

Jordà, Òscar; Schularick, Moritz; Taylor, Alan M.

Working Paper

Disasters everywhere: The costs of business cycles reconsidered

Staff Reports, No. 925

Provided in Cooperation with:
Federal Reserve Bank of New York

Suggested Citation: Jordà, Òscar; Schularick, Moritz; Taylor, Alan M. (2020) : Disasters everywhere: The costs of business cycles reconsidered, Staff Reports, No. 925, Federal Reserve Bank of New York, New York, NY

This Version is available at:

<https://hdl.handle.net/10419/241118>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Federal Reserve Bank of New York
Staff Reports

Disasters Everywhere: The Costs of Business Cycles Reconsidered

Òscar Jordà
Moritz Schularick
Alan M. Taylor

Staff Report No. 925
May 2020



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Disasters Everywhere: The Costs of Business Cycles Reconsidered

Òscar Jordà, Moritz Schularick, and Alan M. Taylor

Federal Reserve Bank of New York Staff Reports, no. 925

May 2020

JEL classification: E13, E21, E22, E32

Abstract

Business cycles are costlier and stabilization policies more beneficial than widely thought. This paper shows 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 all advanced economies since 1870. Focusing on the peacetime sample, we develop a tractable local projection framework to estimate consumption growth paths for normal and financial-crisis recessions. Using random coefficient local projections we get an easy and transparent mapping from the estimates to the calibrated simulation model. 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. In postwar America, households would sacrifice more than 10 percent of consumption to avoid such cyclical fluctuations.

Key words: fluctuations, asymmetry, local projections, random coefficients, macroprudential policy

Schularick: Federal Reserve Bank of New York; University of Bonn; CEPR (email: schularick@uni-bonn.de). Jordà: Federal Reserve Bank of San Francisco; University of California, Davis (email: oscar.jorda@sf.frb.org; ojorda@ucdavis.edu). Taylor: University of California, Davis; NBER; CEPR (email: amtaylor@ucdavis.edu). This work is part of a larger project kindly supported by research grants from the European Research Council (ERC) and the Institute for New Economic Thinking, and the authors are grateful for this support. In addition, they are particularly thankful to Johannes Matschke for outstanding research assistance. All errors are the authors’ own. The views expressed in this paper are those of the authors and do not necessarily represent the position of the Federal Reserve Bank of New York, the Federal Reserve Bank of San Francisco, or the Federal Reserve System.

To view the authors’ disclosure statements, visit
https://www.newyorkfed.org/research/staff_reports/sr925.html.

1. INTRODUCTION

Few questions in the discipline of macroeconomics divide schools of thought as passionately as the cost of business cycles and the gains from stabilization policy. In a highly influential back of the envelope calculation, Robert Lucas made the point in 1987 and prominently again in 2003 in his presidential address to the American Economic Associations that the welfare gains from stabilizing the business cycle were extremely low. The estimate by Lucas (1987, 2003) showed that transitory i.i.d. Gaussian consumption deviations from the postwar U.S. trend, under log utility, deliver very small costs: less than $1/10$ of a percent. Clearly, one might allow more volatile consumption, like the pre-war U.S. or other countries; or allow expected utility and stochastic growth (Obstfeld, 1994). But even then it is hard to get the costs of business cycles and hence potential gains from flattening them to exceed 1–2 percent. The implications, both for research and policy, would be profound. Instead of fine-tuning low-cost fluctuations with monetary and fiscal policies, the profession and policy institutions could devote their time and energy to things that matter.

A challenge to Lucas' findings emerged as an implication of asset pricing models that aimed to explain the equity risk premium. Note that the two questions are intimately linked. While Lucas argued that consumption fluctuations were not very costly, asset pricing was confronted with the puzzle that the compensation for bearing risk over the cycle appeared very high. Why are risk premiums so high if consumption volatility had only minimal welfare costs? "Rare disasters" asset pricing models were developed to address this key macro-finance puzzle (Barro, 2006; Gabaix, 2008; Rietz, 1988). In these models, growth has fat tails: it is skewed by large, low-probability, permanent jump losses, and the welfare costs are much larger. Such losses depend on the size and probability of the disaster but, using historical panel data, Barro (2009) arrives at a welfare cost estimate of 17 percent due to the disaster-driven jumps, ten times larger than the 1.6 percent attributable to Gaussian growth disturbances.¹ The costs can be larger still if utility is recursive, or if disasters have stochastic probability, stochastic size, more persistence, or are permanent rather than transitory (Barro and Jin, 2011; Gourio, 2012; Nakamura, Steinsson, Barro, and Ursúa, 2013; Obstfeld, 1994; Reis, 2009). Fat tails matter, with effectively all of the amplification of the welfare costs driven by higher moments (Martin, 2008).²

From the "rare disaster" perspective, the true welfare cost of economic fluctuations mainly stems from such rare, 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. In fact, when we use long-run historical data to look at growth at risk, we find that business cycle events typically contain the kinds of higher moments that amplify welfare costs. That is, such deviations from the Gaussian benchmark do

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

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

not just appear in the extreme disasters—the wars, revolutions, financial crises, etc. considered by previous research. Instead, we argue that fat tails and persistence also appear in “normal recessions:” there are disasters everywhere.³ Once we embrace this idea, welfare cost judgments change profoundly. In our estimates for post-WW2 U.S. data, households would sacrifice about 10-15 percent of consumption to avoid business cycles. Far from being a side-show for macroeconomics, business cycles are at the heart of the discipline.

The first part of our paper documents this new stylized fact using a comprehensive macro-historical database by (Jordà, Schularick, and Taylor, 2017). The result holds true in advanced economies usually considered exempt from the more frequent economic dislocations seen in emerging markets. Our finding ties into recent research exploring 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 these works, skewness appears as a general phenomenon and not just as a disaster pattern.⁴ Our analysis also meshes with inherently asymmetric macro frameworks, such as de Long 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. Jensen, Petrella, Ravn, and Santoro (2020) make the link between this skewed pattern and the growth of leverage, which we also explore below. But, to our knowledge, we are the first to re-examine the costs of business cycles debate in this setting 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, 2020). LP estimation has been successfully employed in fixed-coefficient form to document the systematic, large, and persistent differences in normal and financial crisis recessions (Jordà, Schularick, and Taylor, 2013). Here we extend the approach using random-coefficient local projections (RCLP). In an

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.

extension, we also explore what happens in this setting when the recession type probability has a conditional mean which depends on covariates, and a natural case to consider is when the financial crisis probability depends on the history of credit growth (Schularick and Taylor, 2012).

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.⁵ 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. A large loss affects the trend in both normal and financial crisis recessions, and 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 2 percent loss; allowing fat tails in normal recessions increases this loss to 5 percent; and allowing fat tails in financial crisis recessions increases it to about 25 percent.

The paper makes two main contributions. The first is in terms of empirical methodology, as we present a new estimation and simulation approach built around the attractive technique of local projections, which we show is particularly suited to the problem of measuring disaster losses over multi-period horizons, and 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.

The second contribution is to the perennial applied macro problem of the cost of business cycles. We argue that *all recessions are disasters*—in the sense that the growth process has fat tails everywhere, not just in Barro-type rare disasters. Looking at the post-1985 era in advanced economies, a time of high credit growth where the mix of disaster types empirically was 50-50, our model says agents would sacrifice 15 percent of consumption to avoid all business cycles; and 10 percent just to get back to a world of only non-financial crisis (i.e., only normal) recessions, a goal which was attainable in the 1950s–1960s era. Thus, even outside of military conflicts, there could be considerable gain to smoothing economic fluctuations, and all the more so in financially unstable economies.

⁵Economics 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.

2. DISASTERS EVERYWHERE? TESTING FOR THE PRESENCE OF DISASTERS

The canonical model of per capita consumption growth used in [Lucas \(1987, 2003\)](#), and also by a large literature in macroeconomics and in finance is a random walk with drift. In such a model, it is easy to see why stabilization policy has few benefits: deviations from trend growth tend to be short-lived and mean-reverting.

