25
$\hat{\alpha}_h^{All}$	-4.40	-5.13	-5.45	-6.26	-7.15	-7.65	-8.20	-8.62	-8.56	-8.67
Difference	-0.16	-0.88***	-1.21***	-2.01***	-2.90***	-3.40***	-3.95***	-4.37***	-4.32***	-4.43***
s.e.	0.17	0.30	0.40	0.50	0.60	0.70	0.78	0.86	0.93	0.99
p -value	0.37	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
N	2237	2220	2203	2186	2169	2152	2135	2118	2101	2084

(b) Real consumption per capita, log c ($\times 100$)										
Horizon h	1	2	3	4	5	6	7	8	9	10
$\alpha_h^{0,All}$	-4.42	-4.42	-4.43	-4.42	-4.42	-4.43	-4.43	-4.42	-4.42	-4.43
$\hat{\alpha}_h^{All}$	-3.39	-4.40	-5.31	-6.26	-7.61	-8.08	-8.29	-8.88	-9.04	-9.91
Difference	1.02***	0.02	-0.88**	-1.83***	-3.19***	-3.65***	-3.87***	-4.46***	-4.62***	-5.49***
s.e.	0.22	0.33	0.44	0.55	0.66	0.76	0.84	0.91	0.99	1.06
p -value	0.00	0.94	0.04	0.00	0.00	0.00	0.00	0.00	0.00	0.00
N	2150	2133	2116	2099	2082	2065	2048	2031	2014	1997

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. See text. Units are in log times 100.

The numerical null values are again negative, but they do not follow a simple dogleg: selection into the disaster bin is not a simple mapping from the one-year post period growth rate, and depends on the peak-to-trough change over a longer, endogenous, horizon.

The strong message from [Figure 2](#) and [Table 2](#) is that [Barro \(2006\)](#) and [Barro and Ursúa \(2008\)](#) were right. At all horizons, the actual historical LP estimates are significantly below the simulated null values, and by a huge margin, often more than minus 10% in level difference. The null deviations are, of course, negative: after all the endogenous dating scheme was designed to capture peaks with the worst quintile subsequent declines. But they are just not as negative as in the actual data. The growth process of the Gaussian null is clearly inadequate to capture the fat tails seen during rare disasters. If we constrain to match empirical consumption volatility, the worst 20% of recessions as simply not as bad as the empirical rare disasters in the observed left tail.

The persistent negative loss of output lasts for many periods and the peak deviation from the Gaussian null path is -14.69% in year 7, with a deviation of -9.43% still present after 10 years. For consumption the peak deviation is -14.25% in year 7, with a deviation of -6.09% after 10 years. This supports the standard view that disaster losses consist, in some large part, of permanent downshifts in the trend path. There is also support for the mutli-period evolving disaster setup in [Nakamura, Steinsson, Barro, and Ursúa \(2013\)](#), since the loss doesn't fully appear in year 1 and slowly builds up, before some transitory part begins to fade away.

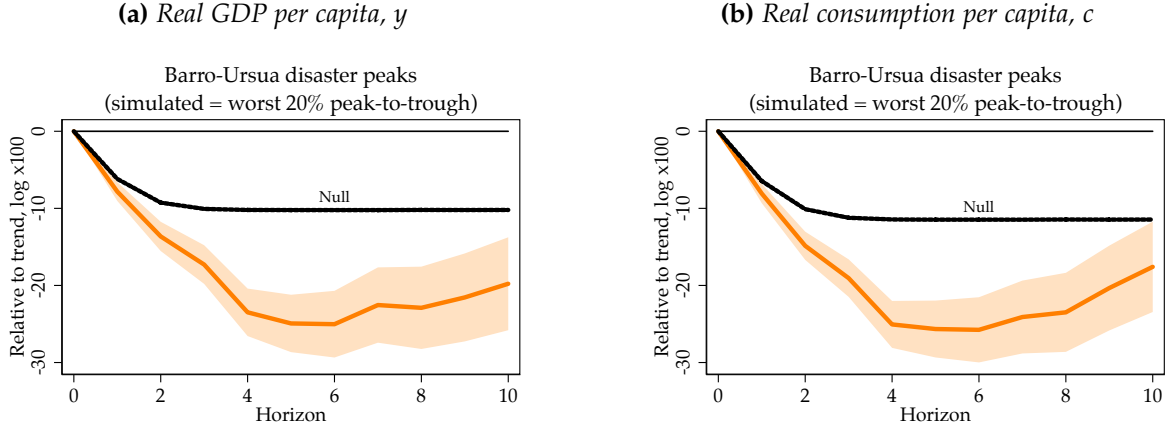
Test 3: All peacetime recessions minus Barro-Ursúa disasters This test applies to the mapping used by an investigator who employs the recession peak criterion used in [Test 1](#), but excludes all of the Barro-Ursúa disaster peaks used in [Test 2](#). We refer to this mapping for the date indicators as $m = All - BU$.

The equally strong message from [Figure 3](#) and [Table 3](#) is that disasters are indeed everywhere, so to speak, and not just in Barro-Ursúa episodes. Every peak in the simulated data is counted in the simulated date indicators, except those simulated as BU disaster peaks above. And every peak in the actual data is counted in the real world date indicators, except BU disaster peaks (for robustness we actually exclude peaks within ± 2 years of a BU disaster peak).

The numerical null values are again negative: all peaks on average have subsequent declines. The null values are also lower than in [Table 2](#), and higher than in [Table 1](#), by construction: we excluded rare disasters in the numerical calculations to focus on other recession events. The historical LP estimates are also lower than in [Table 2](#), and higher than in [Table 1](#), by construction: we excluded the actual rare disasters to focus on other recession events. The null is again no longer a simple dogleg: recoveries from bad shocks that avoid the disaster threshold survive in this sample, but not those that cross the threshold; on average, then, the path bounces up compared to the full sample.

Now the main result appears, as we clearly see, since the very same type evidence of non-Gaussianity emerges. At almost all horizons, the actual LP estimates are again well below the simulated values and the differences are statistically significant. However we should note that, given the exclusion of the worst disaster events, the differences—like the estimates themselves—are

Figure 2: LP test for non-Gaussianity, all BU disaster peaks, all countries, peacetime sample



Notes: See text. The null is a Gaussian process. The LP path and 95% confidence interval are from the data and disaster mapping used.

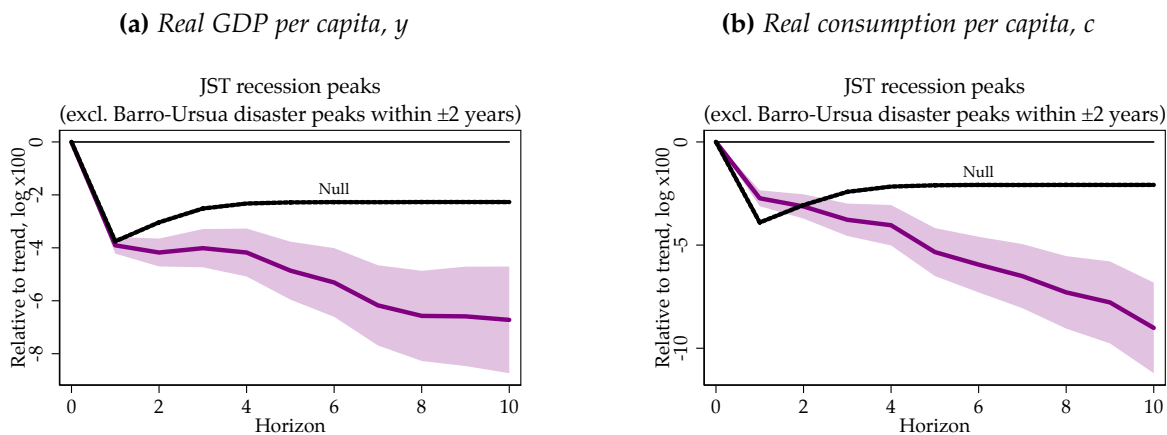
Table 2: LP test for non-Gaussianity, all BU disaster peaks, all countries, peacetime sample

(a) Real GDP per capita, log y ($\times 100$)										
Horizon h	1	2	3	4	5	6	7	8	9	10
$\alpha_h^{0,BU}$	-6.19	-9.33	-10.16	-10.31	-10.34	-10.34	-10.34	-10.34	-10.34	-10.35
$\hat{\alpha}_h^{BU}$	-7.82	-13.64	-17.29	-23.49	-24.93	-25.03	-22.54	-22.90	-21.55	-19.78
Difference	-1.63***	-4.31***	-7.13***	-13.17***	-14.59***	-14.69***	-12.20***	-12.56***	-11.22***	-9.43***
s.e.	0.62	0.97	1.28	1.57	1.90	2.21	2.49	2.72	2.91	3.07
p -value	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
N	2237	2220	2203	2186	2169	2152	2135	2118	2101	2084

(b) Real consumption per capita, log c ($\times 100$)										
Horizon h	1	2	3	4	5	6	7	8	9	10
$\alpha_h^{0,BU}$	-6.45	-10.10	-11.21	-11.44	-11.47	-11.49	-11.49	-11.49	-11.48	-11.48
$\hat{\alpha}_h^{BU}$	-8.08	-14.85	-19.04	-25.05	-25.64	-25.75	-24.08	-23.47	-20.32	-17.57
Difference	-1.63***	-4.75***	-7.83***	-13.60***	-14.17***	-14.25***	-12.59***	-11.98***	-8.83***	-6.09**
s.e.	0.64	0.93	1.24	1.55	1.88	2.16	2.41	2.62	2.81	2.99
p -value	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.04
N	2150	2133	2116	2099	2082	2065	2048	2031	2014	1997

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. See text. Units are in log times 100.

