Cautionary Tales of Fat Tails

Chetan Dave[1], Scott J. Dressler[2,*], Samreen Malik[3]

Abstract

Distributions of GDP fluctuations that exhibit fat tails shed doubt on the suitability of Normal distributions in empirical and theoretical business-cycle analyses. We document: (i) fat tails in US output fluctuations are not a pervading characteristic of the entire post-war sample, and appear in subsamples exhibiting declines in cyclical and trend volatility (e.g., the Great Moderation); (ii) a DSGE environment featuring Normal shocks that match the declines in observed cyclical and trend volatility can explain almost all of the fat-tailed characteristics observed in the data, leaving little support for the role of rare, large shocks delivering fat-tailed distributions.

Keywords: Fat tails, Growth shocks, Real business cycles.

JEL codes: E0, E3

*Corresponding Author

Email addresses: cdave@ualberta.ca (Chetan Dave), scott.dressler@villanova.edu (Scott J. Dressler), samreen.malik@nyu.edu (Samreen Malik)

[1]University of Alberta

[2]Villanova University

[3]New York University Abu Dhabi

1. Introduction

GDP fluctuations exhibit *fat tails* (i.e., *tail risk*, leptokurtosis, or a kurtosis exceeding 3) and are therefore not adequately approximated by a Normal distribution. This stylized fact may be significant given the preponderance of Normality assumptions in both theoretical and empirical macroeconomic modeling. If a model cannot replicate the fat tails observed in the data, then any subsequent analysis pertaining to everything from parameter estimates to welfare assessments may be biased.[4] The general consensus on this issue is a notion of *fat tails in - fat tails out*, meaning that since models driven by Normal shocks are incapable of endogenously generating fat tails as observed in the data, then the underlying shocks themselves must be drawn from fat-tailed distributions. Cúrdia et al. (2014) and Ascari et al. (2015) show that fat-tailed shocks at the aggregate/macro level improve the performance and empirical fit of standard macroeconomic models. Acemoglu et al. (2017) and Atalay et al. (2018) show that sectoral/micro changes help explain aggregate tail risk at the macro level in theory, but the micro shocks themselves still need to be fat-tailed in their quantitative analysis.

We provide cautionary tales, both empirical and theoretical, to the *fat tails in - fat tails out* notion. Empirically, we report slow-moving (i.e., low-frequency) changes in the fat-tails characteristics of cyclical GDP throughout the sample and show that fat tails are not a pervasive feature of the data. In particular, the only significant instances of fat tails in episodes of GDP fluctuations roughly coincide with slow-moving declines in economic volatility at both the cyclical *and* trend levels observed during the Great Moderation. The observation of low-frequency changes in the volatility of the trend component of output is novel to DSGE literatures considering cyclical

[4]These observations as well as their implications have been documented by Kim and Nelson (1999); Blanchard and Simon (2001); Stock and Watson (2002); Fagiolo et al. (2008), and several others. See Christiano (2007) for arguments against the assumption of a Gaussian likelihood function in empirical VAR analyses, and Mishkin (2011) for arguments against the use of Gaussian shocks in quantitative studies of optimal monetary policy.

characteristics.[5] Ultimately, while slow-moving changes in cyclical kurtosis are consistent with the existence of rare and large shocks, they might also be due to the slow-moving changes in cyclical and trend volatility. Theoretically, we employ a traditional DSGE model to determine how much of these observed declines in economic volatility in both cycle and trend shocks can explain the fat tails observed in the full, post-war sample and the slow-moving changes observed throughout.

Our results suggest that the observed slow-moving changes in both cyclical and trend volatility can account for *more than* all of the fat tails in cyclical output fluctuations observed in the full sample as well as a great deal of the slow-moving changes observed throughout. In addition, we show that while the slow-moving decline in cyclical output volatility observed during the Great Moderation is responsible for the spike in cyclical kurtosis observed during the episode, adding the slow-moving declines in trend output volatility bring the predictions of the model closer to the data. The fact that our model-simulated data displays more kurtosis than observed is not surprising considering that the model lacks frictions, and any model lacking internal shock propagation should pass the characteristics of the exogenous shocks to the endogenous variables. What is surprising is that considering changes in the volatility of the underlying cyclical or trend shocks is enough to let exogenous, Normal shocks deliver endogenous, non-Normal distributions similar to those observed in the post-war US.

We believe the significant contributions of this paper are threefold. First is our documentation of slow-moving changes in the fat-tails characteristics of US GDP fluctuations at the business-cycle frequency. This empirical observation must now be accounted for in any explanation of

[5] Grazzini and Massaro (2022) show that the observed decline in the cycle component of economic volatility during the Great Moderation can be explained by shifting proportions and varying standard deviations of various US industries in a multi-sector framework, but they do not examine how these changes impact overall kurtosis. Cúrdia et al. (2014) only consider the low-frequency changes in the cyclical component of economic volatility in their model estimations.

fat tails in aggregate variables. In particular, any explanation for excess kurtosis must explain the Great Moderation *episode*, which is not associated with large, infrequent shocks. Our result that the decline in volatility (at the cycle and trend level) can account for the observed fat tails during the Great Moderation as well as in the full, post-war sample, leaves little support for the role of rare and large shocks in delivering non-Normal business-cycle fluctuations. Second, this analysis suggests that significant changes to underlying *trend* shocks can in general have a significant impact on the characteristics of *cyclical* observations. This notion is related to an extensive literature (e.g., Kahn and Rich, 2007) suggesting that short-run shocks can distort the impact of long-run shocks, but our results suggest that the opposite can occur as well. In other words, an analysis of extreme events that might alter long-run economic trends (e.g., the Great Moderation, the Great Recession, the COVID Pandemic, etc.) might still need to account for any impact on trend even if one is only interested in the economic impact on the cycle. While many RBC analyses simply decompose a time series into cycle and trend and discard the latter, our analysis suggests that the trend information as well as the detrending method itself might be important when considering extreme events. Third, since we show that significant interplay between cycle and trend in macroeconomic data occurs mainly during *extreme* events, researchers can rest assured that traditional DSGE environments relying on Normal disturbances remain valid macroeconomic tools when the question under analysis can reasonably assume constant macroeconomic volatility and economic growth.

We acknowledge that our *cautionary tales* are nuanced and do not deny the existence of large (non-Normal) business-cycle shocks at the macroeconomic level. Regarding the literature that resorts to exogenous fat-tailed shocks at either the micro or macro level, one can argue that the type of cyclical shocks considered in this paper (i.e., Normally distributed with a changing standard deviation) essentially comprise a fat-tailed distribution, and that the trend shocks themselves are direct results of rare and large events.[6] While we are sympathetic to these arguments, we stress

[6]Regarding the first issue, the t-distribution is a Standard Normal divided by the square root of a chi-squared

that the slow-moving changes in kurtosis documented here lend more support to explanations involving sectoral/micro changes in the environment over large shocks pervasive throughout the sample. Furthermore, there is a related literature devoted to documenting possible endogenous sources of fat-tailed shocks. For example, fiscal multipliers are considered in Auerbach and Gorodnichenko (2012) and adaptive learning is a contender as in Benhabib and Dave (2014), Dave and Tsang (2014), and Dave and Malik (2017). Alternatively, retaining rational expectations but making an allowance for sunspots is considered by Ascari et al. (2019) and Dave and Sorge (2020, 2021) as a source of fat tails. Here, we take no stand on the source of these tail events, but our results may help the literature regarding the timing of any endogenous tail events as well as their impact on both economic cycles and trends.

The rest of the paper is organized as follows. We describe our data and empirical observations in Section 2. These observations, drawn using multiple measures of fat tails and detrending methods from the latest time-series econometrics literature, act as inputs as well as targets in the quantitative exercises described in Section 3. Section 4 concludes.