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 [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. [Lucas \(1987, 2003\)](#) proposes a null model that characterizes the growth rate of x (in logs) along the lines of a random walk with drift, such as:

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

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

Next, we expand on this null model. First, define a mapping m that results in 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.

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

We now allow for the possibility that after an event r , the economy takes a hit beyond what the null model with random drift in [Equation 1](#) would prescribe. However, we assume that after at most

H periods, the economy returns to its steady-state growth path. This can be easily accomplished by defining the dummy variable $d_t^h = \mathbb{I}[t(r) < t < t(r+1); t(r) + h = t]$ where $\mathbb{I}[\cdot]$ is equal to 1 when the conditions in brackets are met, and is 0 otherwise. In words, the dummy variable will be 1 at time t if h periods prior was an event, $t(r)$, date. Hence, Equation 1 can be expanded into:

$$\Delta x_t = g + \sum_{h=1}^H \beta_h d_t^h + \epsilon_t. \quad (2)$$

Therefore, $\beta = (\beta_1 \dots \beta_H)$ describes 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.

The cumulative effect of a disaster after h periods is easily seen to be $C_h = \beta_1 + \dots + \beta_h$. One way to see this is by recursive substitution of Equation 2, specifically:

$$\Delta_h x_{t+h} = x_{t+h} - x_{t-1} = gh + \sum_{h=1}^H \left(\beta_h \sum_{j=0}^h d_{t+j}^h \right) + \sum_{j=0}^h \epsilon_{t+j}. \quad (3)$$

Based on these extensions of the null model, it is straightforward to test whether the data support departures from the null model in Equation 1. Based on Equation 2, one would test the null $H_0 : \beta_1 = \dots = \beta_H = 0$ or alternatively $H_0 : C_1 = \dots = C_H = 0$.

However, the null hypothesis can be sharpened by realizing that a given mapping 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 in the next section generate such paths and how to properly test for the absence of event-specific distortions to the null model of Equation 1.

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. These can be estimated using local projections with the sequence of regressions:

$$\Delta_h x_{t+h} = \mu_h + C_h d_t^0 + u_{t+h}; \quad h = 1, \dots, H. \quad (4)$$

Depending on the mapping, one obtains a sequence of null values of the cumulative impulse response, denoted $C_h^{0,m}$, with which to compare sample estimates from Equation 4, \hat{C}_h . Thus, when testing the null that disasters are simply rare events generated under the null model, we will be interested in testing the null hypothesis $H_0 : C_h = C_h^{0,m}$.

To reiterate, different mappings truncate the calendar-time data-generating-process in Equation 1 differently. Events, or disasters, 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 our numerical procedure to construct the null path works. 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: $t = t(r_{All})$ if $x_{t-1} < x_t > x_{t+1}$, that is, a peak in activity defined as in [Bry and Boschan \(1971\)](#).

Rejection of the null in Test 1 could be explained 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, now 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 wonder whether the null would be rejected if we were to focus 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 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 $\bar{\delta} = g$. Given the Gaussian null, the former is $\bar{\delta}^- = g - l$, where $l = \sigma[\phi(-g/\sigma)/\Phi(-g/\sigma)]$ 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: $\Delta_h \hat{x} = gh - ld_h^0$. That is, the null predicts a dogleg jump down of size l 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 gh . 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, $g = 1.96\%$, $\sigma = 3.62\%$, and the implied jump loss under the null is $l = -4.24\%$. For log real consumption per capita, $g = 1.78\%$, $\sigma = 4.02\%$, and the implied null has $l = -4.42\%$. 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.⁶ 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 Consider an observer using the [Barro and Ursúa \(2008\)](#) rare disaster criterion, which we denote $m = BU$. These dates were constructed in such a way as to encompass disasters defined as events with peak-to-trough declines in annual real GDP per capita of at least 15%. This cutoff creates disaster events where the probability that a disaster begins is about $p = 2.4\%$ in the peacetime sample considered above. Accordingly, we construct the Gaussian null numerically by imposing the same conceptual frequency cutoff. That is, we take all simulated peak-to-trough declines, and find the worst 2.4% among them, and designate these as the rare disaster dates.⁷

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.

The persistent negative loss of output lasts for many periods and the peak deviation from the

⁶Rejection 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.

⁷In the numerical approximations we find that 12% of all observations are recession peaks, so to pick out the peaks that occur 2.4% of the time and have the worst peak-to-troughs of all peaks, we proceed by selecting the worst fifth ($20\% = 2.4/12$) of all peaks by this yardstick.

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

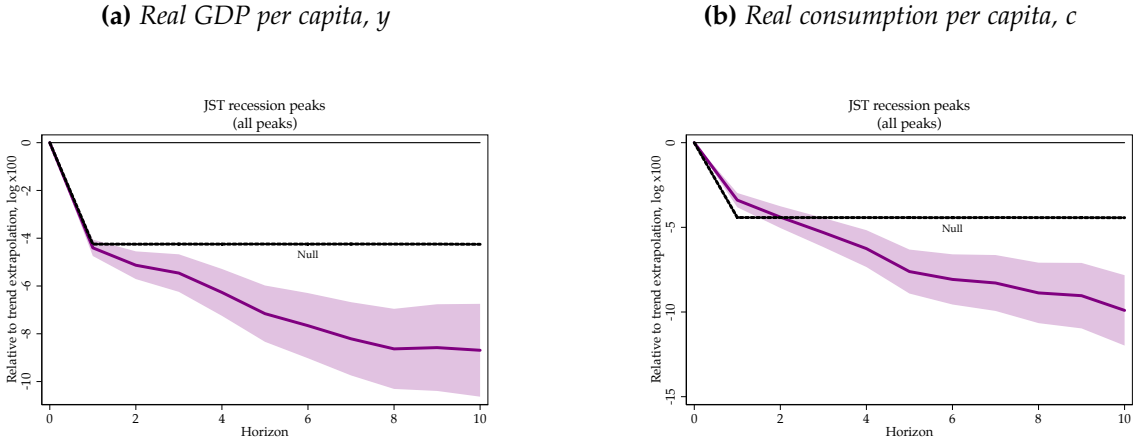


Table 1: LP test for non-Gaussianity, all JST recession peaks, all countries, peacetime sample

(a) Real GDP per capita, $\log y (\times 100)$

	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8	h=9	h=10
$\bar{C}_h^{0,All}$	-4.24	-4.25	-4.25	-4.25	-4.25	-4.25	-4.25	-4.25	-4.25	-4.25
\hat{C}_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, $\log c (\times 100)$

	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8	h=9	h=10
$\bar{C}_h^{0,All}$	-4.42	-4.42	-4.43	-4.42	-4.42	-4.43	-4.43	-4.42	-4.42	-4.43
\hat{C}_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.

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 multi-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 In our third test, we revert to the recession peak criterion used in Test 1, but exclude 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 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.

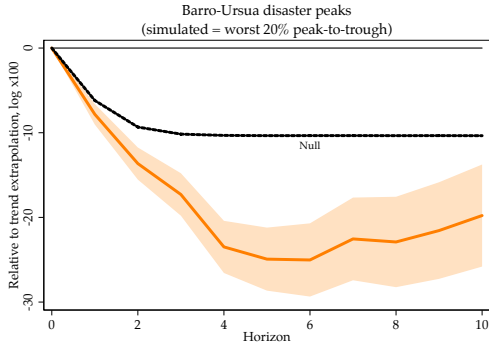
We therefore conclude that the rare disaster approach is valid: those events are still the big disasters. But now even all the other recessions in history can be viewed as mini-disasters of the same ilk. That is, there are disasters everywhere.