Figure 3: LP test for non-Gaussianity, all JST recession peaks, all countries, peacetime sample, excluding BU disaster peaks (± 2 years) from the historical and simulated indicators



Notes: See text. The null is a Gaussian process. The LP path and 95% confidence interval are from the data and disaster mapping used.

Table 3: LP test for non-Gaussianity, all JST recession peaks, all countries, peacetime sample, excluding BU disaster peaks (± 2 years) from the historical and simulated indicators

(a) Real GDP per capita, $\log y$ ($\times 100$)										
Horizon h	1	2	3	4	5	6	7	8	9	10
$\alpha_h^{0, All-BU}$	-3.77	-3.06	-2.55	-2.36	-2.32	-2.31	-2.31	-2.31	-2.31	-2.31
$\hat{\alpha}_h^{All-BU}$	-3.90	-4.18	-4.01	-4.18	-4.86	-5.31	-6.17	-6.57	-6.59	-6.72
Difference	-0.14	-1.12***	-1.46***	-1.82***	-2.55***	-3.00***	-3.86***	-4.26***	-4.28***	-4.42***
s.e.	0.16	0.27	0.37	0.46	0.56	0.66	0.77	0.87	0.96	1.03
p -value	0.39	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
N	2048	2031	2014	1997	1980	1963	1946	1929	1912	1895

(b) Real consumption per capita, $\log c$ ($\times 100$)										
Horizon h	1	2	3	4	5	6	7	8	9	10
$\alpha_h^{0, All-BU}$	-3.90	-3.05	-2.41	-2.15	-2.09	-2.08	-2.07	-2.07	-2.07	-2.07
$\hat{\alpha}_h^{All-BU}$	-2.73	-3.12	-3.77	-4.03	-5.33	-5.93	-6.49	-7.27	-7.76	-8.99
Difference	1.18***	-0.07	-1.36***	-1.88***	-3.24***	-3.86***	-4.42***	-5.20***	-5.69***	-6.92***
s.e.	0.20	0.30	0.40	0.50	0.59	0.69	0.79	0.90	1.01	1.12
p -value	0.00	0.81	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
N	1956	1939	1922	1905	1888	1871	1854	1837	1820	1803

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. See text. Units are in log times 100.

smaller in scale than in [Table 2](#). The deviation from the null prediction is now -4.42% at 10 years for GDP; -6.92% at 10 years for consumption.

We therefore conclude that the rare disaster approach is valid: those events are still the big disasters. But now even all the other recessions in history can be viewed as mini-disasters of the same ilk. That is, there are disasters everywhere.

Test 4: All peacetime normal and financial recessions minus Barro-Ursúa disasters This test applies to the mapping used by an investigator who employs same the recession peak criterion used in Test 3, minus the BU peaks, but the investigator allows different responses for Normal and Financial Crisis recessions. The mapping is the same as $m = All - BU$, but we denote the different responses by labels N and F , respectively.

This is a natural hypothesis to account for the previous result by allowing a role for financial crises in generating these patterns. It is well established that recessions associated with financial crises are more costly on average than other “normal recessions” ([Jordà, Schularick, and Taylor, 2013](#); [Muir, 2017](#)), and this is true even outside of the rare disasters sample. To attack this question we employ the [Jordà, Schularick, and Taylor \(2013\)](#) classification of recessions into N and F types, where the latter means a financial crisis event occurs within ± 2 years of the peak. Now we just repeat the previous test, but examine differences between the null and historical paths in N and F types of recessions separately. The local projection estimates are now denoted by N and F subscripts.

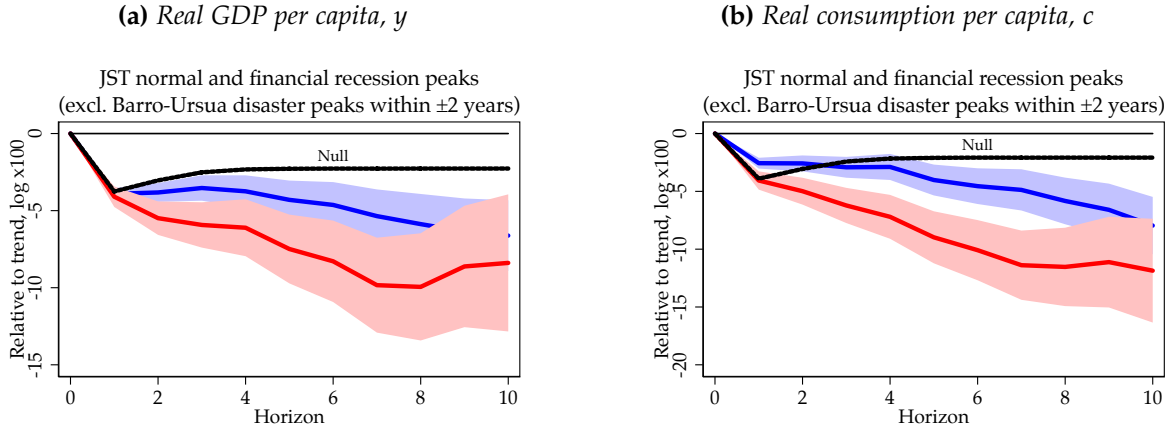
[Figure 4](#) and [Table 4](#) deliver a potentially surprising answer. Financial Crisis type recessions have much worse growth paths than the null, and a good deal of the fat-left-tail growth performance seen in the last test clearly stems from the damage done in financial crisis events. However, even in the Normal recessions we still find strong evidence against the Gaussian null. So the problem is not just rare disaster causing fat tails, nor is it rare disasters and financial crises. It looks like even Normal recessions contribute to the general pattern too.⁹

Implications for welfare: do all disasters matter? If we admit these fat-tailed mini-disasters into a welfare accounting framework, there is prima facie evidence of their quantitative significance. The intuition is clear for one-period disasters. It is well known that the welfare losses, risk premia, and other features of disaster calculus depend—to a first approximation—on expressions of the form $pE(b)$, which is the event probability p times the expected one-shot jump loss $E(b)$. For sure, mini-disasters are on average less painful, about one quarter as severe as rare disasters. However, they are about four times as likely to occur. Overall, then, they are potentially very consequential in welfare terms, and possibly even more than the rare disasters.

Some back-of-the envelope calculations can motivate the rest of the paper. Specifically, in peacetime, using the JST dataset analyzed above—not exactly the same as the BU dataset—the BU

⁹Note that the null is numerically constructed as in Test 3, so by construction includes “big” recessions that would probably better characterize financial crises. Instead of adjusting the null, we left it as is to show that even with this adverse choice of null, Normal recessions still exhibit statistically notable fat-tailed features.

Figure 4: LP test for non-Gaussianity, all JST Normal and Financial recession peaks, all countries, peacetime sample, excluding BU disaster peaks (± 2 years) from the historical and simulated indicators



Notes: See text. The null is a Gaussian process. The LP path and 95% confidence interval are from the data and disaster mapping used.