2. Data

This section presents some new and old stylized facts regarding fat tails in US output. Our data is quarterly, real GDP per capita for the US from 1948:Q1 to 2019:Q4. Selecting output measured in per capita terms allows us to remove any impact of population growth on our empirical volatility and kurtosis calculations, and saves us from needing to include population growth in our theoretical model. Selecting our end date before the tumultuous COVID episode is intentional for several reasons. Primarily, comparability to the literature is a main focus for our calibration exercise, and our goal is to evaluate if our theoretical model can deliver fat-tailed distributions as observed in the existing (pre-COVID) literature without the need for large shocks or tail events.

---

variable. Taking the inverse of the root of the chi-squared variable as the standard deviation, implies that the t-distribution can be interpreted as a Normal with a changing variance.

On a technical side, if the COVID episode were to be included, it would appear at the tail end of our sample and lead to problems regarding the fit of several detrending filters and the reliability of the filtered series themselves.[7]

Decomposing Trend and Cycle. We consider two methods of detrending our data in order to both replicate and add to the stylized facts regarding fat-tailed characteristics in US output. First, we employ the well-known Hodrick-Prescott (HP) filter because it is commonly used and therefore essential for our results to be comparable with earlier studies. Second, we employ an unobserved components model developed by Grant and Chan (2017) which allows for the cycle and trend components of the data to evolve over time. Grant and Chan (2017) show that this filter nests the HP filter as a special case, and that this special case is strongly rejected by the data.[8]

To briefly compare these two filters, consider the decomposition of the log of real GDP per capita y_t into a stationary cyclical component c_t and a nonstationary trend component τ_t:

$$y_t = c_t + \tau_t. \tag{1}$$

The cyclical component can be modeled as a zero mean stationary AR(p) process, whereas the

[7]It is well known that the two-sided, Hodrick-Prescott filter performs worse at the ends of a sample than in the middle. Having an extreme event such as the COVID pandemic appear at the very end of a sample might therefore lead to unreliable results. For example, we show below that the (moment-based) kurtosis of the cyclical component of output under the HP filter is 3.86 in our truncated data. This measure increases by over 50% to 5.84 when the three years including the COVID episode is included.

[8]We thank a referee for pointing out issues with the HP filter such as the imposed link between cycle and trend volatility, and suggesting the use of an unobserved components filter. The filter developed by Grant and Chan (2017) is the latest unobserved components method in the trend-cycle decomposition econometrics literature.

trend can be modeled as a second order Markov process:

$$c_t = \phi_1 c_{t-1} + \cdots + \phi_p c_{t-p} + u_t^c, \tag{2}$$

$$\Delta\tau_t = \Delta\tau_{t-1} + u_t^\tau, \tag{3}$$

where Δ is the first-difference operator such that $\Delta x_t = x_t - x_{t-1}$. Finally, it is assumed that the innovations u_t^c and u_t^τ are jointly Normal.

$$\begin{pmatrix} u_t^c \\ u_t^\tau \end{pmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_c^2 & \rho\sigma_c\sigma_\tau \\ \rho\sigma_c\sigma_\tau & \sigma_\tau^2 \end{pmatrix}\right] \tag{4}$$

Grant and Chan (2017) call this a UCUR-2M filter because it formally captures the correlated unobserved components (UCUR) model of Morley et al. (2003) and allows for a second-order, nonconstant trend growth rate. They further show that a traditional HP trend amounts to the posterior mean of τ in the above system by restricting $\rho = 0$, $\boldsymbol{\phi} = \mathbf{0}$, and $\lambda = \sigma_c^2/\sigma_\tau^2 = 1600$ (for quarterly data).

Table 1: Cyclical Characteristics of US GDP (1948:Q1-2019:Q4)

	HP Filter	UCUR-2M Filter
Std. Dev.	0.0161	0.0247
K(M)	3.86*	2.94
K(Q)	3.52*	2.81
Tailedness	2.05*	1.87*

Note: * indicates rejection of the null hypothesis that the statistic exceeds that of a Normal distribution with 90% confidence

Full-Sample Results. The cyclical characteristics of our full sample are presented in Table 1. The first row presents the full-sample standard deviation of the cyclical component, and shows that the standard deviation under the UCUR-2M filter is over 50% larger than under the HP filter. The moment-based kurtosis results ($K(M)$, second row) suggest the HP filtered series exhibits significant excess kurtosis beyond that of a Normal distribution (>3) and agrees with the related

literature, while the UCUR-2M filtered series fails to report a kurtosis above that of a Normal distribution. The relatively larger average volatility reported using the UCUR-2M filter is consistent with Grant and Chan (2017) who show that the US output gap estimates are larger under UCUR-2M than under HP filtering. Our results further this observation to show that the UCUR-2M filtered series also delivers a relatively smaller average kurtosis due to the series possessing fewer large deviations.

These conflicting results on fat tails point to two shortcomings with traditional, moment-based kurtosis. First, moment-based kurtosis, $K(M) = [(y - \mu)/\sigma]^4$, is inversely related to the standard deviation. This partially explains why the higher standard deviation observed under the UCUR-2M filter is paired with a lower value of $K(M)$. Second, moment-based kurtosis jointly reports the "peakedness" of a distribution (how concentrated data points are around the mean) as well as its "tailedness" (how extreme values are present in the tails) in a single statistic. This implies that the large kurtosis measure reported under the HP filter might be indicating excess peakedness due to a low standard deviation rather than the existence of fat tails.

The remaining rows of Table 1 report results using robust, quantile-based measures of kurtosis and tailedness proposed by Liu (2019).[9] Liu (2019) shows how the quantile-based kurtosis ($K(Q)$, third row of the table) can be expressed as the product of a distribution's quantile-based tailedness (fourth row of the table) and peakedness (not shown). These rows report that while the two filters differ in quantile-based kurtosis being significantly greater than that of a Normal distribution, both filters report that the cyclical component of US output displays significantly more tailedness than a Normal distribution.[10] These results suggest that tailedness is a robust feature of US per-

[9]We thank a referee for pointing out these alternative, quantile-based measures of fat tails that allowed us to present a richer set of stylized facts to examine in our theoretical model.

[10]We refer the interested reader to Liu (2019) for details on this quantile-based measure of kurtosis. Our selection of the specific quantile levels used for our results defines tails to be in the extreme 5% of the distribution and follow Liu (2019). Statistical significance is taken from boot-strapped simulation tables reported by Schmid and Trede (2003).

capita output across multiple filtering methods.

Rolling-Windows Results. While the analysis thus far has restricted attention to full-sample characteristics of US output, Cúrdia et al. (2014) and others have previously documented slow-moving variation in the standard deviation of US output throughout the sample. We replicate the slow-moving variation in standard deviation observed in our data as well as extend this analysis to consider slow-moving variation in our measures of fat tails by considering the following rolling-windows analysis. Starting with the filtered, cyclical component of output, we select an observation size of $n = 60$ and construct overlapping subsamples of the data where the first subsample contains observations 1 through 60, the second contains observations 2 through 61, and so on. This results in 229 $(288 - n + 1)$ subsets for each of our filtered series, where each pair of consecutive subsets only differ by the beginning and ending observations.[11] We then record the same characteristics reported in Table 1 for each subsample. These characteristics are illustrated in Figure 1 at the first observation (i.e., start date) of each rolling-window subset (e.g., the first subset considers observations from 1948Q1 through 1962Q4 and the results appear on the timeline at the start date of 1948:1).