Test 4: All peacetime normal and financial recessions minus Barro-Ursúa disasters

Lastly, we consider a natural hypothesis to account for the previous result, by addressing the role of 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 holds true even outside of the rare disasters sample. To attack this question we just employ the [Jordà, Schularick, and Taylor \(2013\)](#) classification of recessions into Normal (N) and Financial Crisis (F) types, where the latter denotes cases where a financial

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

(a) Real GDP per capita, y



(b) Real consumption per capita, c

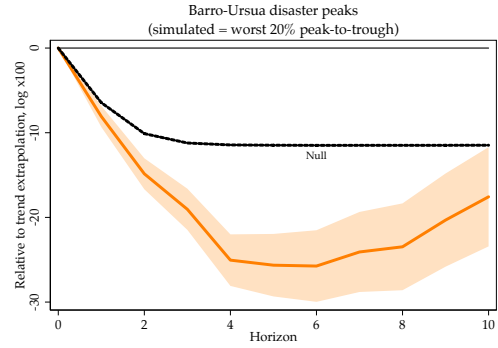


Table 2: LP test for non-Gaussianity, all BU disaster peaks, all countries, peacetime sample

(a) Real GDP per capita, $\log y (\times 100)$

	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8	h=9	h=10
$\overline{C}_h^{0,BU}$	-6.19	-9.33	-10.16	-10.31	-10.34	-10.34	-10.34	-10.34	-10.34	-10.35
\hat{C}_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, $\log c (\times 100)$

	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8	h=9	h=10
$\overline{C}_h^{0,BU}$	-6.45	-10.10	-11.21	-11.44	-11.47	-11.49	-11.49	-11.49	-11.48	-11.48
\hat{C}_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

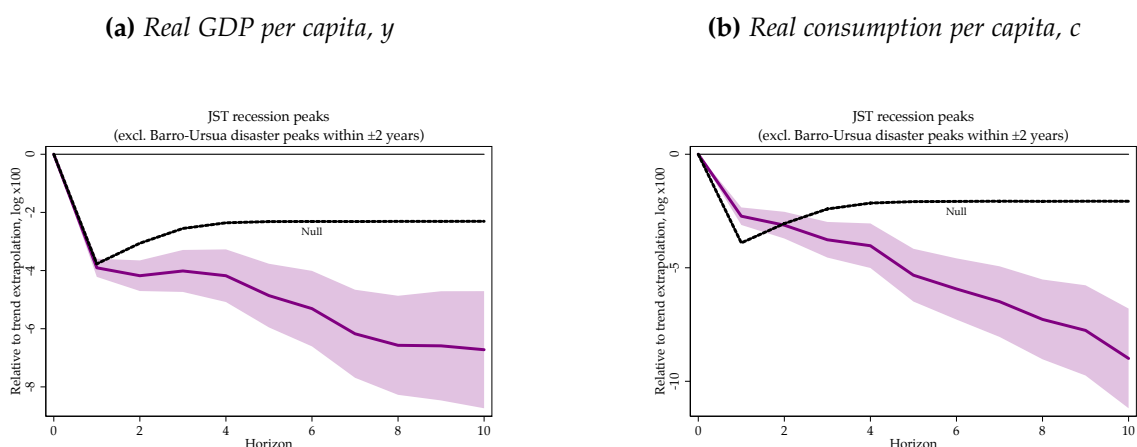


Table 3: LP test for non-Gaussianity, all JST recession peaks, all countries, peacetime sample, excluding BU disaster peaks (± 2 years) from the historical and simulated indicators

(a) Real GDP per capita, log y ($\times 100$)

	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8	h=9	h=10
$C_h^{0,All-BU}$	-3.77	-3.06	-2.55	-2.36	-2.32	-2.31	-2.31	-2.31	-2.31	-2.31
\hat{C}_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 ($\times 100$)

	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8	h=9	h=10
$C_h^{0,All-BU}$	-3.90	-3.05	-2.41	-2.15	-2.09	-2.08	-2.07	-2.07	-2.07	-2.07
\hat{C}_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.

crisis event occurs within ± 2 years of the recession 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.⁸

Implications for welfare: do all disasters matter? If we admit these fat-tailed mini-disasters into the welfare accounting framework, there is prima facie evidence of their quantitative significance. The intuition for this result is straightforward. Welfare calculations, risk premia, and other features of the disaster calculus center on an expression like $pE(l)$, the event probability p times the expected jump loss in an event $E(l)$. 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 have the potential to be very consequential in welfare terms, and possibly even more than the rare disasters.

Some back-of-the envelope calculations motivate the rest of the paper. Specifically, in the peacetime JST panel data we analyzed above, the BU rare disasters have $p = 0.024$ and from the analysis $E(l) = 16\%$, which would be larger if wars were included in our sample, as in Barro (2006). Thus, $pE(L) = 0.384\%$. However, non-disaster Financial Crisis recessions have $p = 0.0024$ and from the charts $E(L) = 8\%$, for $pE(b) = 0.19\%$. And non-disaster Normal recessions have $p = 0.072$ and from the charts $E(L) = 4\%$, for $pE(b) = 0.29\%$. Thus the expected jump loss from recessions *other than* rare disasters has a larger magnitude than that associated with the rare disasters themselves.

This highlights the key concern 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 calculations 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.

⁸Note that the null is numerically constructed as in Test 3 and therefore by construction includes “big” recessions that would probably would better characterize financial crises. Instead of adjusting the null, we left the figure as is to highlight that even in this adverse choice of null, Normal recessions still exhibit statistically measurable 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

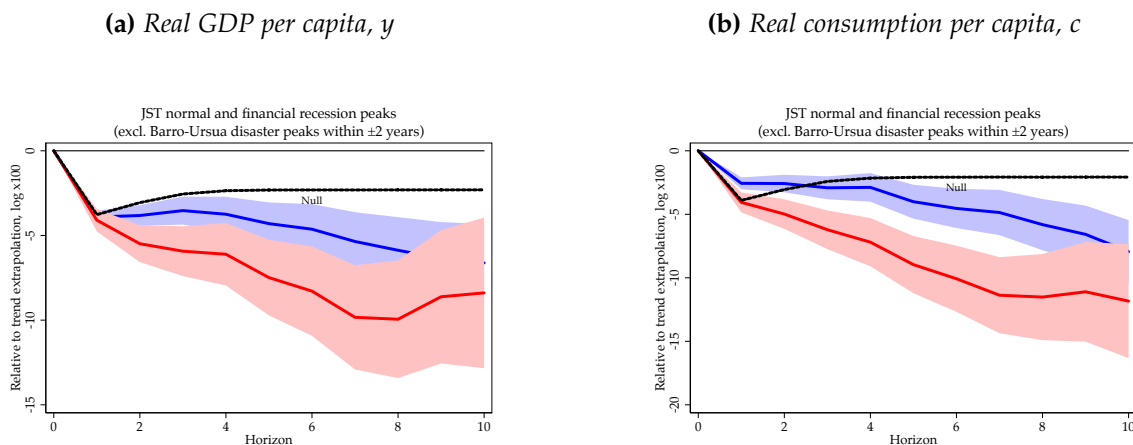


Table 4: LP test for non-Gaussianity, all JST Normal and Financial recession peaks, all countries, peacetime sample, excluding BU disaster peaks (± 2 years) from the historical and simulated indicators

(a) Real GDP per capita, $\log y$ ($\times 100$)										
	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8	h=9	h=10
$C_h^{0,N}$	-3.77	-3.06	-2.55	-2.36	-2.32	-2.31	-2.31	-2.31	-2.31	-2.31
\hat{C}_h^N	-0.14	-0.76	-0.98	-1.39	-1.99	-2.32	-3.04	-3.57	-4.05	-4.31
\hat{C}_h^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, $\log c$ ($\times 100$)										
	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8	h=9	h=10
$C_h^{0,N}$	-3.90	-3.05	-2.41	-2.15	-2.09	-2.08	-2.07	-2.07	-2.07	-2.07
\hat{C}_h^N	1.34	0.47	-0.50	-0.73	-1.92	-2.46	-2.79	-3.74	-4.51	-5.87
\hat{C}_h^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.