Table 4: LP test for non-Gaussianity, all JST Normal and Financial recession peaks, all countries, peacetime sample, excluding BU disaster peaks (± 2 years) from the historical and simulated indicators

(a) Real GDP per capita, $\log y$ ($\times 100$)										
Horizon h	1	2	3	4	5	6	7	8	9	10
$\alpha_{h,N}^{0,N}$	-3.77	-3.06	-2.55	-2.36	-2.32	-2.31	-2.31	-2.31	-2.31	-2.31
$\hat{\alpha}_{h,N}^N$	-0.14	-0.76	-0.98	-1.39	-1.99	-2.32	-3.04	-3.57	-4.05	-4.31
$\hat{\alpha}_{h,N}^F$	-0.33	-2.43	-3.38	-3.75	-5.17	-5.97	-7.52	-7.64	-6.31	-6.08
Difference N	-0.14	-0.76**	-0.98**	-1.39***	-1.99***	-2.32***	-3.04***	-3.57***	-4.05***	-4.31***
s.e.	0.19	0.31	0.42	0.53	0.64	0.76	0.88	1.00	1.09	1.17
p -value	0.45	0.01	0.02	0.01	0.00	0.00	0.00	0.00	0.00	0.00
Difference F	-0.33	-2.43***	-3.38***	-3.75***	-5.17***	-5.97***	-7.52***	-7.64***	-6.31***	-6.08***
s.e.	0.34	0.55	0.75	0.94	1.14	1.35	1.57	1.77	2.01	2.27
p -value	0.34	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
N	2077	2060	2043	2026	2009	1992	1975	1958	1941	1924

(b) Real consumption per capita, $\log c$ ($\times 100$)										
Horizon h	1	2	3	4	5	6	7	8	9	10
$\alpha_{h,N}^{0,N}$	-3.90	-3.05	-2.41	-2.15	-2.09	-2.08	-2.07	-2.07	-2.07	-2.07
$\hat{\alpha}_{h,N}^N$	1.34	0.47	-0.50	-0.73	-1.92	-2.46	-2.79	-3.74	-4.51	-5.87
$\hat{\alpha}_{h,N}^F$	-0.16	-1.93	-3.80	-5.04	-6.88	-8.00	-9.30	-9.44	-9.03	-9.77
Difference N	1.34***	0.47	-0.50	-0.73	-1.92***	-2.46***	-2.79***	-3.74***	-4.51***	-5.87***
s.e.	0.24	0.35	0.46	0.57	0.68	0.79	0.91	1.03	1.16	1.26
p -value	0.00	0.18	0.27	0.20	0.00	0.00	0.00	0.00	0.00	0.00
Difference F	-0.16	-1.93***	-3.80***	-5.04***	-6.88***	-8.00***	-9.30***	-9.44***	-9.03***	-9.77***
s.e.	0.40	0.59	0.78	0.97	1.15	1.33	1.53	1.73	2.00	2.29
p -value	0.70	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
N	1974	1957	1940	1923	1906	1889	1872	1855	1838	1821

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. See text. Units are in log times 100.

rare disasters have a frequency $p = 0.016$; then, from the LP charts, we might calibrate $E(b) = 16\%$, which would be larger if wars were in our sample, as in BU. Thus, for the BU “rare” disasters the expected loss is $pE(b) = 0.26\%$.

What about our proposed “non-rare” mini-disasters using the JST recession classification? Our Financial Crisis recessions above have $p = 0.036$ and from the charts we calibrate $E(b) = 8\%$, for an expected loss of $pE(b) = 0.29\%$. And for Normal recessions we get $p = 0.106$ and from the charts $E(b) = 4\%$, for an expected loss of $pE(b) = 0.42\%$. Adding these last two expected losses we get 0.71% for the total expected disaster losses of the “non-rare” or non-BU type. Thus the expected jump loss from recessions *other than* rare disasters (0.71%) has a much larger magnitude than that associated with the rare disasters themselves (0.26%).

Of course, these are first approximations, and neglect convex utility. Welfare losses from fat-tailed disasters, rare or frequent, will be amplified under risk aversion: an agent prefers two small losses to a single loss of twice the size, so to speak. So further analysis is needed, and the point of the rest of this paper is to make the correct formal calculations. But the intuition holds up and leads to the key finding of our paper. Rare disasters hurt a lot, but happen once in a blue moon; but the smaller-but-frequent left-tail pain of other recessions really adds up. To the best of our knowledge, no research has considered the existence of the latter type of non-Gaussianity, nor measured the relative importance of these different flavors of disasters in the context of the literature on the welfare costs of business cycles. Setting aside the rough motivation in the last paragraph, we now need a carefully estimated growth process, a well calibrated model, and a set of meaningful counterfactuals. To this we now turn.

3. DISASTER PATHS WITH RANDOMLY VARYING SEVERITY

The previous section establishes that, on average, even plain vanilla recessions are like mini-disasters. However, it does not quite tell us how much consumption the representative agent would forgo to avoid them, a critical calculation to assess welfare costs. Motivated by the methods used above, we introduce an estimation technique that will allow us to calibrate disaster dynamics directly from local projections. Later we incorporate these estimates to conduct the proper welfare calculation.

We begin by conceiving of the economy as evolving under two different regimes. In normal times the economy evolves as the random walk with drift of [Equation 1](#). When a disaster event hits, the economy is shunted onto a different regime whose dynamics we wish to characterize next.

Formally, consider the empirical process corresponding to all calendar dates associated with an event r , that is $\{\tau(r_m)\}_{r_m=1}^{R_m}$. This is a *point process* with a long tradition in statistics (see, e.g., [Cox and Isham, 1980](#), for an introduction to the topic).

In particular, if we assume—as we will assume in our simulations later—that at each calendar time period a Bernoulli draw with constant success probability determines whether an event takes place, then the duration elapsed between events, $H(r) = \tau(r) - \tau(r - 1)$, is approximately distributed as an exponential random variable whose mean is the inverse of the Bernoulli probability

of success. Also note that we will in practice usually refer to these events as peaks (i.e., of the cycle).

Applications of point processes in economics include models describing the arrival of policy rate changes (Hamilton and Jordà, 2002; Hu and Phillips, 2004), descriptions of the arrival of trades in financial markets (Bowsher, 2007; Engle and Lunde, 2003; Engle and Russell, 1998), and models of trades in commodity markets (Davis and Hamilton, 2004), among others.

In order to potentially allow the severity of each disaster to vary randomly, we expand on our previous specification. Therefore, at each event date $\tau(r)$, we assume that there is an associated stochastic process characterizing the depth of the recession associated to the r^{th} event. We denote the path at each event r as $s(r) = (s_1(r), \dots, s_{H(r)}(r))$ (we omit reference to the mapping for simplicity). The reason we use the notation s_h rather than ψ_h as before is to differentiate estimates based on our random coefficients model from the fixed coefficient models discussed earlier. This will become clear momentarily.

Before becoming more specific, it is important to note that the stochastic process given by $\{\tau(r), s(r)\}_{r=1}^R$ for $r \in \mathbb{Z}$ now becomes a *marked point process* (see Cox, 1972; Last and Brandt, 1995, for two classical references). That means that the intensity or *hazard* rate with which events arrive and the values that the marks take could be related, in principle.

Our strategy is more modest motivated by the evidence and to facilitate the simulations in later sections. First, we assume that there is no duration dependence in the point process, that is, the duration between two disasters is uninformative about when the next disaster will hit, or as stated earlier, $H(r) \sim \mathcal{E}(\lambda)$, that is, durations are exponentially distributed with constant hazard $\lambda = 1/p$ where p is the success probability in the Bernoulli trial. This seems consistent with findings reported in the literature on recession prediction (see, e.g., Berge and Jordà, 2011; Chauvet and Hamilton, 2006; Diebold and Rudebusch, 1990). Second, the point process and the process for the marks are assumed to be independent from each other. In other words, the severity a disaster is not a function of how long ago the previous disaster happened. Here again the evidence does not strongly suggest otherwise. Potential violations of these two assumptions in other contexts are an interesting area of research that we reserve for a different paper.

A simple way to extend the framework of marked point processes is to build out from our earlier local projections specification in Equation 2 as follows:

$$\Delta x_t = g + \left[s_1 D_t^1 + \dots + s_H D_t^H \right] \exp(\zeta_{\tau(r)}) + \epsilon_t; \quad \zeta_{\tau(r)} \sim \mathcal{N}(0, \sigma_\zeta^2); \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2); \quad (5)$$

$$t = \tau(r) + h; \quad r = 1, \dots, R; \quad h = 1, \dots, H(r); \quad H = \max\{H(r)\}_{r=1}^R,$$

where ζ and ϵ are independent from each other and the D_t^h are the indicator variables defined in the previous section. In Equation 5 the simplifying assumption that we make is that all impulse response coefficients s_h are shifted up or down by the same amount within event, but by different amounts across events. Note that $\mathbb{E}(\exp(\zeta)) = \exp(\sigma_\zeta^2/2)$ and hence, evaluated at the mean, the average impulse response is $s_h \exp(\sigma_\zeta^2/2) = \psi_h$, which clarifies why we switched to the notation s_h . Also note that we present the model by omitting the panel dimension of our application to avoid

notational clutter. Finally, it is important to understand that this estimation only applies at horizons within a specific cycle, given the last restriction on h .