Before detailing the results of our rolling-windows analysis, we stress that this exercise is an *informal* way of illustrating the evolution of our cyclical characteristics over the full sample. For example, note that two consecutive subsamples differ in only one observation at the beginning / end of the range and share in the remaining 59 observations. Therefore, if one were to observe a large drop in the standard deviation of output between two consecutive subsamples, then one can attribute this drop to the replacement of a single observation possessing a relatively large deviation from the mean at the beginning of the first subset with an observation possessing a relatively small deviation from the mean added to the end of the second subset. We explored more formal options for estimating time-varying volatility and kurtosis via the GARCH-K model of

[11] We also consider alternative window sizes of 40 and 80 quarters consisting of 249 and 209 overlapping subsets of data, respectively. Our results presented below are robust to these alternative window sizes.

Brooks et al. (2005). For our HP-filtered data, we found significant conditional heteroskedasticity (e.g., GARCH) effects, but no significant autoregressive kurtosis. Since the GARCH-K only tests for significant first-order autoregressive conditional kurtosis, failure to reject the null could imply a constant kurtosis over time, or a variable kurtosis that is either conditional on past observations of kurtosis beyond the first quarter (i.e., a higher-order GARCH-K model) or not conditional on past observations of kurtosis whatsoever. Since our rolling-window analysis suggests that the kurtosis of our HP-filtered series is not constant, we can conclude that failing to reject the null is due to one of the other two potential outcomes. We were unable to estimate the GARCH-K model using our UCUR-2M filtered data because the data does not meet the model requirement of a full-sample kurtosis being greater than 3 (see Table 1).[12]

The top and bottom panels of Figure 1 illustrate the rolling-window characteristics of our HP and UCUR-2M filtered series respectively. The dotted line in each panel illustrates the slow-moving evolution of the standard deviation of the filtered series, and both series suggest a gradual decline in standard deviation beginning in the early 1970s as well as a steep decline in the early 1980s. These steep declines are particularly associated with drastic increases in the moment-based kurtosis of each filtered series (the solid lines). A dot along each line indicates that the kurtosis of the subsample beginning at that date is significantly greater than that of a Normal distribution with 90% confidence. Although the full, HP-filtered sample possessed a kurtosis significantly greater than 3 and the UCUR-2M filtered series did not, it is interesting to note that both filtered series indicate kurtosis measures significantly greater than 3 in the subsamples occurring around the extreme drop in standard deviation and the recessions of the early 1980s. The remaining two lines in the panels illustrate the slow-moving evolution of quantile-based kurtosis and tailedness across the subsamples. Recall the HP-filtered series exhibited significant

[12]The GARCH-K model establishes a link between the conditional kurtosis of the data and the time-varying degrees of freedom of a t distribution (see Equation 12 of Brooks et al., 2005) and restricts that the conditional kurtosis of the data must always exceed 3 for the degrees of freedom to be positive.

quantile-based kurtosis and tailedness beyond that of a Normal distribution in the full sample. For the HP-filtered subsamples (top panel), significant quantile-based kurtosis is only exhibited in the episode associated with the decline in standard deviation and at the end of the range. Significant tailedness is also exhibited in these regions, along with two other regions earlier in the range. The UCUR-2M filtered series did not exhibit significant quantile-based kurtosis beyond that of a Normal distribution in the full sample but did exhibit significant tailedness. For the UCUR-2M filtered subsamples (bottom panel), the episode associated with declines in rolling-window standard deviation is the only episode where there is significant quantile-based kurtosis as well as tailedness beyond that of a Normal distribution.

Trend Analysis. Our rolling window analysis illustrates that any significant changes in either moment-based or quantile-based kurtosis are associated with large changes in the standard deviation of either the HP or UCUR-2M filtered series. Since our filters differ in how they identify the trend component of the data in order to arrive at the cycle component, we extend this analysis one step further and examine the slow-moving characteristics of the trend. To accomplish this, we take the trend extracted (for example) by the HP-filter, denote it τ_t^{HP}, and estimate an autoregressive model in order to identify the trend shocks.[13]

$$\Delta\tau_t^{HP} = \sum_{i=1}^{4} \rho_i \Delta\tau_{t-i}^{HP} + u_t^{\tau} \tag{5}$$

The standard deviation of the resulting trend shocks (u_t^{τ}) were then examined using our rolling-windows analysis to observe any slow-moving changes. The results for both filters are illustrated along side the previously-illustrated, slow-moving changes in cyclical standard deviation in Figure 2. The results suggest that there is a steep drop in the standard deviation of the trend shocks that occurs slightly before the drop in the standard deviation of the cyclical component for both

[13] A sample partial autocorrelation analysis suggested that four lags were sufficient to uncover a random residual series.

filtered series. While this relationship is not surprising for the HP-filtered series since the HP filter imposes a relationship between the standard deviations of the trend and cycle by imposing a fixed smoothing parameter, it is a surprising result of the UCUR-2M filtered series because that filter makes no such imposition.

Summary. We summarize the key observations reported in this section as follows:

1. With respect to the full-sample, cyclical characteristics observed under the HP filter, we find all measures of fat tails to be significantly larger than that of a Normal distribution. With respect to the full-sample, cyclical characteristics observed under the UCUR-2M filter, the data suggests that tailedness is significantly greater than that of a Normal distribution.
2. Regardless of the full-sample characteristics, we find subsamples of the detrended series that exhibit measures of fat tails to be significantly larger than that of a Normal distribution. This suggests that while not all measures of fat tails are robust to methods of detrending the full sample, there is one particular episode (i.e., the early 1980s) where every measure of fat tails considered reports results significantly different than that of a Normal distribution.
3. The slow-moving increase in fat tails of the cyclical component of US output coincided with steep declines in volatility for both the cycle *and trend* components. This observation is robust to the method of detrending.

This last observation formally motivates the quantitative question of our theoretical DSGE analysis and lays out the discipline of our calibration exercise. Our question is, how much of the fat tails in the data can be explained by the slow-moving declines in cycle and trend volatility? Answering this quantitative question requires a model that can replicate the slow-moving cyclical and trend volatility in output as observed in Figure 2. With the model in hand, we can then analyze the simulated data for full sample and slow-moving fat tails, and compare these predictions with the data. The next section lays out a theoretical environment capable of answering this quantitative question.

3. Model

3.1. Overview

Our model is a standard, neoclassical framework following Pakko (2002) and Fisher (2006). It features shocks to labor-augmenting total-factor productivity (TFP), where the shock process contains both a short-run component that is stationary in levels and a long-run component that is stationary in growth rates. Our model allows us to perform several experiments. First, since our model assumes exogenous growth, we can feed in the trend shocks identified in the data directly into the model in order to assess how they impact the fat tails in the detrended data. Second, we can feed in a series of cyclical (TFP) shocks featuring slow-moving changes in volatility such that the predicted slow-moving changes in output volatility match what is observed in the data, and again assess the impact on cyclical fat tails. Finally, we can include both the trend and cycle shocks to assess the full impact.

3.2. Preferences and Technologies

The preferences of an infinitely-lived representative household are described by the expected utility function

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[ln(C_t) - \Psi H_t \right], \tag{6}$$

where β is the discount factor and C_t and H_t respectively denote consumption and labor hours supplied at time t. Utility is logarithmic in consumption and linear in leisure to make the model consistent with balanced growth.[14]

[14]While some non-separable utility specifications consistent with balanced growth could be used, the specification in (6) follows Hansen (1985) and Rogerson (1988) by assuming that the economy consists of a large number of individual households, each of which includes a potential employee who either works full time or not at all during any given period.

The representative household chooses consumption, labor hours, investment (I_t), and next period's capital stock (K_{t+1}) to maximize (6) subject to a budget constraint and a law of motion for the capital stock.

$$C_t + I_t \leq K_t^{\alpha}(\tilde{Z}_t H_t)^{1-\alpha} \tag{7}$$

$$K_{t+1} \leq (1-\delta)K_t + I_t \tag{8}$$

The TFP shock $\tilde{Z}_t$ in (7) is labor-augmenting in order for balanced growth to be feasible. Technology shocks of this kind are interpreted as disturbances that directly influence general production possibilities such as changes in taxes, regulation, or market structure. The Cobb-Douglas share parameter α in (7) and the depreciation rate of the capital stock δ in (8) both lie between zero and one.