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 hits, the economy is shunted onto a different regime whose dynamics we wish to characterize next. Thus, the empirical process corresponding to all calendar dates associated with an event r , that is $\{t(r_m)\}_{r_m=1}^{R_m}$, 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, $\tau(r) = t(r) - t(r-1)$, is approximately distributed as an exponential random variable.

Applications in economics include models describing the arrival of policy rate changes ([Hamilton and Jorda, 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 allow the severity of each disaster to vary randomly, we expand on our previous specification. Therefore, at each event date $t(r)$, 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 $C(r) = (C_1(r), \dots, C_H(r))$ where the $C(r)$ are indexed by r (we omit reference to the mapping for simplicity). Before becoming more specific, it is important to note that the stochastic process given by $\{t(r), C(r)\}_{r=1}^R$ 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 recession will hit, or as stated earlier, $\tau(r) \sim \mathcal{E}(\lambda)$, 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 augment our local projections specification as follows:

$$\begin{aligned} \Delta_h x_{t(r)+h} &= \mu_h + \alpha_h \exp(\zeta_{t(r)}) d_{t(r)}^0 + u_{t(r)+h}; & \zeta_{t(r)} &\sim N(0, \sigma_\zeta^2); & u_{t(r)+h} &\sim N(0, \sigma_h^2); \\ r &= 1, \dots, R; & h &= 1, \dots, H, \end{aligned} \quad (5)$$

where ζ and u are independent from each other. Call this model *random coefficient local projections* or RCLP. We also present the model by omitting the panel dimension of our application to avoid notational clutter. Thus, note that $E(\exp(\zeta)) = \exp(\sigma_\zeta^2/2)$ and hence, when evaluated at the mean, $C_h \approx \alpha_h \exp(\sigma_\zeta^2/2)$ where C_h refers to the non-random version of the local projections in [Equation 4](#).

In order to get a tangible sense of what the RCLP specification buys us, notice that $\exp(\zeta) \in (0, \infty)$ since $\zeta \in (-\infty, \infty)$, meaning that as $\zeta \rightarrow -\infty$, the disaster approaches the null model proposed by [Lucas \(1987, 2003\)](#) 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 $\sigma_\zeta \approx 0.5$. This means that the recession penalty coefficients, α_h are scaled up or down as follows. Take a centered 95% probability range for ζ . At the low end, the coefficient α_h is scaled down by a factor of 0.4, approximately. At the other end, it gets scaled by a factor of approximately 2.7. In other words, the worst recession is about 7 times worse than the weakest recession 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, a mild recession will result in a decline of about 0.8%, whereas a severe recession is in the order of 5.4%. 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. When ζ is at its zero mean, the coefficient α_h is scaled by a factor of 1.3, meaning that our estimates of α_h are approximately such that $C_h \approx \alpha_h$ since $e^{0.5^2/2} = 1.01$. This provides a convenient rule of thumb to compare the fixed and random coefficient model estimates.

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_τ is common to all horizons and the system of equations can be set up accordingly. Although this is the recommended approach for general applications of the model, the particular application that we pursue here allows for a simpler, single equation estimation procedure as follows:

$$\Delta_h x_{t(r)+h} = \mu_h + \left(\sum_{h=1}^H \alpha_h d_{t(r)}^0 \right) \exp(\zeta_{t(r)}) + \eta_{t(r)+h}. \quad (6)$$

Because we have worked so far with cumulated responses, it is worth remarking that a similar

approach can be used when working with typical impulse responses. In that case, Equation 5 can be rewritten as:

$$\Delta x_{t(r)+h} = g_h + \pi_h \exp(\zeta_{t(r)}) d_{t(r)}^0 + v_{t(r)+h}; \quad \zeta_{t(r)} \sim N(0, \sigma_\zeta^2); \quad v_{t+h} \sim N(0, \sigma_h^2). \quad (7)$$

Similarly, Equation 7 would then be expressed instead as:

$$\Delta x_{t(r)+h} = \gamma_h + \left(\sum_{h=1}^H \pi_h d_{t(r)}^0 \right) \exp(\zeta_{t(r)}) + \xi_{t(r)+h}. \quad (8)$$

In our application we will report results based on both approaches by further noting that the α_h coefficients in Equation 6 simply accumulate the coefficients π_h in Equation 8, that is, $\alpha_h = \pi_1 + \dots + \pi_h$. This makes it easy to compare one set of estimates with the other as we do below.

3.1. Estimating the baseline RCLP model

Using the methods just introduced, we estimate baseline values on data that we will then use for the welfare calculation of the cost of business cycles. For this reason we focus only on real consumption per capita in what follows. 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). Accordingly, let $F_{i,t} = 1$ for a business cycle peak of the F type, and $N_{i,t} = 1$ an analogous indicator variable for N type peaks. In due course, the occurrence of these peaks will be modeled as Bernoulli draws, as in standard disaster models.

The sample used for estimation is limited by business cycle peak episodes, $r = 1, \dots, R$. The r th peak in the sample corresponds to the country-year observation $\{i(r), t(r)\}$. The next recession peak does not occur until observation $\{i(r+1), t(r+1)\}$. Hence, the sample used for estimation is $\mathcal{S} = \{i, t, h | i = i(r); t = t(r); h = 1, \dots, H(r); r = 1, \dots, R\}$, that is, we use all country-year pairs during recession episodes only.

We index $H(r)$ by r . Because the length of expansion varies across recessions r . In any case, as a practical matter, we limit included observation to those which satisfy $H(r) \leq 10 = H_{max}$ since few expansion last more than 10 years, and this subsample of observations is too small to afford precise path estimates beyond this horizon. Of course, allowing $H_{max} = 10$ does not constraint the recession to revert to its growth path sooner. We simply let the data determine that.

We now describe the counterpart specification to Equation 6 in cumulative form. Hence, we specify the LHS variable of the LP in similar fashion as $\Delta_h c_{i(r), t(r)+h}$. In parallel, we also estimate the counterpart specification in growth rates given by Equation 8, in which case the LHS is defined as $\Delta c_{i(r), t(r)+h}$. As we will see shortly, the results are essentially the same. We now discuss the specifics of the RCLP specification in cumulative terms. The specification in growth rates is analogous as we show momentarily.

As discussed in the previous section, to allow for the depth of each recession episode to vary randomly, we extend the standard LP set-up to the RCLP specification introduced earlier as follows for the counterpart to Equation 6:

$$\begin{aligned} \Delta_h c_{i(r),t(r)+h} &= \mu_h + \left(\sum_{h=1}^{H(r)} [\alpha_h^f F_{i(r),t(r)} + \alpha_h^n N_{i(r),t(r)}] \right) e^{\tilde{\zeta}_{i(r),t(r)}} \\ &+ \left(\alpha_{LR}^f \mathbb{I}(h > H_{max}) F_{i(r),t(r)} + \alpha_{LR}^n \mathbb{I}(h > H_{max}) N_{i(r),t(r)} \right) e^{\tilde{\zeta}_{i(r),t(r)}} \\ &+ \epsilon_{i(r),t(r)}^h; \quad \forall i, t, h \in \mathcal{S}; H(r) \leq 10 = H_{max}, \end{aligned} \quad (9)$$