We refer to this model the *random coefficient local projections* or RCLP model. To get a tangible sense of what the RCLP specification buys us, notice that $\exp(\zeta) \in (0, \infty)$ since $\zeta \in (-\infty, \infty)$. So, as $\zeta \rightarrow -\infty$, the disaster approaches the Gaussian null model and shown in Equation 1. This is why we do not scale the constant term by $\exp(\zeta)$ as we do with the other coefficients. As $\zeta \rightarrow +\infty$, the disaster becomes increasingly severe.

As we will show momentarily, we estimate $\sigma_\zeta \approx 0.375$. This means that the recession penalty coefficients, s_h are scaled up or down as follows. Take a centered 95% probability range for ζ . At the low end, the coefficient s_h is scaled down by a factor of 0.48, approximately. At the other end, it gets scaled by a factor of approximately 2.1. In other words, the worst recession is about 4.3 times worse than the mildest when comparing the 2.5% quantile to the 97.5% quantile of the distribution of the latent variable ζ .

As an example, if the typical recession results in a 2% decline in GDP in the first year, the mild recession above will see a decline of about 1%, whereas the severe recession will see a decline on the order of 4%. Thus, this simple extension of the model permits quite a range of variation in the severity of disaster events, which can easily bracket the mild recession of 2001 and the Great Recession in the U.S., for example. Note also that $e^{0.375^2/2} = 1.07$ so the median responses s_h should be inflated by a factor 1.07 to compute mean responses.

The model in Equation 5 can be extended in a number of ways that we leave unexplored here as those extensions are not our primary focus in this paper. In more general applications, a typical local projection exercise would include additional controls. However, it is easy to see how these extensions could be accommodated. The specification in Equation 5 can be estimated by maximum likelihood for each horizon $h = 1, \dots, H$ noting that because we work in event time, ζ_t is common to all horizons and the system of equations can be set up accordingly. Our specific environment allows for a simpler approach that relies on single equation maximum likelihood estimation.

Of course, the counterpart to Equation 3 using a similar RCLP approach as in Equation 5 is:

$$\Delta_h x_t = gh + \left[a_1 D_t^1 + \dots + a_H D_t^H \right] \exp(\zeta_{i,\tau(r)}) + u_t, \quad (6)$$

where $a_h \exp(\sigma_\zeta^2/2) = \alpha_h$, $a_h = (s_1 + \dots + s_h)$, and $u_t = (\epsilon_t + \dots + \epsilon_{\tau(r)+1})$ since $\tau(r) + h = t$. Hence, note that the variance of u_t is $h\sigma_\epsilon$, that is, the residual u_t is mechanically heteroskedastic.

3.1. Estimating the baseline RCLP model

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 c_{it} denote the log of annual real consumption per capita across years $t = 1, \dots, T$ and countries $i = 1, \dots, I$. The relevant sample will be drawn from the

same peacetime dataset as above, with recession peaks, r , further sorted into two types: recessions associated with financial crises (F) and normal recessions (N). In due course, the occurrence of these peaks will be modeled as Bernoulli draws, as in standard disaster models. We set a maximal horizon for estimation $H_{max} = 10$ 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.

Estimation proceeds by adapting the specification in Equation 5 for our panel of consumption data. Hence, we specify the LHS variable of the LP in similar fashion as $\Delta c_{i,t} = c_{i,t} - c_{i,t-1}$. Specifically, we estimate the following expression over every cycle:

$$\Delta c_{i,t} = g + \left[s_1^n N_{i,t}^1 + \dots + s_H^n N_{i,t}^H + s_1^f F_{i,t}^1 + \dots + s_H^f F_{i,t}^H \right] \exp(\zeta_{i,\tau(r)}) + \epsilon_{i,t}, \quad (7)$$

where the sample is for all i , all r , and all t such that $\tau(r) < t \leq \tau(r+1)$.

Here, g is (conditional) trend drift in growth form. For each i and r , and $h = 1, \dots, \max(H(r), H_{max})$, and $t = \tau(r) + h$, we then define the indicator variable $N_{i,t}^h = 1$ if r corresponds to a normal recession, and $t \leq \tau(r+1)$, 0 otherwise. $F_{i,t}^h = 1$ is defined similarly for r corresponding to a financial recession. Note that once a given cycle r ends, the indicators for that cycle do not carry forward. We will examine the importance of this choice below.

Here, the s_h^n and s_h^f measure disaster-type growth losses within a cycle. In particular, s_h^f measures the growth impact of a financial crisis recession at h -years, relative to trend. The interpretation of the coefficient s_h^n is analogous for normal recessions. Both are modulated by $e^{\zeta_{i,\tau(r)}}$ to allow for potentially random variation in the severity of the episode, as explained in the previous section.

Next, as noted, we cannot achieve reasonable sample sizes of outcomes at large H . Thus, pragmatically, we truncate LP estimations at a reasonable horizon H_{max} . 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^n and s_h^f parameters will cease at that point, and no further growth losses are incurred: the disaster is over.

Finally, to avoid multicollinearity with the drift g , we impose the restriction $s_1^n + \dots + s_{10}^n = 0$, that is, we assume that normal recessions do not have long-run permanent level effects on consumption growth. We view this as a conservative assumption. Of course, we expect, and will obtain, estimates $s_1^f + \dots + s_{10}^f < 0$, in line with the consensus that financial recession do have long-run permanent level effects. The hypothesis that a financial crisis recession leads to a permanent loss has become the consensus view since at least the work of Cerra and Saxena (2008). Assuming that a normal recession also leads to a permanent loss, however, might be controversial, and is avoided.¹⁰

¹⁰That is, this type of constraint effectively adjusts the trend drift component and therefore also rescales all financial recession coefficients. In other words, if financial recessions had the same long-run impact as normal recessions, the financial recession coefficients would sum up to zero as well. We will discover that this is clearly not the case, in line with consensus. Of course, all else equal, these restrictions limit the welfare losses in the normal recessions, and lower estimates of welfare costs, and in that sense are conservative.

3.2. Results

We now show how to estimate the parameters of the data generating process for consumption growth given by estimating equation [Equation 7](#). Recall that to get to this point, the first results in this paper have established that recession peaks are associated with subsequent non-Gaussianity in the process for a considerable time, and that the occurrence of persistent left-tail is not just a feature of the so-called rare disaster events but is also seen in peacetime financial crisis and even normal recessions. We argue that our estimating equation [Equation 7](#) is a flexible yet parsimonious way to capture the evolving growth penalty that comes in the wake of our broader class of mini disasters, and that the results are robust. We process in steps, building up to our preferred specification.

Fixed coefficients A first step is to present an estimate of [Equation 7](#) but restricted to fixed coefficients rather than random coefficients. That is, we drop the random scaling terms $\exp(\zeta_{i\tau(r)})$. This leaves us with a standard LP estimation which can be performed using OLS methods:

$$\Delta c_{i,t} = g + \left[s_1^n N_{i,t}^1 + \dots + s_H^n N_{i,t}^H + s_1^f F_{i,t}^1 + \dots + s_H^f F_{i,t}^H \right] + \epsilon_{i,t}. \quad (8)$$

Note here that the $N_{i,t}^h$ and $F_{i,t}^h$ indicators pick up growth effects when an event occurred at time $\tau(r) = t - h$ within the same cycle r . As stated earlier, we impose the restriction $s_1^n + \dots + s_{10}^n = 0$.

That is, we start off by estimating a simpler data generating process where the disasters are uniform rather than variable. All else equal, this process of course would be expected to yield smaller welfare losses—since a source of volatility has been suppressed—for the same reasons as in [Barro \(2006\)](#), and we can indeed confirm this intuition later.