3.3. Equilibrium Allocations

The welfare theorems apply in this model, so the problem of a representative household is the same as that of the social planner: choose contingency plans for C_t, H_t, I_t, and K_{t+1} for all $t = 0, 1, 2, \ldots$ to maximize the utility function (6), subject to the constraints imposed by (7) and (8) for all t. Letting Λ_t denote the non-negative multiplier on the budget constraint (7), and using (8) to remove I_t from the problem, the first-order conditions can be written as

$$\frac{1}{C_t} = \Lambda_t, \tag{9}$$

$$\Psi = \Lambda_t \left[(1-\alpha)K_t^{\alpha} H_t^{-\alpha} \tilde{Z}_t^{1-\alpha}\right], \tag{10}$$

$$\Lambda_t = \beta E_t \Lambda_{t+1} \left[\alpha K_{t+1}^{\alpha-1}(\tilde{Z}_{t+1}H_{t+1})^{1-\alpha} + 1 - \delta\right], \tag{11}$$

and (7) with equality for all t.

Intuitively, while (7) shows how the TFP shock $\tilde{Z}_t$ directly impacts the household's budget constraint, (9) indicates how Λ_t measures the marginal utility of consumption for the household.

3.4. Driving Processes

As stated previously, the TFP shock possesses a short-run component stationary in levels, and a long-run component stationary in growth rates. These components are formally defined as $\tilde{Z}_t = e^{z_t} Z_t$, where z_t and Z_t are respectively the short and long-run components. The short-run component follows an AR(1) process

$$z_t = \rho z_{t-1} + \varepsilon_t, \tag{12}$$

with $|\rho| < 1$ and ε_t representing iid draws from a Normal distribution with zero mean and *time-varying* standard deviation σ_t. The long-run component Z_t represents the cumulative product of growth shocks to technology where

$$Z_t = e^{g_t} Z_{t-1} = \prod_{s=0}^{t} e^{g_s}, \tag{13}$$

$$g_t = (1-\tau)g + \sum_{i=1}^{4} \tau_i g_{t-i} + \nu_t, \tag{14}$$

with ν_t representing iid draws from a Normal distribution with zero mean and standard deviation σ^g. The term g in (14) denotes the long-run growth rate of technology, and $\tau = \sum_{i=1}^{4} \tau_i$.

In the short-run, shocks to both level and trend will impact the model variables. The purpose of the quantitative exercises below is to determine if a short-run data characteristic like fat tails can be attributable to either changing the volatility of the short run shocks alone or changing the volatility of both the short and long-run shocks. If so, which shock sources are most plausible?

In the long run, only shock components stationary in growth rates can account for the non-stationary behavior of the model variables. Output, consumption, and investment must all grow by the same rate according to the budget constraint (7). This rate is simply

$$G_{t-1} = Z_{t-1}, \tag{15}$$

so $c_t = C_t/G_{t-1}$, $i_t = I_t/G_{t-1}$, and $y_t = Y_t/G_{t-1}$ where $Y_t = C_t + I_t$. These transformations can be confirmed to deliver stationary variables along a balanced growth path.

4. Quantitative Analysis and Results

4.1. Overview

We subject our model to three experiments. In our first *Long Run (LR) Only* experiment, we assume no short-run, cyclical shocks and simply feed the trend shocks extracted from the data (via Equation (5)) into our model. This will establish a baseline case delivering the extent to which observed trend shocks can impact short-run output characteristics. In our second *Short Run (SR) Only* experiment, we assume no long-run trend shocks and perform a *dynamic calibration* where we back out the slow-moving standard deviation series of our theoretical TFP shock needed to match the volatility of the cyclical component of output observed in the rolling-windows analysis (see Figure 2). We then simulate the model and examine the implications for the observed volatility on fat-tailed output characteristics. In our third *Long Run and Short Run (LR + SR)* experiment, we repeat the dynamic calibration in the second experiment but add the trend shocks considered in the first experiment. Taken together, these experiments allow for a complete picture of how much the slow-moving changes in (short-run) cycle and (long-run) trend volatility can explain the fat tails observed in output in both the full sample and rolling-windows subsamples. The remainder of this section details the steady-state calibration of the model which is shared by all experiments, followed by more detailed descriptions of each experiment. The section concludes with a presentation and comparison of all model results.

4.2. Calibration

Our calibration strategy assigns parameter values following the business-cycle literature (e.g., Cooley and Hansen, 1989) so the resulting steady state of the model matches particular long-run properties of the US economy. Once these parameters have been determined, they are fixed for all model specifications.

A model period is one quarter of a year. The parameter Ψ is set to 2.71, so a household's average allocation of time to market activity (net of sleep and personal care) is one-third, which is in line with the estimates of Ghez and Becker (1975). The depreciation rate $(\delta = 0.024)$ is set to a 10-percent annual rate, and capital's share of national income is set to $\alpha = 0.36$.

The persistence parameter in the cyclical shock process (Equation (12)) is set to 0.95. For Equation (14), the steady-state growth rate g is set to 2 percent annually, and the four AR coefficients were taken from the estimates obtained in Section 2. Finally, $\beta = 0.99$, so the capital stock to annual output ratio of the model is 2.5. It should be noted that although the value of g impacts the size of the discount parameter β, the quantitative results presented below are not sensitive to the steady-state growth rate.

4.3. Experiments

Experiment 1: Long Run Only (LR)

Our first (LR) experiment shuts down any short-run, cyclical volatility (i.e., $\sigma_t = 0 \ \forall \ t$) and determines how much of the fat tails observed in the cyclical component of real GDP fluctuations can be accounted for by the observed shocks to trend. Since our model assumes exogenous growth, the trend shocks extracted from the data are simply fed into the model. The resulting simulated data is detrended and the cyclical component is compared to the actual data.

Experiment 2: Short Run Only (SR)

Our second (SR) experiment shuts down any long-run, trend volatility (i.e., $\sigma^g = 0$) and determines how much of the fat tails observed in the cyclical component of real GDP fluctuations can be accounted for by the observed, slow-moving changes in its standard deviation. This is accomplished via a *dynamic calibration* of the model where a series of standard deviations for the TFP shock (using rolling windows) is determined so the predicted series of standard deviations of the cyclical component of output matches the target illustrated in Figure 2. This is stated in more detail as follows.

1. Establish a full sample of 288 observations and rolling-window size of 60 quarters as we

have in the data. Establish the vector of 227 different standard deviations of output from our data, one for each 60-quarter subset of our full sample. The elements of this vector serve as the 227 targets of the dynamic calibration. Denote this vector $\boldsymbol{\sigma_Y}$.