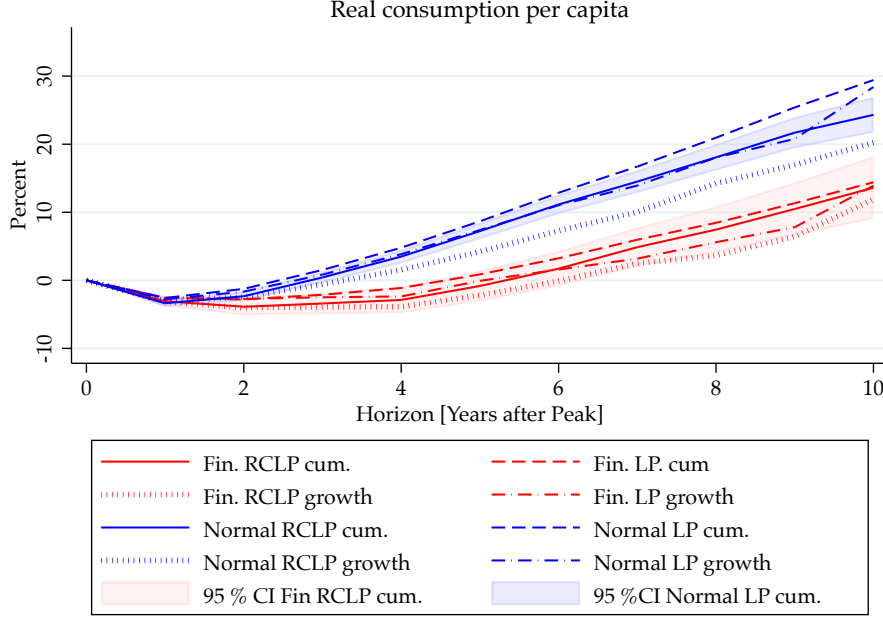
where the terms in the second line of this expression capture unspecified long-run effects sometimes lasting beyond H_{max} . We explain this feature in more detail momentarily. First note that the coefficients α_h^f and α_h^n measure disaster-type losses in financial crises and normal recessions, respectively. In particular, the coefficient α_h^f measures the (one-period or cumulative growth) impact of a financial crisis recession at h -years, relative to trend. The interpretation of the coefficient α_h^n is analogous for normal recessions. Both are modulated by $e^{\tilde{\zeta}_{i(r),t(r)}}$ to allow for variation in the severity of the episode, as explained in the previous section.

Returning to the second line in Equation 9, note that in finite samples—where cycles of arbitrarily long length H are never observed—we cannot achieve reasonable sample sizes of outcomes to measure growth or level deviations at large H . Thus, pragmatically, we truncate the LP estimation at a reasonable horizon H_{max} . This allows estimation using a reasonable number of observations from which to compute growth or cumulative growth. Second, beyond that horizon we implicitly assume in the growth rate regression that the process reverts to the random walk with drift for $h > H_{max}$. That is, the cumulation of all growth losses captured by the α_h^n and α_h^f parameters will cease at that point, and no further permanent loss is incurred: the disaster is over. To make the regression consistent with this formulation, we need to impose permanent losses via the long-run permanent loss terms α_{LR}^f and α_{LR}^n .

Assuming that a financial recession leads to a permanent loss has become the consensus view since at least Cerra and Saxena (2008). Assuming that a normal recession also leads to a permanent loss, however, might be controversial. Therefore, in our baseline level regression results to be conservative, we will impose the restriction $\alpha_{LR}^n = 0$ and assume there is no permanent loss from a normal recession, and only estimate α_{LR}^f . For the baseline growth regression we will do the same, and therefore impose the restriction $\sum_{h=1}^H \alpha_h^n = 0$.

Note that this type of constraint effectively adjusts the trend 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.

Figure 5: Recession paths for Normal and Financial Crisis Recessions



Notes: The figure shows *Normal* versus financial crisis recession (*Fin*) paths by displaying estimates of the α_h^n and α_h^f coefficients in Equation 9 in the cumulative (*cum.*) version in that specification as well as the *growth* rate version (not shown). In addition, the figure displays the same coefficients when estimated with traditional fixed coefficient local projection versions of Equation 9. 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 deviations from the peak of activity. 95 % confidence intervals from the RCLP estimates are provided as shaded regions. The recession paths are scaled to ensure that normal recessions have no long run effect, i.e., the constraint $\sum_{h=1}^{10} \hat{\pi}^h = 0$ is imposed. See text.

This specification makes the connection with the empirical results in Section 2 direct. However, in the next section we conduct our welfare calculations using the specification in growth rates parallel to Equation 8. Thus, the version of Equation 9 in growth rates is easily seen to be:

$$\Delta c_{i(r),t(r)+h} = \mu_h + \left(\sum_{h=1}^{H(r)+1} [\pi_h^f F_{i(r),t(r)} + \pi_h^n N_{i(r),t(r)}] \right) e^{\zeta_{i(r),t(r)}} + \epsilon_{i(r),t(r)}^h; \quad \forall i, t, h \in \mathcal{S}, \quad (10)$$

As we show next, the results from using the cumulative or growth specifications in equations 9 and 10 generate essentially the same results, as one would expect. However, Equation 10 makes it easier to simulate the data for our purposes.

3.2. Results

Based on the previous section, we estimate a battery of specifications that serve as baseline for the welfare calculations in later sections. In particular, we estimate models in cumulative and growth forms, using fixed coefficient local projections (LP) and random coefficient local projections (RCLP), and provide separate results for normal recessions (N) and financial crisis recessions (F). Figure 5 and Table 5 summarize the estimated paths based on all these variants.

Several results deserve comment. First, it is encouraging (though not unexpected) to see that

Table 5: Recession paths for Normal and Financial Crisis Recessions

	(1)	(2)		(3)	(4)
	LP cumulative	RCLP cumulative		LP Growth	RCLP Growth
α_1^n	-5.6*** (0.5)	-4.1*** (0.4)	π_1^n	-5.5*** (0.2)	-5.2*** (0.2)
α_2^n	-7.3*** (0.6)	-5.3*** (0.5)	π_2^n	-1.6*** (0.3)	-1.3*** (0.2)
α_3^n	-7.7*** (0.6)	-5.4*** (0.5)	π_3^n	-0.2 (0.3)	0.3 (0.2)
α_4^n	-7.5*** (0.7)	-5.4*** (0.6)	π_4^n	0.3 (0.3)	0.6** (0.3)
α_5^n	-6.8*** (0.8)	-4.8*** (0.7)	π_5^n	0.9** (0.4)	1.2*** (0.3)
α_6^n	-6.0*** (0.9)	-4.0*** (0.7)	π_6^n	1.3*** (0.4)	1.3*** (0.4)
α_7^n	-6.0*** (1.0)	-3.4*** (0.8)	π_7^n	0.9** (0.5)	0.8** (0.4)
α_8^n	-4.7*** (1.1)	-1.6* (0.9)	π_8^n	1.3** (0.5)	1.0** (0.5)
α_9^n	-4.8*** (1.2)	-1.0 (1.0)	π_9^n	1.5*** (0.6)	1.1** (0.5)
α_{10}^n	0.0 (.)	0.0 (.)	π_{10}^n	1.1* (0.6)	0.2 (0.6)
α_1^f	-5.4*** (0.7)	-4.2*** (0.7)	π_1^f	-5.5*** (0.4)	-5.0*** (0.4)
α_2^f	-8.4*** (0.8)	-6.6*** (0.8)	π_2^f	-3.2*** (0.4)	-2.9*** (0.3)
α_3^f	-11.0*** (1.0)	-8.2*** (1.0)	π_3^f	-2.3*** (0.5)	-1.8*** (0.4)
α_4^f	-13.7*** (1.2)	-9.8*** (1.2)	π_4^f	-1.9*** (0.6)	-1.7*** (0.4)
α_5^f	-14.3*** (1.3)	-10.1*** (1.4)	π_5^f	-0.9 (0.6)	-0.3 (0.5)
α_6^f	-15.5*** (1.6)	-10.0*** (1.6)	π_6^f	-0.6 (0.8)	0.1 (0.6)
α_7^f	-16.7*** (1.8)	-9.5*** (1.6)	π_7^f	-0.2 (0.9)	0.7 (0.6)
α_8^f	-17.1*** (2.0)	-10.2*** (2.0)	π_8^f	-0.5 (1.0)	0.1 (0.8)
α_9^f	-17.8*** (2.3)	-9.6*** (2.0)	π_9^f	-0.1 (1.1)	0.5 (0.9)
α_{10}^f	-14.5*** (3.0)	-6.8*** (2.6)	π_{10}^f	0.1 (1.5)	0.6 (1.1)
α_{LR}^f	-15.9*** (1.2)	-3.8*** (1.1)			
σ_ζ		0.63			0.45
Observations	1915	1891		2157	1915