Our estimate of the fixed coefficient local projection (FCLP) model in [Equation 8](#) is shown in [Table 5](#), column (1) and [Figure 5](#), panel (a). In the figure, the cumulated estimated \hat{a}^n and \hat{a}^f coefficients are shown where blue denotes normal and red denotes financial, $a_h^n = (s_1^n + \dots + s_h^n)$ and $a_h^f = (s_1^f + \dots + s_h^f)$. We also show 95% confidence intervals. For reference, the pure conditional drift path gh is also shown. We see that normal recessions last about one year. By year two the economy regains the peak level, and continues to grow. Financial recessions are deeper, longer, and even 10 years out, sit about 10 percentage points below the normal recession path.

These preliminary results are not so new, but serve as a sense check. They reassuringly conform to previous research using local projection methods. As is well documented in large samples, (e.g. [Jordà, Schularick, and Taylor, 2013, 2017](#); [Reinhart and Rogoff, 2014](#)), there are significant differences between normal and financial crisis recessions. And the local projections in [Jordà et al. \(e.g. 2013\)](#) are similar to these paths expressed as deviations from trend.

As a second step we pause to ask whether the data structure used in [Equation 8](#) is innocent. Note that these LPs are estimated only using *within* cycle observed paths for each cycle r : the data traces out a path for the normal and financial subsets starting at time $t = \tau(r)$, and the s_h coefficients in [Equation 8](#) simply recover the unconditional mean change from there up to time $t = \tau(r) + h$.


Table 5: Fixed and Random Coefficient Local Projections

	(1) FCLP	(2) FCLP (non-truncated)	(3) RCLP	(4) RCLP (GBF)
s_1^n	-5.17*** (0.23)	-4.58*** (0.24)	-5.49*** (0.26)	-5.23*** (0.24)
s_2^n	-0.76*** (0.23)	-0.28 (0.24)	-0.74*** (0.21)	-0.96*** (0.15)
s_3^n	0.02 (0.26)	0.20 (0.24)	0.08 (0.22)	-0.20 (0.13)
s_4^n	0.40 (0.29)	0.50** (0.25)	0.46* (0.25)	0.63*** (0.16)
s_5^n	1.10*** (0.33)	0.69*** (0.25)	1.45*** (0.30)	1.23*** (0.16)
s_6^n	1.00*** (0.37)	0.67*** (0.25)	1.01*** (0.36)	1.44*** (0.12)
s_7^n	0.60 (0.41)	1.03*** (0.25)	0.67* (0.40)	1.29*** (0.11)
s_8^n	1.26*** (0.45)	0.80*** (0.26)	1.27*** (0.43)	0.94*** (0.15)
s_9^n	0.99** (0.49)	0.90*** (0.26)	0.88* (0.47)	0.57*** (0.15)
s_{10}^n	0.56 (0.53)	0.07 (0.26)	0.41 (0.55)	0.30*** (0.12)
s_1^f	-5.37*** (0.36)	-5.24*** (0.37)	-5.48*** (0.40)	-6.20*** (0.39)
s_2^f	-2.12*** (0.36)	-2.03*** (0.37)	-2.55*** (0.35)	-1.75*** (0.21)
s_3^f	-1.71*** (0.41)	-1.42*** (0.38)	-1.73*** (0.37)	-1.38*** (0.21)
s_4^f	-1.10** (0.47)	-0.15 (0.38)	-1.46*** (0.41)	-0.96*** (0.24)
s_5^f	-0.31 (0.57)	-0.07 (0.39)	0.14 (0.46)	-0.67** (0.29)
s_6^f	-0.34 (0.67)	0.24 (0.39)	-0.36 (0.60)	-0.53 (0.33)
s_7^f	-0.11 (0.79)	0.00 (0.41)	0.11 (0.71)	-0.43 (0.39)
s_8^f	-0.33 (0.89)	-0.31 (0.42)	-0.31 (0.78)	-0.17 (0.48)
s_9^f	-0.13 (1.04)	-0.01 (0.44)	-0.14 (0.98)	0.32 (0.59)
s_{10}^f	-0.22 (1.22)	-0.44 (0.48)	-0.26 (1.15)	0.96 (0.66)
g	2.80*** (0.10)	2.38*** (0.09)	2.88*** (0.09)	2.86*** (0.09)
Unconditional drift, actual	1.89	1.89	1.89	1.89
Unconditional drift, fitted	1.89	1.89	1.89	1.89
σ_ξ^2	11.71	11.81	10.84	10.99
σ_ζ^2			0.14 (0.02)	0.14 (0.02)
R^2	0.26	0.22	0.35	0.34
RMSE	3.42	3.51	3.21	3.24
Observations	2143	2143	2143	2143


Notes: Dependent variable: $\Delta c_{i,t} \times 100$. The table displays regression coefficients for the FCLP and RCLP models with $H=10$. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. We impose $s_1^n + \dots + s_{10}^n = 0$. The row σ_ζ^2 displays the estimate of the variance of the latent variable ζ in the RCLP specifications. See text.

Figure 5: Recession paths for Normal and Financial Crisis Recessions: Fixed Coefficient Local Projections

(a) FCLP estimates



(b) FCLP estimates (non-truncated)



Notes: The figure shows Normal versus Financial Crisis recession paths based on Fixed Coefficient Local Projections. 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. Note that coefficients are restricted such that normal recessions have no long run effect. See text.

Information from events before the peak at time $t = \tau(r)$ are ignored. This means that if one or previous cycles terminated early at a horizon $h < H_{max} = 10$ then potential catch-up growth due to the later years of that cycle at horizons $h + 1, \dots, 10$ is ignored.

We therefore formally consider the natural alternative where the growth effects from all past cycles are always carried forward into future periods, whether those are in the same cycle r or not. To do this, we simply re-estimate Equation 8 with the added requirement that the $N_{i,t}^h$ and $F_{i,t}^h$ indicators are now redefined, only for this estimation, to denote an event that occurred at time t , whether or not that event was within the same cycle or not, unlike before.

Our estimate of this “non-truncated” LP model is shown in Table 5, column (2) and Figure 5, panel (b). It is clear from the figure that overall the change in the recession paths is fairly small. The only notable change, as expected, is in the estimate of the conditional drift g , which is 0.42% per annum lower in this case (2.38 versus 2.80). The intuition is straightforward: in the FCLP in column (1), growth catching up in the late stages of cycles is truncated. When a new recession peak event hits, no cumulation of any past cyclical catch up is carried forward. Thus when a cycle of either type ends at a horizon $h < 10$ a permanent level loss is baked in going forward. To offset this, and for the model to still match the unconditional mean as it should, a higher drift intercept must be estimated. In column (2), full recovery and catch up out to $h = 10$ is allowed and so no permanent losses materialize for normal recessions, and smaller permanent losses materialize for financial recessions. Thus, a higher drift intercept must be estimated. Either way, as the table shows, the models still always match the unconditional mean.


Still the LP path difference are not that great, and in what follows as our preferred method we use the truncated estimation method for three reasons. First, judged by measures of fit like R^2 and RMSE, the results in Table 5 favor the model in column (1) versus (2). Second, as we shall see below, the differences in estimated welfare that obtain when we use these two methods are negligible. Thirdly, when we move to random coefficient estimation the technical difficulties of keeping track of multiple draws over past cycles of the latent state variable ζ in the RCLP would prove computationally very costly, for minimal gain of insight, as compared to simply modeling a single draw for the current cycle at each observation.

Random coefficients We now move to our preferred baseline approach where a variable-disaster data generating process is estimated using random coefficients local projections (RCLP) based on the specification in Equation 7 and we impose $\tau(r) < t \leq \tau(r) + H(r)$, so estimation of the growth effects is again within-cycle. As always, we continue to impose the restriction $s_1^n + \dots + s_{10}^n = 0$.


Our estimate of this RCLP model is shown in Table 5, column (3) and Figure 6, panel (a). In the table, the estimate of σ_ζ^2 is added, along with its standard error. In the figure, as before, estimated \hat{a}^n and \hat{a}^f coefficients are again shown, with 95% confidence intervals. Recall that these coefficients are now median growth effects, corresponding to a draw of $\zeta = 0$. To give a sense of the range of variability of the growth effects the figure also displays a fan chart on dotted lines corresponding to every decile of the distribution of ζ , where blue denotes normal and red denotes financial.

Figure 6: Recession paths for Normal and Financial Crisis Recessions: Random Coefficient Local Projections

(a) RCLP estimates



(b) RCLP estimates (GBF)



Notes: The figure shows Normal versus Financial Crisis recession paths based on Random Coefficient Local Projections. 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. Note that coefficients are restricted such that normal recessions have no long run effect. See text.