2. Select a vector of 288 standard deviations of the TFP shock for the model, one standard deviation for each period of the data, and denote this vector $\boldsymbol{\sigma_z}$.
3. To simulate the model, draw 3000 series of TFP shocks, each 288 quarters long, where each shock series contains iid draws of a Normal distribution with a standard deviation of one. Multiplying each series by $\boldsymbol{\sigma_z}$ element-by-element delivers 3000 draws of a shock process with a time-varying volatility depicted by $\boldsymbol{\sigma_z}$. The model is simulated for each of our 3000 series, and the cyclical characteristics of our detrended output predictions are stored. Denote the vector of 227 different standard deviations of output from our data, one for each 60-quarter subset of our full sample averaged across all 3000 series, as $\hat{\boldsymbol{\sigma}}_{\boldsymbol{Y}}$.
4. The *dynamic* calibration then selects 288 individual standard deviations of the exogenous shock process in $\boldsymbol{\sigma_z}$ to solve the following minimization problem.

$$min(\hat{\boldsymbol{\sigma}}_{\boldsymbol{Y}} - \boldsymbol{\sigma_Y})^T(\hat{\boldsymbol{\sigma}}_{\boldsymbol{Y}} - \boldsymbol{\sigma_Y}) \tag{16}$$

The exercise detailed above is similar to a simulated method of moments (SMM) estimation, but is considered calibration since it does not utilize the inverse of the variance-covariance matrix of the moment conditions as an optimal weighting matrix. It should be noted that while we would have liked to perform this dynamic calibration using both filters, the HP-filter was chosen over the UCUR-2M filter strictly on the grounds of computational feasibility. The UCUR-2M filter, being a Bayesian estimation algorithm, takes upwards of 1.5 minutes to filter a data series, and this filter would need to be applied to each of the 3000 simulated output series every time the model was simulated. This would amount to 4 days of computing time in order to arrive at one $\hat{\boldsymbol{\sigma}}_{\boldsymbol{Y}}$ for each function iteration of our dynamic calibrations algorithm (i.e., for each selection of $\boldsymbol{\sigma_z}$). Since our dynamic calibration algorithm required between 20,000 and 40,000 function iterations when

using the HP-filter, we conservatively estimate that the dynamic calibration algorithm using the UCUR-2M filter would take over 200 years to complete.

Experiment 3: Long Run and Short Run (LR+SR)

Our third (LR+SR) experiment combines the previous two experiments. It determines how much of the fat tails observed in real GDP fluctuations can be accounted for by simultaneously considering long-run, trend shocks and low-frequency changes in the standard deviation of the cycle component. This is accomplished by repeating the dynamic calibration performed in the SR experiment, only now simultaneously feeding in the observed, long-run trend shocks extracted from the data for each simulation of the model.

4.4. Results

Full-Sample Results. The full-sample results from our three experiments are compared with the data in Table 2. As stated above, the objective of the dynamic calibration performed in the SR and LR+SR experiments was to match the slow-moving changes in standard deviation observed in the rolling-windows analysis for HP-filtered output. While the dynamic calibration does not consider the full sample standard deviation as a target, the first row of results indicates that the predicted standard deviation of HP-filtered output in the SR and LR+SR results come very close to the observed standard deviation of output observed in the full sample. The predicted standard deviation of output in the LR results is much smaller than the other results, but the LR experiment only considers growth shocks prior to the data being detrended. Therefore, it is surprising that a detrended data series in the LR experiment exhibits any volatility at all.

Rows two through four of results in the table compares the fat-tailed characteristics of the HP-filtered output data with the predictions from our three experiments. In general, the table reports that the predictions under our SR and LR+SR experiments display levels of moment-based kurtosis, quantile-based kurtosis, and tailedness that significantly exceed that of a Normal distribution and are in-line with the data. While there are unsurprisingly no significant levels of fat tails reported in our LR experiment, it is surprising that combining the long-run and short-

run shocks together in our model (LR+SR) brings the predictions of the model closer to the actual data than when just considering the short-run shocks alone (SR). Note that this applies for the quantile-based measures of kurtosis and tailedness as well as the moment-based measure, so the explanation is not simply the addition of additional (long-run) volatility resulting in less (short-run) kurtosis.

Rows five through eight in the table compare the predictions of the model with the data when using the UCUR-2M filter for detrending. Note that although this filter was not used in the dynamic calibration, there are some interesting takeaways. First, the full-sample standard deviations predicted in the SR and LR+SR experiments are in line with the data similar to the results obtained using the HP filter. Second, while the SR experiment predicts significant fat tails when considering moment-based kurtosis, quantile-based kurtosis, and tailedness, the LR+SR experiment only predicts significant tailedness which coincides with the data. This suggests that considering trend shocks when comparing cyclical characteristics can bring the predictions of the model closer to the data under a variety of detrending filters.

Table 2: Full-Sample Characteristics

HP Filter				
	Data	LR	SR	LR+SR
Std.Dev	0.0161	0.007	0.0166	0.0160
K(M)	3.86*	2.58	6.19*	5.52*
K(Q)	3.52*	2.66	4.44*	4.05*
Tailedness	2.05*	1.59	2.25*	2.16*
UCUR-2M Filter				
	Data	LR	SR	LR+SR
Std.Dev	0.0247	0.0034	0.0256	0.0288
K(M)	2.94	2.71	5.12*	3.46
K(Q)	2.81	2.84	4.09*	3.09
Tailedness	1.87*	1.62	2.15*	1.84*

Note: * indicates rejection of null hypothesis that the statistic exceeds that of a Normal distribution with 90% confidence

Rolling-Windows Results. Analogous to our rolling-windows analysis of the data, we constructed overlapping subsamples of each of our simulated output series and recorded the cyclical characteristics of each subsample detrended using either the HP or UCUR-2M filters. We then averaged these characteristics across the 3000 simulated series to arrive at our results. Our results using the HP filter are compared with the data in Figure 3. The top panel of the figure compares the slow-moving standard deviation of output in the data with that predicted by the model under our three experiments. Since this slow-moving series was the target of our dynamic calibration, this panel is simply a visualization of how well the models were able to hit the target. Note how the model under both the SR and LR+SR experiments was able to match the slow-moving changes in standard deviation observed in the data quite closely, including the rapid decline observed in the early 1980s. This rapid decline allows the model to mimic the dramatic increase in moment-based kurtosis observed in the rolling-windows data (second panel). However, it is again interesting to see that the inclusion of trend shocks (LR+SR experiment) brings the predictions of the model closer to the data. Note that this reduction in moment-based kurtosis is not due to an increase in standard deviation, since the amount of output standard deviation predicted by the SR and LR+SR experiments are almost identical.

The third and fourth panels of Figure 3 illustrate the results for quantile-based kurtosis and tailedness, respectively. Again, the model predictions under the LR+SR experiment outperform those of the SR experiment, specifically when it comes to quantile-based kurtosis. This result is important given that the quantile-based measures of fat tails are robust refinements to the moment-based measures (as discussed in the Data section).

Finally, our rolling-windows results using the UCUR-2M filter are compared with the data in Figure 4. The first panel illustrates that the predicted slow-moving changes in standard deviation for our experiments do not match the data as closely using the UCUR-2M filter as with the HP filter. In particular, the steep drop in standard deviation observed in the data during the early 1980s is not closely captured in any of the model simulations. This failure to closely account for the slow-moving changes in standard deviation directly relates to how and when the model fails

to account for the slow-moving fat tails observed in the data. Panel 2 of the figure illustrates that the moment-based kurtosis predicted under the SR model does best at capturing the spike observed in the data, because the SR model delivers the steepest decline in slow-moving standard deviation at that time. Panels 3 and 4 illustrate that the simulated data under the UCUR-2M filter fails to predict any of the spike in the quantile-based measures of fat tails during the early 1980s. However, it is interesting to note that the LR+SR model does a superior job at matching the slow-moving cyclical characteristics outside of the episode featuring the extreme declines in standard deviation. In particular, the LR+SR model does a surprisingly good job at matching the quantile-based kurtosis and tailedness of the rolling window observations before 1973 and after 1984, and these predictions outperform those of the SR model.

Summary. We summarize the key results reported in this section as follows:

1. Our simulated data suggests that allowing for traditional, Normal shocks with observed, slow-moving changes in standard deviation can account for more than all of the full-sample, fat tail characteristics of detrended output. When considering observed shocks to both cycle and trend, the full-sample characteristics of the simulated data match the those of the observed data remarkably well under using both the HP and UCUR-2M filters.
2. When considering slow-moving (rolling windows) characteristics of the simulated data using the HP-filter, the ability of the model to closely capture the slow-moving changes in the standard deviation of output allows the model to predict the extreme increase in moment-based kurtosis observed in the 1980s. The overprediction made when considering only short-run shocks and ignoring trend shocks is reduced when accounting for trend shocks.
3. Despite the model making decent full-sample predictions when using the UCUR-2M filter, the slow-moving predictions with respect to fat tails under this filter do not match the data and are due to the model being unable to match the slow-moving standard deviation of output.