Notes: Dependent variable for columns (1) and (2): $\Delta_h c_{i(r),t(r)+h} \times 100$. Dependent variable in columns (3) and (4): $\Delta c_{i(r),t(r)+h} \times 100$. Thus, the cumulative sum of the coefficients in columns (3) and (4) is approximately equal to the coefficients in columns (1) and (2). The Table displays regression coefficients for the LP and RCLP models with H=10. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. We impose $\sum_{h=1}^H \alpha_h^n = 0$. The row σ_ζ displays the estimate of the standard deviation of the latent variable ζ in the RCLP specification. See text.

estimates based on our new random coefficient local projection (RCLP) specification match pretty closely the usual fixed coefficient LP counterparts. Second, it is also useful to see that, whether one estimates the model in growth or cumulated terms, essentially makes no difference and that is as true for LPs as it is for RCLPs. The confidence regions plotted alongside the RCLP estimates show that all paths are rather close to each other in the economic and statistical sense. Finally and as is well documented in, e.g. [Jordà, Schularick, and Taylor \(2013, 2017\)](#) there are significant differences between normal and financial crisis recessions. Normal recessions tend to last about one year. By year two the economy returns to the level at the peak, and continues on its growth path thereafter. Financial recessions are deeper, last longer, and even 10 years out, sit about 10 percentage points lower than the path following a normal recession. All in all, these results provide a great deal of comfort regarding the new methods that we introduced as they match previous, well-documented features of economic downturns.

However, the more interesting feature of the RCLP estimates perhaps lies in estimates of the variance of the latent variable ζ , which modulates the strength of the recession. In order to get a closer look, we turn to [Table 5](#). The estimates reported in columns (1) and (2) correspond directly to those shown in [Figure 5](#). The estimates reported in columns (3) and (4) correspond to estimates based on the growth version of the specification. When sequentially cumulated, they are approximately the same as those in columns (1) and (2) as [Figure 5](#) corroborates. Importantly, estimates of σ_ζ whether for the cumulative or growth specifications are very close to each other: 0.63 vs. 0.45, or approximately 0.5 as we claimed earlier.

As we discussed in previous sections, one way to conceptualize what this estimate of σ_ζ means is to explore the implied variation in the depth of a recession at both extremes of a centered 95% probability range. This range puts a severe recession about 7 times worse than a mild recession. Therefore, at an intuitive level, the RCLP estimates make economic sense and will therefore be helpful in constructing the welfare simulations of the coming sections.

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 null model in ways as shown above—disasters are everywhere. If so, how much is a representative consumer willing to pay to such randomness? In this section we take the lessons from our empirical strategy seriously and put them to work in a welfare calculation exercise.

Our starting point will be to model consumption growth as a two-state process. In good no-disaster states of the world, consumption follows the familiar random walk with drift model. However, with some probability drawn from a Bernoulli distribution calibrated to match the frequency of the events investigated earlier, a disaster 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 the disaster using our estimated LP coefficients in the previous section, and we also adjust its severity by taking a draw calibrated to the latent process ζ in the previous section. Once on a disaster path, the economy remains on it until horizon H_{max} is encountered and it returns to the no-disaster state, or until another disaster draw is encountered.

Thus, the consumer faces uncertainty from a variety of quarters. First, at any point a disaster may take place. Second, given that a disaster takes place, its type and severity is uncertain. Third, the duration of the disaster itself is uncertain. Faced with these multiple sources of uncertainty, the consumer will be willing to pay a price to insure the stability of its consumption stream. How much consumption is she willing to forgo? This is big question we try to answer.

Model simulations To do this, 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. As described, consumption follows a two-state process given by:

$$\begin{aligned} \Delta c_t &= \hat{\mu} + \epsilon_t, && \text{No Disaster Period, } h > H_{max} \\ \Delta c_t &= \hat{\mu} + \left(\sum_{h=1}^{H(r)+1} [\hat{\pi}_h^f F_{t(r)} + \hat{\pi}_h^n N_{t(r)}] \right) e^{\zeta_{t(r)}} + \epsilon_t && \text{Disaster Period, } 1 \leq h \leq H_{max} \end{aligned}$$

The estimates $\hat{\mu}$, $\hat{\pi}_h^f$, and $\hat{\pi}_h^n$ come from the estimates reported in [Table 5](#) in column (4). $\zeta_{t(r)}$ is drawn from $N(0, \hat{\sigma}_\zeta^2)$ once per recession cycle at $t(r)$, where $\hat{\sigma}_\zeta^2$ also comes from [Table 5](#). Similarly, $\epsilon_t \sim N(0, \hat{\sigma}_\epsilon^2)$. The dummies indicating when there is a recession (normal or financial) are $R_t = \{N_t, F_t\}$ just as in the empirical model. What remains is to specify the Bernoulli draw probabilities: the top-level probability draw for entering a disaster period, and the subdraw for whether the type is normal or financial.

Our baseline probabilities are chosen as follows, based on empirical frequency in the data in our full sample, excluding wars:

- **Disaster event:** using earlier notation, $p = P(d_t^0 = 1 | d_{t-1}^0 = 0) = xx.x\%$, that is, the baseline probability of a disaster today conditional on no disaster the previous year.
- **Peak type:** baseline probability given that a disaster occurs, $d_t^0 = 1$, what type will it be. With probability $q_f = 19.3\%$ it will be a financial recession (type F) and with probability $q_n = 81.7\%$ it will be a normal recession (type 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:* $q_f = 0, q_n = 1$, approximate empirical frequencies of the 1950s–60s era.
- *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.

These results will provide guidance as to the welfare implications of varying financial crisis risk.

This completely specifies the simulated consumption process. Simulation begins at time $t = 0$ and $h = 0$ 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, h is reset to 1, and ζ is drawn afresh for the newly-starting Disaster Period. Initially, and at any time where $h > H_{max}$, we draw from the No Disaster Period process, otherwise the Disaster Period process. To recover the path in log levels, c_t , we simply cumulate growth draws, and then we convert to levels.⁹

To afford welfare comparisons with other benchmarks from the literature, we complement our simulated models above with two additional simulated models of consumption growth: a deterministic trend, $c_t = \hat{\mu}^0 t$, to be used as a welfare baseline; and a ‘‘Lucas’’ path, $c_t = \hat{\mu}^0 t + \epsilon_t^0$, which is a deterministic trend plus i.i.d. Gaussian shocks. Formally, we denote these

$$\begin{aligned} \Delta c_t &= \hat{\mu}^0, && \text{Deterministic Trend} \\ \Delta c_t &= \hat{\mu}^0 + \epsilon_t^0 - \epsilon_{t-1}^0, && \text{Lucas Experiment} \end{aligned}$$

where $\hat{\mu}_0$ refers to the unconditional annual growth rate in our sample.¹⁰ The variance of the Lucas process, $\text{Var}(\epsilon_t^0 - \epsilon_{t-1}^0)$ is normalized to match the unconditional variance of consumption growth in the data. We also consider a minor modification, an Adjusted Lucas experiment, in which we assume that the variance equals to $\hat{\sigma}_\epsilon^2$, i.e., the remaining consumption growth variance once variable recession paths are controlled for.