Note that, as a robustness check, we have estimated variants (not shown) where we allow σ_ζ to differ between normal and financial recession events. However, the results indicated no significant difference between these two σ_ζ estimates, so we kept the more parsimonious specification above. It may seem surprising given that the fan charts show wider ranges for financial versus normal events, but this is just a result of the fact that growth deviations on the financial are larger, and so scale up to create much larger ranges when multiplied by the random draws of $\exp(\zeta_{i,\tau(r)})$.

The main message from these estimates is that the RCLP model is the preferred specification so far. The estimated s and a coefficients are quite similar to those in the FCLP models. But we can see that measures of fit are better for this model as compared to the FCLP models, with the R^2 rising substantially to 0.35, versus 0.26 and 0.22. The variance of the random draw parameter σ_ζ is not close to zero, is precisely estimated, and the fan charts show that the induced variability is meaningfully large. The data seem to be much better described by a mini-disaster model that allows for variability in disaster intensity.

Finally, we estimate a variant of our preferred RCLP model that allows for dimension reduction in the description of the data generating process. It is well understood that an important occupational hazard for researchers studying rare disasters is the temptation to try to estimate too many parameters from too few observed disaster states, as noted by [Nakamura, Steinsson, Barro, and Ursúa \(2013\)](#). In our RCLP model we are estimating 22 parameters: the 20 fixed coefficients s with one restriction, the drift g , the residual variance σ_ϵ^2 , and a random draw variance σ_ζ^2 . We now show how to employ a standard “smoothed LP” technique to reduce the 20 fixed coefficients s down to 6.

Formally, in any setting one can restrict LP impulse response coefficients to be described by a suitable parametric function of the horizon h . A useful choice in many applications is to approximate impulse responses by using one or more Gaussian Basis Functions of the form $G(h; \theta_l, \theta_m, \theta_n) = G(h; \theta) = \theta_l \exp((h - \theta_m)/\theta_n)^2$, which was proposed by [Barnichon and Matthes \(2018\)](#). We therefore re-estimate the RCLP model subject to the following functional approximations of the cumulative impulse responses:

$$\begin{aligned} a_h^n &= G(h; \theta^1) - G(10; \theta^1), & h = 1, \dots, 10; \\ a_h^n &= 0, & h > 10; \\ a_h^f - a_h^n &= G(h; \theta^2), & h = 1, \dots, 10; \\ a_h^f - a_h^n &= G(10; \theta^2), & h > 10. \end{aligned}$$

These describe the cumulative impulse responses a , which can be straightforwardly differenced to obtain the noncumulative impulse responses s . Here, in the first two equations, the normal recession response a_h^n is described by a Gaussian Basis Function up to horizon 10, and the intercept shift term is imposed in order to force a zero cumulative level effect at horizon 10 and beyond. In the second two equations, the additional financial recession response $a_h^f - a_h^n$ is described by a second Gaussian Basis Function up to horizon 10, and the level difference then becomes a persistent cumulative level effect which stays constant after horizon 10.

Our estimate of this RCLP-GBF model is shown in [Table 5](#), column (4) and [Figure 6](#), panel (b). With fewer underlying parameters, the precision of each LP coefficient is sharpened, but at some cost in fit. Yet the paths are almost identical, and the R^2 in the GBF approximation falls to only 0.34 versus 0.35 in the preferred RCLP. The main message from these estimates is that the RCLP model and its GBF approximation differ very little. The approximation is good, and it could be employed as a parsimonious description of mini-disaster cycles if concerns about over-parameterization prove troubling, or if the researcher wishes to impose a smoothed response.

Summing up, our estimates of the data generating process appear robust and stable. With these estimates in hand we now turn to simulating these models as the basis of counterfactual welfare analysis.

4. 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.

Our starting point will be to simulate our baseline two-state model of consumption growth. In “good” no-disaster states consumption follows the familiar Gaussian random walk with drift model. However, with some probability drawn from a Bernoulli distribution, a “bad” recession disaster event takes place, although we rule out a disaster in consecutive years to match the empirical constraint that a peak in output or consumption by definition cannot occur in two consecutive years. If such a disaster occurs, then with some probability taken as a sub-draw from another Bernoulli distribution, the normal or financial type of disaster is decided. When such draws yield a disaster, we then characterize the average path of consumption using our preferred RCLP estimates \hat{s}_h above, adjusting its severity by taking a draw calibrated to the latent process ζ . Thus, growth deviations imposed on the path are $\exp(\zeta)\hat{s}_h$ as $h \geq 1$ counts upwards within a given cycle.

Once on a disaster path, the economy remains on it until either (i) horizon $h = H_{max}$ is encountered and the economy returns to the no-disaster state, or (ii) another disaster draw is encountered at horizon $h < H_{max}$ and a new disaster path begins immediately, truncating the old path. In the latter case, no further growth deviations from the old path are cumulated in the baseline model. As a robustness check we also simulate versions of the other three data generating processes we estimated above, although we find only small difference in their welfare implications.

To sum up, the consumer in our simulation faces uncertainty from various quarters. First, at any point a disaster may take place. Second, when a disaster happens, its type and severity are uncertain. Third, the duration of the disaster is also variable. Faced with these multiple sources of uncertainty, the consumer is willing to pay a price to insure against the instability of the consumption stream. How much consumption is she willing to forgo? This is big question we try to answer.

4.1. Model simulations

To make progress, we simulate an artificial consumption series given by the following process 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 two-state process given by

$$\Delta c_t = \hat{g} + \epsilon_t, \quad (\text{RCLP: No-disaster period})$$

$$\Delta c_t = \hat{g} + \exp(\zeta_t) \sum_{h=1}^{H_{max}} \left(\hat{s}_h^f F_t^h + \hat{s}_h^n N_t^h \right) + \epsilon_t. \quad (\text{RCLP: Disaster period})$$

where ϵ_t is drawn from $\mathcal{N}(0, \hat{\sigma}_\epsilon^2)$, and ζ_t is drawn from $\mathcal{N}(0, \hat{\sigma}_\zeta^2)$ once per recession cycle at $\tau(r)$. The parameters \hat{g} , \hat{s}_h^f , \hat{s}_h^n , $\hat{\sigma}_\epsilon^2$, and $\hat{\sigma}_\zeta^2$ are taken from the RCLP estimates reported in [Table 5](#).

To fully specify the simulation, it remains to define the process governing the evolution of the disaster event dummies indicating when there is a recession (normal or financial), namely $R_t = \{N_t, F_t\}$. So what remains is to set the probabilities for the two Bernoulli draws: the top-level draw probability (p) for entering a disaster period (excluding in years with $h = 1$ as noted), and the sub-draw probability for whether the type is normal (q_n) or financial (q_f).

Candidates for baseline probabilities are chosen as follows, based on empirical frequency in the data in the full sample of the JST dataset, excluding wars:

- **Disaster events** For the top-level draw, using earlier notation, we set $p = P(D_t^0 = 1 | D_{t-1}^0 = 0) = 16.3\%$, the sample frequency of a disaster today given no disaster the previous year.¹¹
- **Peak types** Given a disaster occurs, $D_t^0 = 1$, the sub-draw probability determines type. The sample frequency is $q_f = 28.6\%$ for a financial recession (F) and $q_n = 71.4\%$ for a normal recession (N), with $q_n + q_f = 1$.

To illustrate how varying the sub-draw probabilities (q_n, q_f) affects welfare outcomes, holding p fixed, we also consider alternative simulations using a range of other calibrations as follows:

- **Zero financial crisis risk** $q_f = 0, q_n = 1$, empirical frequencies for the “quiet” 1950s–60s.
- **Medium financial crisis risk** $q_f = 0.25, q_n = 0.75$, approximate empirical frequencies of the pre-WW1 era.
- **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.

Here, empirical frequencies are based on the dataset and classifications above. We explore a range of scenarios to provide guidance on the welfare cost implications of varying financial crisis risk.

¹¹Note that this probability is higher than the unconditional JST recession probabilities used earlier in this paper. It must be defined here in a different way given the simulated data-generating process.

We can now completely specify the simulated consumption process. Simulation begins at time $t = 0$ and (wlog) $h = H_{max} + 1$ with the economy in a No Disaster Period. The relevant draw for consumption growth is then made and t and h are stepped forward one year. But if a peak was drawn, a ζ is drawn afresh for the newly-starting Disaster Period, and h will be forced back to 1 next period. Two peaks cannot happen in a row, so no draws happen when $h = 1$, but after that draws restart. A new Disaster Period may be drawn when $1 < h \leq H_{max}$ in which case the current cycle Disaster Period ends, and a new cycle Disaster Period starts up. If at any time we reach $h > H_{max}$, we revert to the No Disaster process until a draw dictates a new Disaster Period. To recover the path in log levels, c_t , we simply cumulate growth draws, and then we can convert to absolute levels C_t .