One key takeaway from these results rest on the comparison of the LR+SR model predictions across the two methods for detrending the data (HP versus UCUR-2M). When using the HP-filter, the LR+SR model predictions closely resemble the data for the full sample and slow-moving fat tail measures, and this is a direct result of closely matching the observed slow-moving standard deviation of output. When using the UCUR-2M filter, the LR+SR model predictions closely resemble the full sample fat tail measures, but not the slow-moving ones, and the latter result is a direct result of the model not closely matching the observed slow-moving standard deviation of output. These particular results suggest that the full-sample fat tail characteristics of the data can be reasonably captured by simply considering slow-moving changes in shock volatility, while the slow-moving fat tails observed in business cycle data is the more ambitious and arduous stylized fact in need of capturing. Our results suggest that sufficiently capturing the slow-moving changes in output standard deviation is necessary to explain the slow-moving changes in fat tails. Until this is accomplished, one should not resort to rare shocks or tail events in any empirical or theoretical analysis.

5. Conclusion

While GDP fluctuations exhibit fat tails, one should proceed with caution when assuming that this is due to rare and large exogenous shocks. Our paper augments the stylized fact of fat tails using several detrending methods, moment and quantile-based measures, and a rolling-window subset analysis, and shows that the only instances of fat tails in US GDP fluctuations between 1948 and 2019 were observed around the Great Moderation. This episode in US history is not known for rare and large shocks (i.e., tail events), but for declines in shock volatility at both the cycle and trend level. We show that a simple DSGE environment matching the observed low-frequency changes in the standard deviation of economic shocks at both the cycle and trend level can account for all of the observed fat tails (and then some) in the full sample as well as in the subsamples. While not intended to represent conclusive evidence against the existence of tail events (e.g., the COVID pandemic), our results leave little room for the notion of *fat tails in - fat*

tails out in aggregate, pre-COVID, US GDP.

Our results deliver three contributions to future macroeconomic analyses. For any theory attempting to explain the existence of fat tails in output, one should consider the implications of slow-moving changes in volatility at both cycle and trend levels, and assess the validity of the explanation by matching the slow-moving changes in kurtosis. For quantitative analyses considering cyclical economies in general, we show that simply discarding the trend may not be innocuous when there are significant changes to trend characteristics. Finally, if one can reasonably assume constant trend characteristics, then our results suggest that the assumption of Normal shocks is valid for empirical and theoretical modeling.

Figure 1: Rolling Windows Analysis

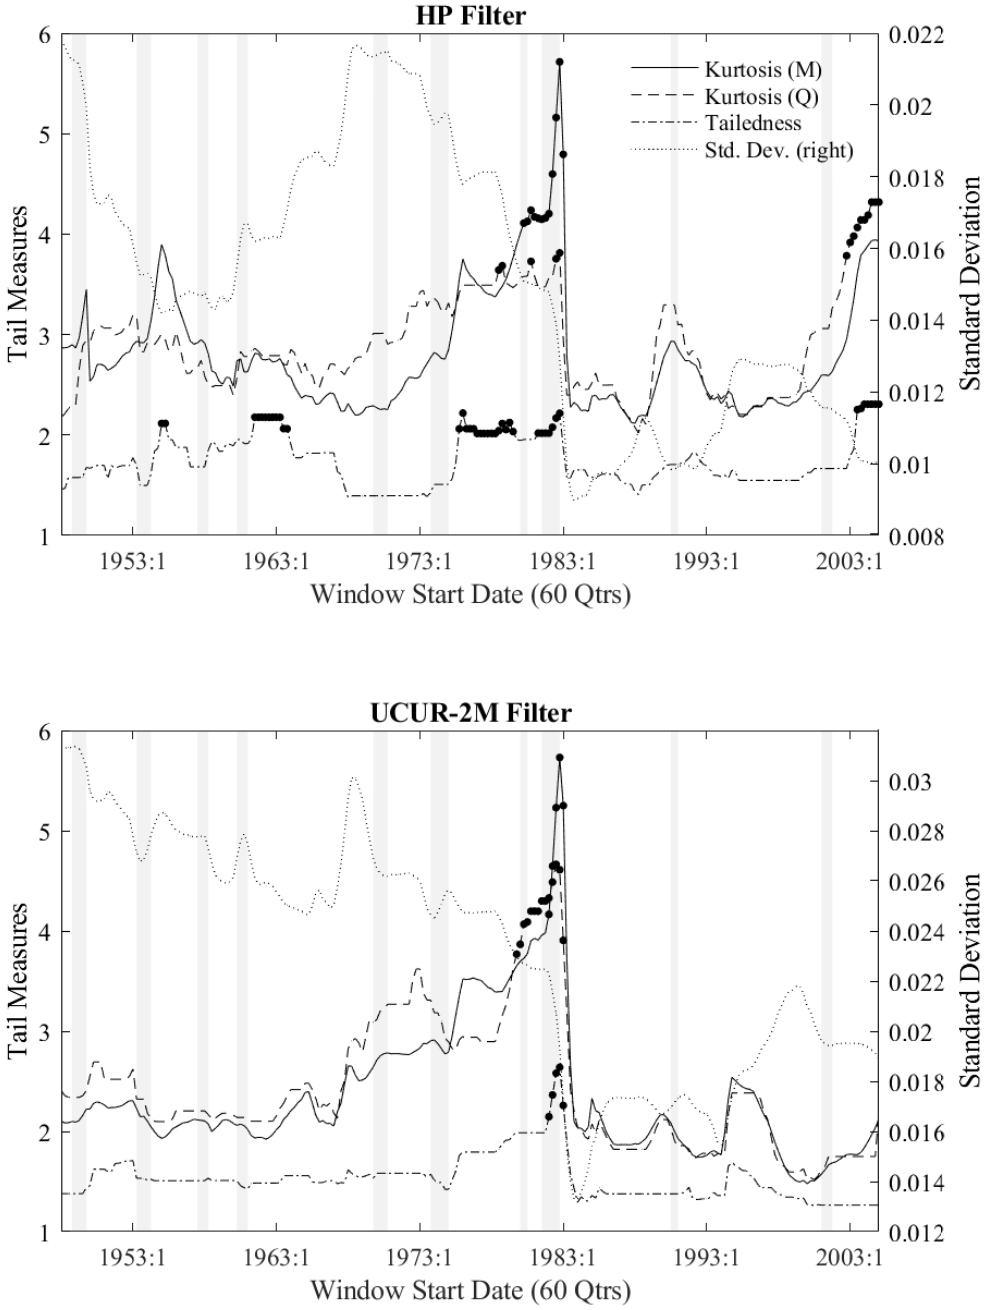

Notes: The full data sample is configured into an exhaustive set of overlapping subsets of 60 observations each. The various lines represent the specific characteristic of the cycle component of logged, quarterly, real GDP per capita (i.e., the data) for each 60-observation subsample at its starting date. A dot indicates that that specific measure is significantly different than that of a Normal distribution with 90% confidence.