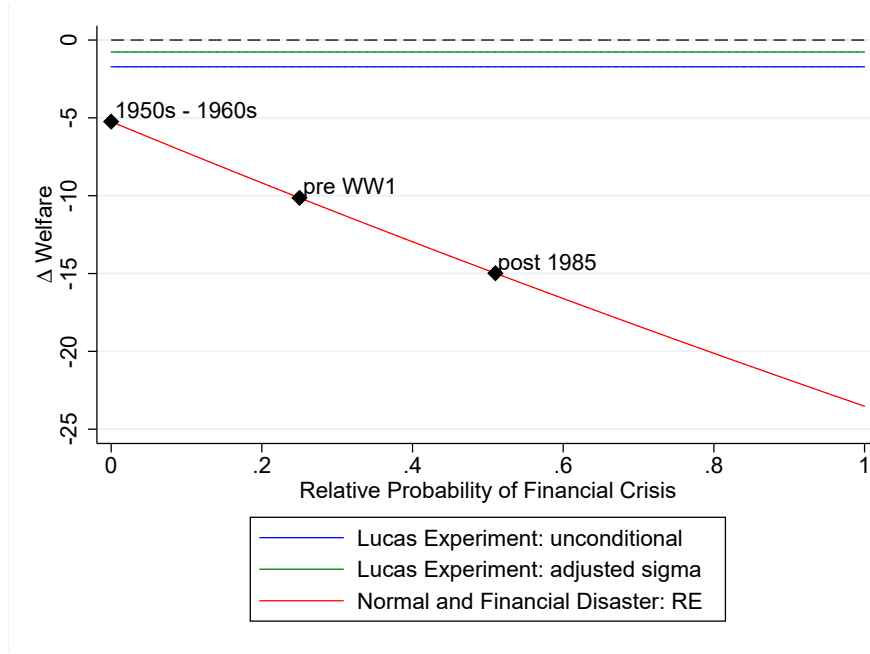
In terms of the numerical implementation, we simulate all series over $T = 5,000$ years for each of the five models $m \in \{\text{Financial, Normal, Lucas unconditional, Lucas adjusted, deterministic}\}$. We then randomly choose $S = 1,000$ starting points based on which we compute 1,000 distinct consumption series $c_{m,s,t}$ for each model. We then obtain a welfare estimate for each model by averaging the discounted lifetime utility for all the related 1,000 sub-series. We assume CRRA preferences with risk aversion parameter $\theta = 2$ and a discount factor of $\beta = 0.96$ for annual data. By today’s standards the choice of $\theta = 2$ may seem conservative, but it is midway between the values picked in seminal papers, namely $\theta = 1$ (Lucas, 1987, 2003) and $\theta = 4$ (Barro, 2006, 2009).

Note that, In 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.

⁹In fact, we implement our simulation in a simpler way. To capture extreme cases, and even all cases between, we actually start by simulating only two consumption growth series: one in which all disasters are financial type recessions ($q_f = 1, q_n = 0$) and one in which all disasters are normal type recessions ($q_f = 0, q_n = 1$). Then any other sub-draw probability ($0 < q_n, q_f < 1$) can easily be simulated by switching between these two extreme growth process, for any sub-draw probabilities which we wish to assign, including the empirical frequencies above, and then re-cumulating to recover levels as necessary.

¹⁰The unconditional mean is lower than the conditional mean $\hat{\mu}$ which abstracts from recession periods.

Figure 6: Simulation: Main Results



Notes: The graph portrays welfare losses associated with the disaster approach and the Lucas experiment relative to the deterministic trend in %. The recession simulation is based on the annual growth rate regression with random effects and $H=10$. We provide two estimates for the Lucas experiment. In the unconditional experiment (blue), the variance of the disturbance term matches the unconditional variance of consumption growth in the data. The second Lucas experiment (green) assumes that the variance equals the remaining consumption growth variance once variable recession paths are controlled for. The black dots refer to observed relative financial crisis probabilities in the data during various periods.

However, as we show momentarily, even in our CRRA environment with a conservative choice of $\theta = 2$, our set up manages to generate considerable welfare losses by being careful in modeling the types of recession consumers face in practice: in our setup, disaster-style fat tails are a feature of all recessions, and so weigh much more heavily in the full welfare cost.

Thus, the welfare level for model m is given by

$$\text{Welfare}_m = \frac{\sum_{j=1}^S \sum_{t=s(j)}^T \beta^{t-s} \frac{C_{m,s,t}^{1-\theta}}{1-\theta}}{S},$$

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

Comparing welfare under actual and counterfactual histories Our main result is shown in Figure 6. Using the deterministic model as a baseline, the vertical axis shows welfare relative to the baseline, in percent, computed as $\Delta \text{Welfare} = (\text{Welfare}/\text{Welfare}_{\text{deterministic}}) - 1$. The horizontal axis shows the relative probability of a financial crisis for the subdraw f which varies between zero and one. The reference level for the deterministic model is the thin dotted line at $\Delta \text{Welfare} = 0$.

Consider first the two Lucas models. Welfare losses are shown by the dashed blue (unconditional) and dotted green (adjusted) lines. These deliver welfare losses of a fixed amount, of about 1 to 2 percent, in line with the prior literature once the more volatile histories of consumption growth for

samples that are broader than just the tranquil post-WW2 United States period.

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 5 and 25 percent. The line shows how the variable financial crisis disaster frequency f impacts the welfare losses relative to the deterministic counterfactual. Three particular points on the line are picked out with black dots 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.¹¹

These calculations are staggering, given that we are using a CRRA parameter of just $\theta = 2$. They speak not only to the massive welfare losses associated with frequent financial crises, but also to the hitherto ignored but nontrivial welfare losses associated with normal recessions. The latter effects have not been captured by traditional models Barro (2006); Lucas (1987, 2003) which treat normal recessions as deriving from Gaussian draws. Instead, as our LP tests have shown, even normal recessions are non-Gaussian, as indicated by significant fat tails in the consumption path relative to the Gaussian null. Figure 6 shows that, in welfare terms, this really matters.

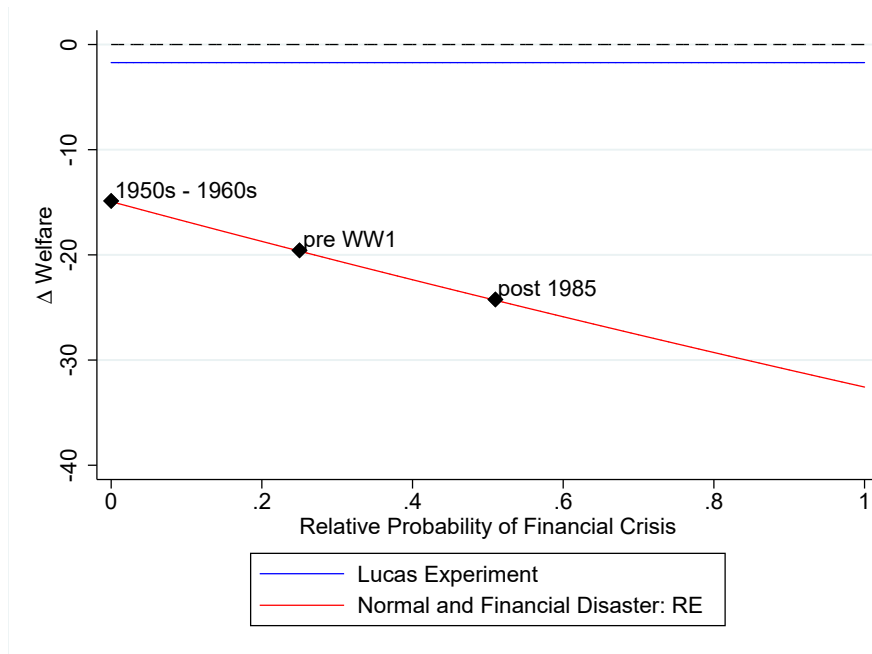
Why? A representative agent living in a world of high financial crisis risk ($q_f = 0.5$, post-1985) would enjoy a welfare gain of 18% from moving to the deterministic path ($1/0.85$), up to 18 times the Lucas estimates. In a world of medium financial crisis risk ($q_f = 0.2$, pre-WW1) they would gain 11% from moving to the deterministic path ($1/0.90$), up to 10 times the Lucas estimates. But even with *zero* financial crisis risk ($q_f = 0.00$, 1950s–1960s) they would gain 5% from moving to the deterministic path ($1/0.95$), up to 5 times the Lucas estimates

We also provide some robustness checks to show how these findings are affected by the choice of local projection estimation technique. Figure 8 shows that switching from our RCLP estimates in changes to our RCLP estimates in levels makes the welfare losses even larger, by a factor of two or more. Now even with zero financial crisis risk ($q_f = 0.00$, 1950s–1960s), the agent would gain 17% from moving to the deterministic path, and is 33% in the high financial crisis risk case ($q_f = 0.5$, post-1985). These high values obtain even though we are again using a CRRA parameter of just $\theta = 2$.