In fact, we implement our simulation in a simpler way. We actually start by simulating only two consumption growth series: one in which all disasters are Financial recessions ($q_f = 1, q_n = 0$) and one in which all disasters are Normal recessions ($q_f = 0, q_n = 1$). Any other sub-draw probability ($0 < q_n, q_f < 1$) can easily be simulated by switching between these two extreme processes, for any sub-draw probabilities which we wish to assign, including the empirical frequencies above, and then re-cumulating to recover levels as necessary.

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, $c_t = \hat{g} t$, to be used as a welfare baseline; a ‘‘Lucas’’ path, $c_t = \hat{g} t + v_t$, which is a deterministic trend plus i.i.d. Gaussian shocks; and an ‘‘Obstfeld’’ path, $\Delta c_t = \hat{g} + \epsilon_t$, which is a random walk with drift and i.i.d. Gaussian shocks. Formally, we can write these as follows using first differences, imposing $\epsilon_t = v_t - v_{t-1}$, so

$$\Delta c_t = \hat{g}, \quad (\text{Deterministic})$$

$$\Delta c_t = \hat{g} + v_t - v_{t-1}, \quad (\text{Lucas})$$

$$\Delta c_t = \hat{g} + \epsilon_t. \quad (\text{Obstfeld})$$

Here, \hat{g} refers to the conditional annual growth rate term in our baseline RCLP model estimation, and $\hat{\sigma}_\epsilon^2$ is also from the same estimation. We set $\hat{\sigma}_v^2 = \hat{\sigma}_\epsilon^2/2$ to match residual moments.

What is being assumed in these counterfactuals? Following standard practice in the disasters literature, we assume that all permanent losses and temporary skew arising from the extra disaster terms are completely switched off in the counterfactuals (see, e.g., Barro, 2006).¹² In the case of paths with permanent losses (i.e., our financial recessions) this means that the counterfactual process will inherit a higher unconditional mean growth, as well as being stripped of negative-skewed terms. For paths with no permanent losses (i.e., our normal recessions) the counterfactual process will have unchanged unconditional mean growth but will still be stripped of negative-skewed terms. In the counterfactuals, these changes will contribute to higher welfare for the representative consumer.

¹²An 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).

For numerical implementation, we simulate a large history of all series over $T = 110,000$ years for the five models indexed by $M \in \{\text{Financial, Normal, Lucas, Deterministic, Obstfeld}\}$. For each model history we then randomly draw starting points for $S = 1,000$ episodes, each of 1,000 years length, based on which we compute 1,000 distinct consumption series $c_{M,s,t}$ for each model. Using the re-weighting method above we can use the Financial and Normal type draws to derive a path for any $0 < q_f < 1$. We then obtain a welfare estimate for each model by averaging the discounted lifetime utility for all the related 1,000 sub-series draws. We assume CRRA preferences with risk aversion parameter $\gamma = 2$ and a discount factor of $\beta = 0.96$ for annual data.

The welfare level for model M is then given by

$$\text{Welfare}_M = \frac{\sum_{j=1}^S \sum_{t=s(j)}^T \beta^{t-s} \frac{c_{M,s,t}^{1-\gamma}}{1-\gamma}}{S},$$

where $s(j)$ refers to the initial period of the j th sub-series.

Our utility choices are conservative. By today's standards the choice of $\gamma = 2$ may seem low, but it is midway between the values picked in seminal papers, namely $\gamma = 1$ (Lucas, 1987, 2003) and $\gamma = 4$ (Barro, 2006, 2009). We will, however, consider other γ values as a sensitivity check. As for our functional form, Barro (2006); Lucas (1987, 2003) preferences are CRRA, but Barro (2009) employs Epstein-Zin-Weil preferences. EZW preferences tend to magnify welfare losses, all else equal. However, as we show, our CRRA environment with a conservative choice of $\gamma = 2$, generates considerable welfare losses: in our setup, disaster-style fat tails are a feature of all recessions, and so weigh much more heavily in the full welfare cost.


4.2. Comparing welfare under actual and counterfactual histories

Our main result is shown in Figure 7. Using the deterministic model as a baseline, the vertical axis shows welfare relative to the baseline, in $\log \times 100$ units, computed as $\Delta \text{Welfare}_M = \log(\text{Welfare}_M / \text{Welfare}_{\text{deterministic}})$. 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 line at $\Delta \text{Welfare}_M = 0$.

Consider first the Lucas reference model, shown by the thin solid blue line. The welfare loss is a fixed amount and miniscule (about 0.05%), a small fraction of one percent, in line with the prior literature even with the more volatile histories of consumption growth for samples that are broader than just the tranquil post-WW2 United States period. Next consider the Obstfeld model, shown by the thin green dot-dash line. The welfare loss is still a fixed amount but notably larger, about 1%, which is about 20 times the Lucas loss, and as expected for the random walk case where shocks are permanent rather than transitory.

Now we move on and compare with our disaster model. In contrast, the solid red line shows the welfare losses for our simulated model, using the RCLP estimates, which are roughly and order of magnitude larger, ranging between about 10 and 25 percent. The line shows how the

Figure 7: Simulation: Main Results



Notes: The graph portrays welfare losses associated with our baseline disaster model and also for the Lucas and Obstfeld reference cases, relative to the deterministic trend. Units are $\log \times 100$. Our baseline simulation is shown by the thick red line and is based on the annual growth rate regression using RCLP and $H = 10$. The Lucas reference model is shown by the thin solid blue line. The Obstfeld reference model is shown by the thin solid green line. The diamonds refer to relative financial recession probabilities q_f . The hollow black diamonds correspond to the zero, medium, and high frequency baselines. The solid black diamond corresponds to the empirical value of q_f in our full non-war JST sample data. See text.

variable financial crisis disaster frequency q_f impacts the welfare losses relative to the deterministic counterfactual. Three particular points on the line are picked out with hollow black diamonds corresponding to the zero, medium, and high frequency baselines described above. Note that these illustrative cases hold fixed the top-level Bernoulli recession disaster probability p , and change only the sub-draw probabilities (q_n, q_f) . In reality, the top-level p value may also have varied across eras, but we do not attempt to recalibrate the parameter here for an illustrative exercise.¹³ Finally, we denote with a solid black diamond the welfare loss given the empirical value of q_f in our full non-war JST sample data. At the latter point, the welfare loss is about 15%, which is 15 times the Obstfeld loss, and 300 times the Lucas loss.

These calculations are staggering, given that we are using a CRRA parameter of just $\gamma = 2$. They 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 Barro (2006); Lucas (1987, 2003) 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-Gaussian, with significant fat tails in the consumption path relative to the null. Figure 7 shows that, in welfare terms, this really matters.

¹³For example, recessions were more frequent before WW1 and less frequent after WW2, compared to the full sample, which would pull down the middle dot, and lift up the left and right dots somewhat, if the values of p were also adjusted.

Why? As the figure shows, a representative agent living in a world of high financial crisis risk, like the recent decades, ($q_f = 0.50$, post-1985) would enjoy a welfare gain of about 20% from moving to the deterministic path. In a world of medium financial crisis risk ($q_f = 0.25$, pre-WW1) they would gain about 15% from moving to the deterministic path. But even with *zero* financial crisis risk, and all recessions constrained to be of the normal type ($q_f = 0.00$, 1950s–1960s), they would gain 10% from moving to the deterministic path. Taking a rough post-WW2 average, an agent living in time like our current peacetime era would also gain about 10% from moving to the deterministic path. These are very large permanent-consumption equivalent gains, and they arise with wartime rare disasters not present at all in the analysis, and with no rare disasters of any kind in the latter simulation restricted to purely normal peacetime recessions.