Figure 2: Rolling Windows Analysis: Trend and Cycle Volatility

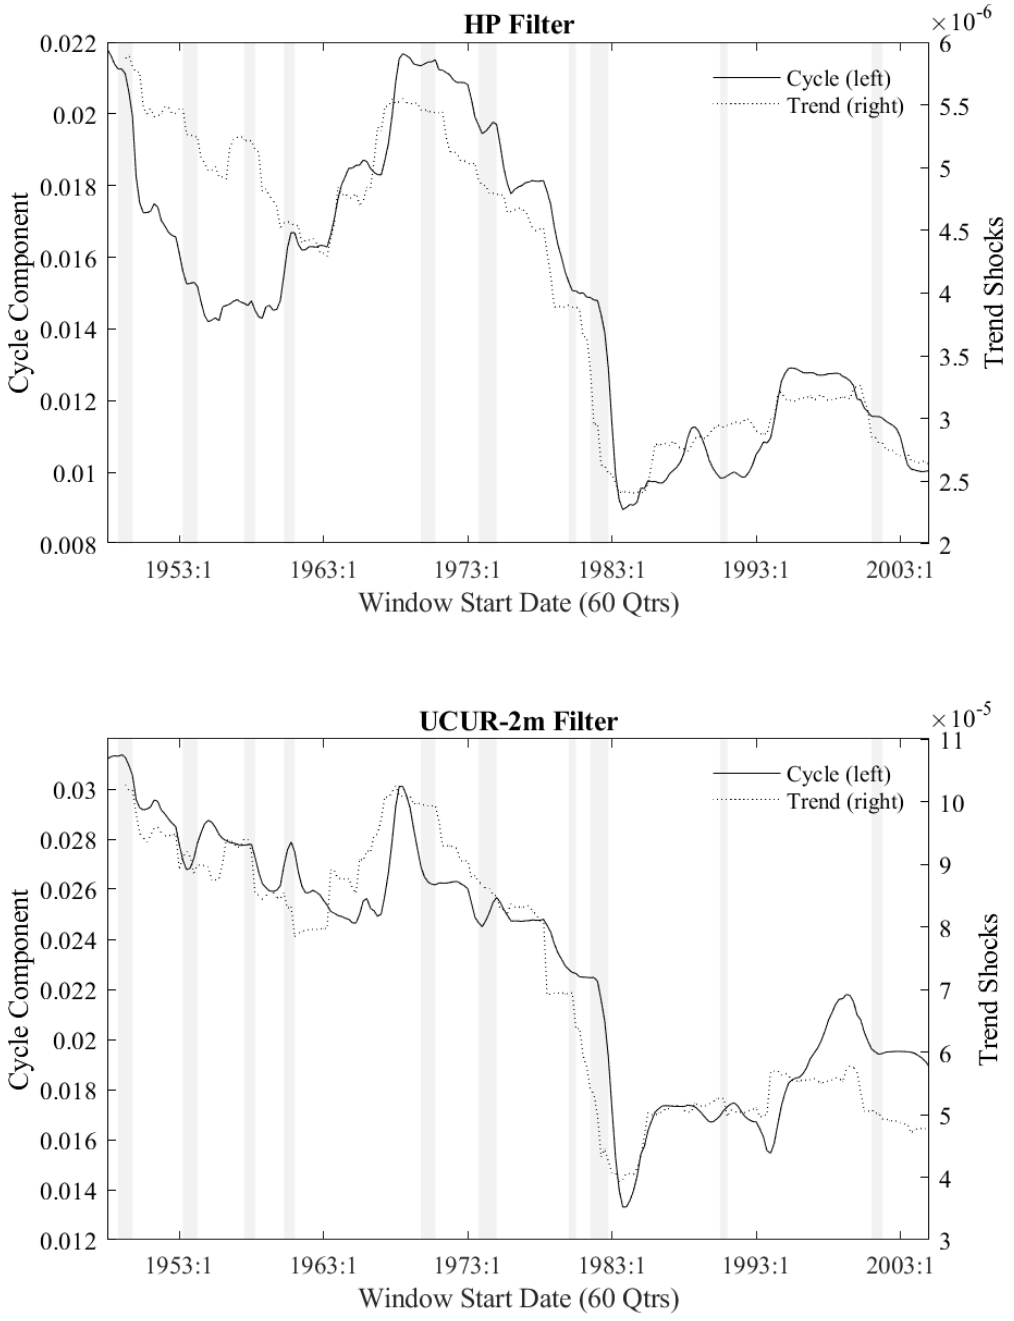

Notes: The full data sample is decomposed into cycle and trend according to the indicated filter. A sample partial autocorrelation analysis is performed on the trend component in order to identify the trend shocks. Both the cyclical component of the data and the trend shocks were configured into an exhaustive set of overlapping subsets of 60 observations each. The various lines represent the specific characteristic for each 60-observation subsample at its starting date.

Figure 3: Rolling Windows Analysis (HP Filter)

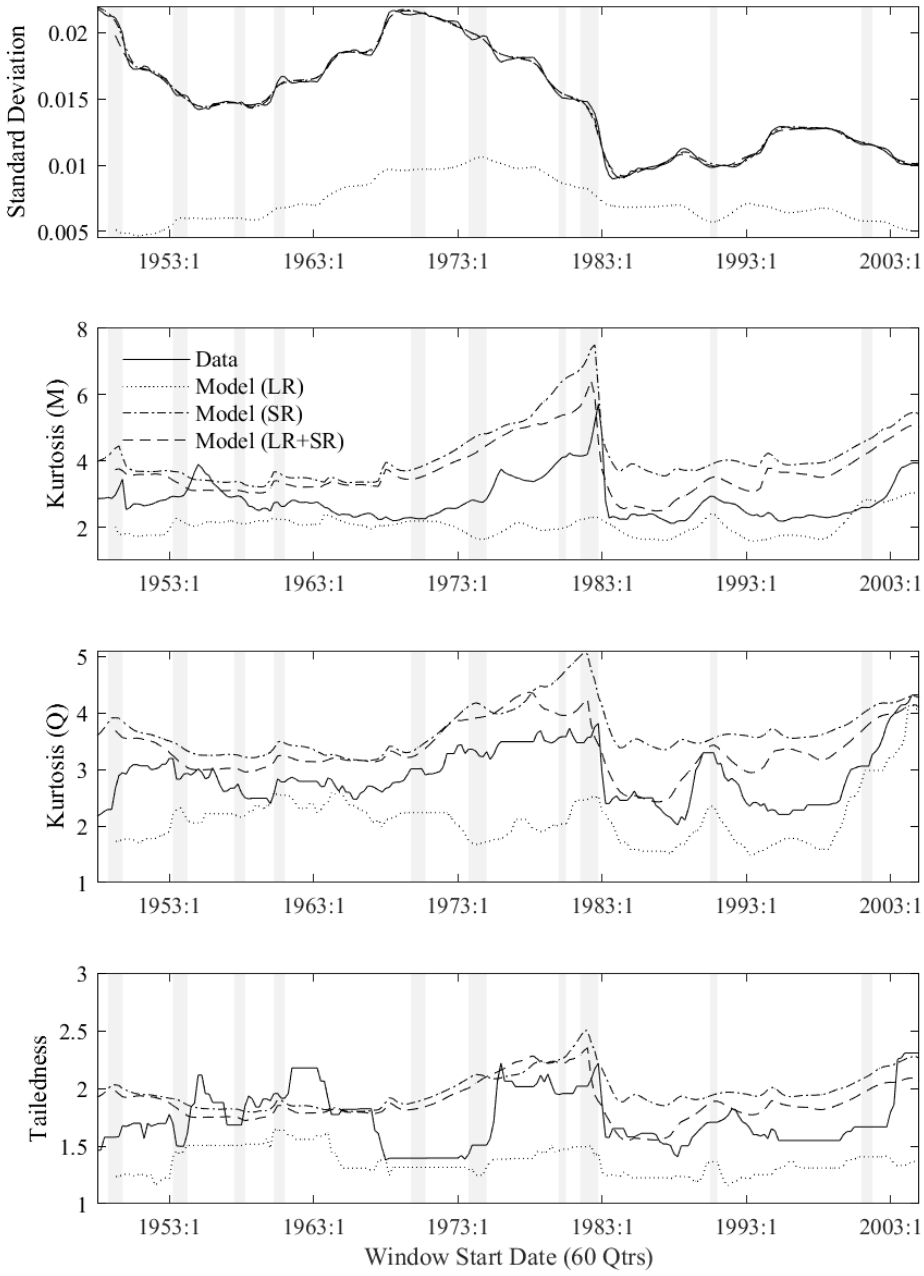

Notes: The simulated data sample is configured into an exhaustive set of overlapping subsets, each containing 60 observations, which are then processed using the HP-filter. The panels present various statistics for the empirical data (denoted by "Data"), the model capturing only long-run trends (denoted by "Model (LR)"), the model capturing slow-moving series (denoted by "Model (SR)"), and a combination of the two models (denoted by "Model (SR+LR)"). Panel 1, 2, 3, and 4 illustrate the standard deviation, the moment-based kurtosis, the quantile-based kurtosis, and the Tailedness statistics, calculated using Liu (2019).