Figure 7 shows that switching to the FCLP fixed coefficient estimates would substantially lower the welfare losses, but this is not a surprise and is inappropriate given that Barro (2006, 2009) demonstrated the importance of accounting for variable-sized disasters so as not to underestimate their importance. However, shifting to different variants of the RCLP model associated with different horizon cutoffs H_{max} makes very little difference.

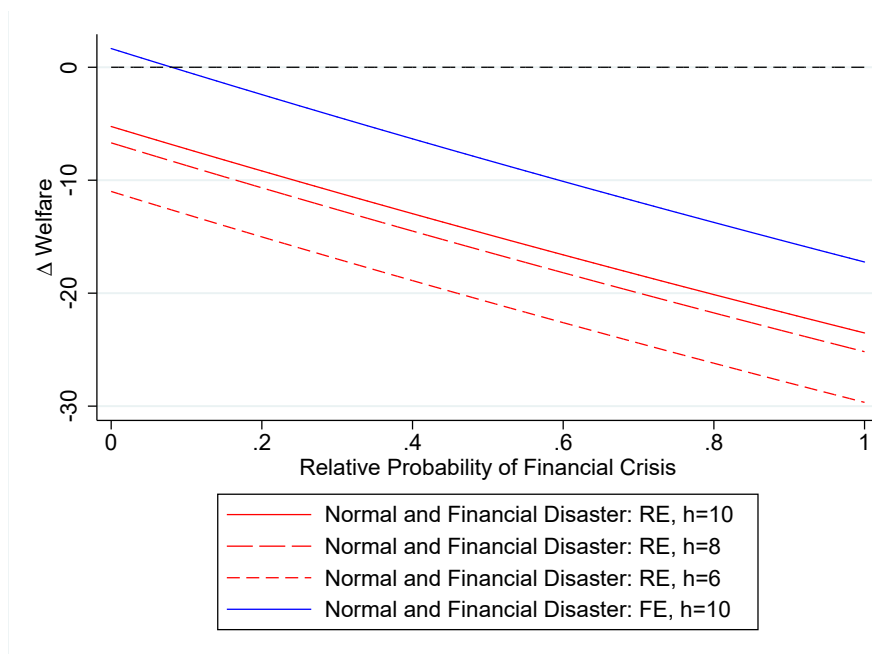
¹¹For example, recessions were more frequent before WW1 and less frequent after WW2, compared to the full sample, which would pull down the middle dot, and life up the left and right dots somewhat, if the values of p were also adjusted.

Figure 7: Simulation: Level Regression



Notes: The graph presents welfare losses associated with the disaster approach relative to the deterministic trend in %. The recession simulation is based on the level regression with random effects and $H=10$. The blue line corresponds to the baseline Lucas experiment. The black dots refer to observed relative financial crisis probabilities in the data during various periods.

Figure 8: Simulation: RE vs FE, different H



Notes: The graph provides robustness checks. It shows welfare losses associated with the disaster approach relative to the deterministic trend in % for various specifications of the growth rate regression. The blue line refers to a deterministic FE-LP with $H=10$. The red lines include random effects with $H = \{6, 8, 10\}$.

5. CONCLUSIONS

At the time of writing this paper, the world's economies are experiencing one of the largest and most sudden declines in stock markets and output due to the outbreak of the coronavirus epidemic—the type of disaster that Barro (2006, 2009) had in mind in justifying the large equity premiums observed in the data and, parenthetically, the most obvious justification for stabilization policy.

However, our argument in this paper is that such extreme and costly events are not the only reason for stabilization policy. This is because disasters are everywhere. We do not live in Gaussian world. All business cycles are highly asymmetric and resemble mini “disasters”, to put it in the language of the asset pricing literature. Consumers experience considerable welfare losses from peacetime recessions, only some of which are associated with financial crises. Because the depth and duration of recessions are unknown and cause deviations from trend growth that can last for extended periods, households would be happy to insure against them.

The paper re-calculates the costs of business cycles in this setting of frequent fat tails. The size of our the welfare loss that we report – up to 15% of consumption for cycles in the past three decades – is magnitudes above Lucas' original calculations. Moreover, the welfare gains have increased in recent decades of more frequent financial recessions. Substantial gains in welfare can be made from better stabilization policies. More than ever, depression prevention and stabilization policies are central to the discipline of macroeconomics.

REFERENCES

- Acemoglu, Daron, and Andrew Scott. 1994. Asymmetries in the Cyclical Behaviour of UK Labour Markets. *Economic Journal* 104(427): 1303–1323.
- Atkeson, Andrew, and Christopher Phelan. 1994. Reconsidering the Costs of Business Cycles with Incomplete Markets. *NBER Macroeconomics Annual* 9: 187–207.
- Barro, Robert J. 2006. Rare Disasters and Asset Markets in the Twentieth Century. *The Quarterly Journal of Economics* 121(3): 823–866.
- Barro, Robert J. 2009. Rare Disasters, Asset Prices, and Welfare Costs. *The 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. Front matter to “Cyclical Analysis of Time Series: Selected Procedures and Computer Programs”. In *Cyclical analysis of time series: Selected procedures and computer programs*, 13–2. NBER.
- Busch, Christopher, David Domeij, Fatih Guvenen, and Rocio Madera. 2018. Asymmetric Business-Cycle Risk and Social Insurance. NBER Working Papers 24569, National Bureau of Economic Research, Inc.
- 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. *The Review of Financial Studies* 29(8): 2069–2109.
- Cox, David R. 1972. The statistical analysis of dependencies in point processes. *Stochastic Point Processes*. Wiley: New York 55–66.
- Cox, D.R., and V. Isham. 1980. *Point Processes*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis.
- 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.

- de Long, 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 Papers 26351, National Bureau of Economic Research, Inc.
- 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* 1127–1162.
- Fatás, Antonio, and Ilian Mihov. 2013. Recoveries. CEPR Discussion Papers 9551, C.E.P.R. Discussion Papers.
- 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. *The American Economic Review* 102(6): 2734–2766.
- Hamilton, James D., and Oscar Jorda. 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.
- Jensen, Henrik, Ivan Petrella, Søren Hove Ravn, and Emiliano Santoro. 2020. Leverage and Deepening Business-Cycle Skewness. *American Economic Journal: Macroeconomics* 12(1): 245–281.
- Jordà, Òscar. 2005. Estimation and Inference of Impulse Responses by Local Projections. *The 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*, NBER Chapters, 213–263. National Bureau of Economic Research, Inc.
- Jordà, Òscar, Sanjay R. Singh, and Alan M. Taylor. 2020. The Long-Run Effects of Monetary Policy. Working Paper 26666, National Bureau of Economic Research.
- 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, G., and A. Brandt. 1995. Marked point processes on the real line. the dynamic approach, 1995.
- Lucas, Robert E., Jr. 1987. *Models of business cycles*. Oxford: Basil Blackwell.
- Lucas, Robert E., Jr. 2003. Macroeconomic Priorities. *The American Economic Review* 93(1): 1–14.
- Martin, Ian W. R. 2008. Disasters and the Welfare Cost of Uncertainty. *The American Economic Review* 98(2): 74–78.
- Morley, James, and Jeremy Piger. 2012. The Asymmetric Business Cycle. *The Review of Economics and Statistics* 94(1): 208–221.
- Muir, Tyler. 2017. Financial Crises and Risk Premia. *The 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.
- 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. Working Paper 26565, National Bureau of Economic Research.
- 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.