4.3. Robustness

Sensitivity to Other Fixed and Random Coefficients Models A robustness check shown in Figure 8 examines shutting down variable disasters by constraining the randomly drawn scaling factor $e^{\tilde{\zeta}(t)} = 1$ (equivalently, setting $\hat{\sigma}_{\tilde{\zeta}}^2 = 0$). For this we replace our baseline RCLP estimates with the FCLP estimates in Table 5. The corresponding welfare results are also shown by the blue dash-dot line in and there is no surprise here. At the empirical value of q_f , baseline welfare costs of about 15% under RCLP estimation decline to about 13% when we switch to fixed-coefficient LP estimation. Interestingly, the effect of variable disasters is of the expected direction, but not as large as in the rare disasters model. As noted by Barro (2006, 2009), the additional kick from variable disasters matters profoundly for risk premium and welfare analysis as compared to fixed disasters. The same is of course true here in our “disasters everywhere” setup. But whereas, say, Barro (2006) rare disasters had a widely dispersed empirical distribution between 15% and over 60%, our latent scaling shock produces much less dispersion in the impulse response paths, as we saw from the decile fan charts seen earlier. Hence, adding variability to our disasters matters less for welfare.

Two other robustness checks are shown in Figure 8 using simulations of the other two models in Table 5. The FCLP estimation without truncation is simulated like the plain FCLP, but growth effects are fully cumulated in all cycles out to $h = H_{max} = 10$, even if a new disaster event occurs. The RCLP with GBF estimation is exactly as in the baseline, except that the coefficients \hat{s}_h are replaced with their GBF approximated versions. Neither of these models radically change the welfare implications of our analysis. The FCLP without truncation produces smaller welfare costs as expected: it approximates cycles with lower conditional growth but also less negatively skewed \hat{s}_h . Thus, the simulated path is a bit smoother than the plain FCLP. But the effect on welfare is minimal, about 1%–2% at any level of q_f . The effect of switching from plain RCLP to RCLP with GBF are smaller still. The latter makes welfare losses a little smaller, but by less than 1% at any level of q_f .


Sensitivity to risk aversion parameter γ A robustness check shown in Figure 9 examines the sensitivity of calculations to different values of the CRRA risk aversion parameter γ . As noted, we

Figure 8: *Simulation: Sensitivity to Other Fixed and Random Coefficients Models*



Notes: The graph portrays welfare losses associated with our baseline disaster model and alternative models, relative to the deterministic trend. Units are $\log \times 100$. Our baseline simulation is shown by the thick red line and is based on the annual growth rate regression using RCLP and $H = 10$. The RCLP with GBF results for $H = 10$ are shown by red dash line. The FCLP results for $H = 10$ are shown by blue dash-dot line. The FCLP with non-truncation results for $H = 10$ are shown by blue dot line. The diamonds refer to relative financial recession probabilities q_f . The hollow black diamonds correspond to the zero, medium, and high frequency baselines. The solid black diamond corresponds to the empirical value of q_f in our full non-war JST sample data. See text.

Figure 9: *Simulation: Sensitivity to risk aversion parameter*



Notes: The graph portrays welfare losses associated with our baseline disaster model and for alternative CRRA choices, relative to the deterministic trend. Units are $\log \times 100$. Our baseline simulation with $\gamma = 2$ is shown by the thick red line and is based on the annual growth rate regression using RCLP and $H = 10$. The results for $\gamma = 3$ are shown by the red dash line. The results for $\gamma = 3$ are shown by the red dot line. The hollow black diamonds correspond to the zero, medium, and high frequency baselines. The solid black diamond corresponds to the empirical value of q_f in our full non-war JST sample data. See text.

made a conservative choice with $\gamma = 2$. Lucas (2003) considers values in the range of 1 to 4 plausible and we are in the middle of that range, close to the highest value 2.5 Lucas entertains.¹⁴ Barro (2006, 2009) has a baseline of $\gamma = 4$, which is low relative to values of $\gamma = 10$ or more in macro asset pricing research. Here, as alternatives to our baseline $\gamma = 2$, we re-compute welfare losses for the cases $\gamma = 3$ and $\gamma = 4$.

The results are shown in Figure 9 and the losses can be huge for even these ostensibly moderate risk aversion parameters, rising above 20% (at the empirical q_f) when $\gamma = 3$, and over 30% when $\gamma = 4$. These are very large welfare losses, and of the same order of magnitude as the rare disasters (wars and Great Depression events) in Barro (2009). Although our disasters are smaller, less skewed, and happen in peacetime, they occur much more often and the welfare cost really adds up.

5. CONCLUSIONS

At the time of writing this paper, the world's economies were experiencing 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. Absent from standard stochastic growth models, such extreme left-tail events present the most obvious justification for stabilization policy.

However, our argument in this paper is that such extreme and costly rare events are not the only tail events that matter. This is because disasters are everywhere and we do not live in a Gaussian world. We find that all business cycles are asymmetric and resemble "mini-disasters". Consumers 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 be happy to pay a nontrivial cost to insure against them.

Our paper re-calculates the costs of business cycles in this setting of frequent fat tails. The size of the welfare loss that we find is large: almost 20% of permanent consumption for cycles in the crisis-prone past three decades under very moderate risk aversion, maybe 15% over the full sample. This loss is more than two magnitudes above that in a Lucas (1987) model, and more than one magnitude above that in an Obstfeld (1994) model. Moreover, the welfare costs 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. More than ever, depression prevention and stabilization policies are central to the discipline of macroeconomics.

¹⁴Low values close to one were preferred to avert the risk-free rate puzzle, but subsequent research has proposed alternative theories to solve that separate but distinct conundrum.

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.
- Barnichon, Regis, and Christian Matthes. 2018. Functional Approximation of Impulse Responses. *Journal of Monetary Economics* 99(C): 41–55.
- 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. 2008. Macroeconomic Crises since 1870. *Brookings Papers on Economic Activity* 2008: 255–335.
- Berge, Travis J, and Òscar Jordà. 2011. Evaluating the classification of economic activity into recessions and expansions. *American Economic Journal: Macroeconomics* 3(2): 246–77.
- Bowsher, Clive G. 2007. Modelling security market events in continuous time: Intensity based, multivariate point process models. *Journal of Econometrics* 141(2): 876–912.
- Bry, Gerhard, and Charlotte Boschan. 1971. *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs*. New York: National Bureau of Economic Research.
- Busch, Christopher, David Domeij, Fatih Guvenen, and Rocio Madera. 2018. Asymmetric Business-Cycle Risk and Social Insurance. NBER Working Papers 24569.
- Cerra, Valerie, and Sweta Chaman Saxena. 2008. Growth Dynamics: The Myth of Economic Recovery. *American Economic Review* 98(1): 439–57.
- Chauvet, Marcelle, and James D Hamilton. 2006. Dating business cycle turning points. *Contributions to Economic Analysis* 276: 1–54.
- 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.
- Cox, David R. 1972. The statistical analysis of dependencies in point processes. In *Stochastic Point Processes: Statistical Analysis, Theory, and Applications*, edited by Lewis, Peter A. W., 55–66. New York: Wiley.
- Cox, David R., and Valerie Isham. 1980. *Point Processes*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Boca Raton, Fla.: Chapman & Hall/CRC.

- Davis, Michael C., and James D. Hamilton. 2004. Why Are Prices Sticky? The Dynamics of Wholesale Gasoline Prices. *Journal of Money, Credit and Banking* 36(1): 17–37.
- 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.
- Diebold, Francis X., and Glenn D. Rudebusch. 1990. A Nonparametric Investigation of Duration Dependence in the American Business Cycle. *Journal of Political Economy* 98(3): 596–616.
- Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson. 2019. A Plucking Model of Business Cycles. NBER Working Paper 26351.
- Engle, Robert F., and Asger Lunde. 2003. Trades and Quotes: A Bivariate Point Process. *Journal of Financial Econometrics* 1(2): 159–188.
- Engle, Robert F., and Jeffrey R. Russell. 1998. Autoregressive conditional duration: a new model for irregularly spaced transaction data. *Econometrica* 66(5): 1127–1162.
- 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.
- Hamilton, James D., and Òscar Jordà. 2002. A model of the federal funds rate target. *Journal of Political Economy* 110(5): 1135–1167.
- Hu, Ling, and Peter C. B. Phillips. 2004. Dynamics of the federal funds target rate: a nonstationary discrete choice approach. *Journal of Applied Econometrics* 19(7): 851–867.
- 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.
- Jordà, Òscar, Sanjay R. Singh, and Alan M. Taylor. 2020. The Long-Run Effects of Monetary Policy. NBER Working Paper 26666.
- 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.
- Last, Günter, and Andreas Brandt. 1995. *Marked Point Processes on the Real Line: The Dynamical Approach*. New York: Springer.
- 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.
- Montiel Olea, José Luis, and Mikkel Plagborg-Møller. 2020. Local Projection Inference is Simpler and More Robust Than You Think. Columbia University. Unpublished.
- 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.
- 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.
- Wachter, Jessica A. 2013. Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance* 68(3): 987–1035.