Figure 4: Rolling Windows Analysis (UCUR-2M Filter)

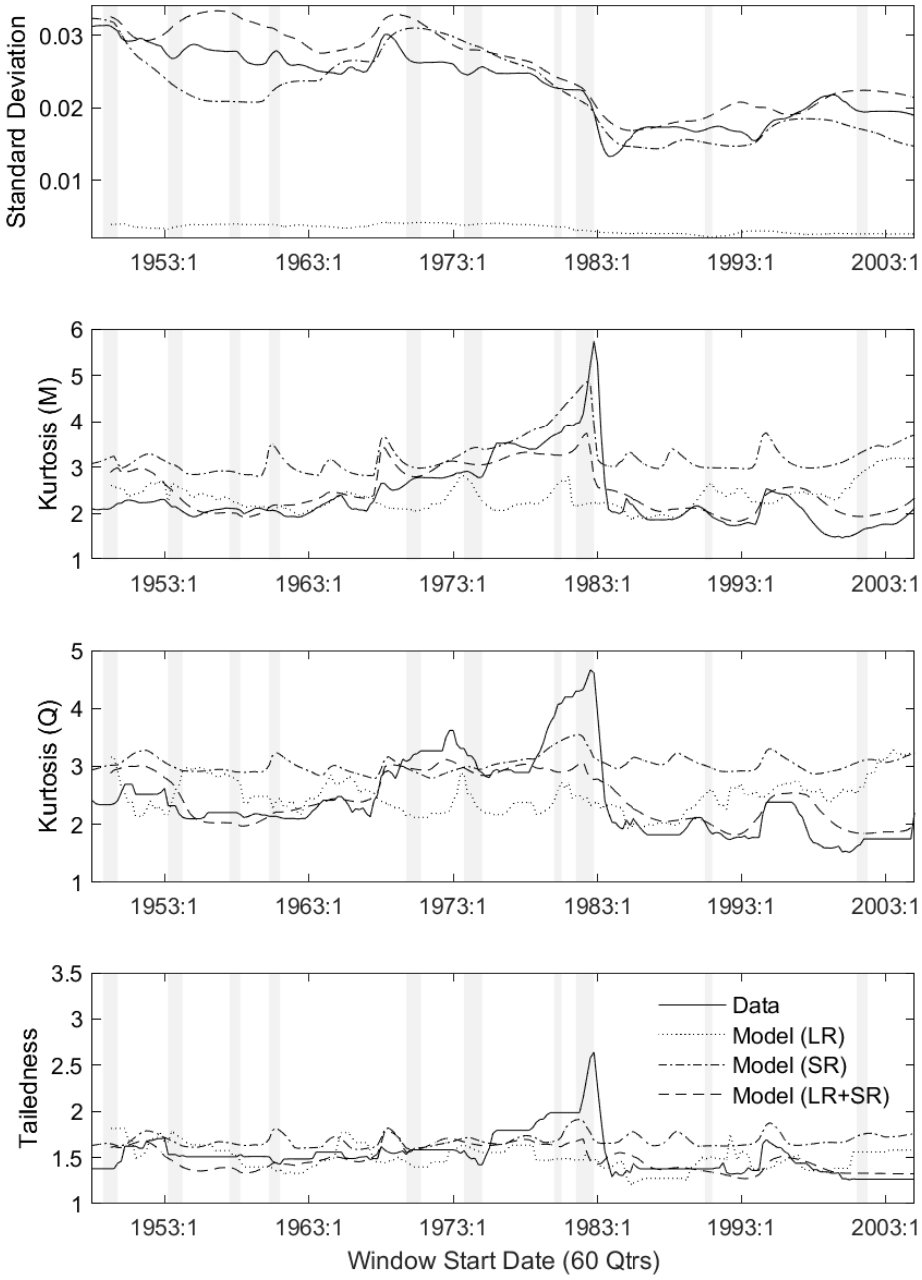

Notes: The simulated data sample is configured into an exhaustive set of overlapping subsets, each containing 60 observations, which are then processed using the UCUR-2M filter developed by Grant and Chan (2017). The panels present various statistics for the empirical data (denoted by "Data"), the model capturing only long-run trends (denoted by "Model (LR)"), the model capturing slow-moving series (denoted by "Model (SR)"), and a combination of the two models (denoted by "Model (SR+LR)"). Panel 1, 2, 3, and 4 illustrate the standard deviation, the moment-based kurtosis, the quantile-based kurtosis, and the Tailedness statistics, calculated using Liu (2019).

References

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2017): "Microeconomic origins of macroeconomic tail risks," *American Economic Review*, 107, 54–108.

Ascari, G., P. Bonomolo, and H. F. Lopes (2019): "Walk on the Wild Side: Temporarily Unstable Paths and Multiplicative Sunspots," *The American economic review*, 109, 1805–1842.

Ascari, G., G. Fagiolo, and A. Roventini (2015): "Fat-Tail Distributions and Business-Cycle Models," *Macroeconomic dynamics*, 19, 465–476.

Atalay, E., T. Drautzburg, and Z. Wang (2018): "Accounting for the sources of macroeconomic tail risks," *Economics letters*, 165, 65–69.

Auerbach, A. J. and Y. Gorodnichenko (2012): "Measuring the output responses to fiscal policy," *American Economic Journal: Economic Policy*, 4, 1–27.

Benhabib, J. and C. Dave (2014): "Learning, large deviations and rare events," *Review of economic dynamics*, 17, 367–382.

Blanchard, O. O. J. and J. Simon (2001): "The Long and Large Decline in U.S. Output Volatility," *Brookings papers on economic activity*, 2001, 135–164.

Brooks, C., S. P. Burke, S. Heravi, and G. Persand (2005): "Autoregressive Conditional Kurtosis," *Journal of Financial Econometrics*, 3, 399–421.

Christiano, L. J. (2007): "Comment on 'On the Fit of New Keynesian Models' by Del Negro, Schorfheide, Smets and Wouters," *Journal of Business and Economic Statistics*, 25, 143–151.

Cooley, T. F. and G. D. Hansen (1989): "The Inflation Tax in a Real Business Cycle Model," *American Economic Review*, 79, 733–748.

Cúrdia, V., M. del Negro, and D. L. Greenwald (2014): "RARE SHOCKS, GREAT RECESSIONS," *Journal of applied econometrics (Chichester, England)*, 29, 1031–1052.

Dave, C. and S. Malik (2017): "A tale of fat tails," *European economic review*, 100, 293–317.

Dave, C. and M. M. Sorge (2020): "Sunspot-driven fat tails: A note," *Economics letters*, 193, 109304.

——— (2021): "Equilibrium indeterminacy and sunspot tales," *European Economic Review*, 140, 103933.

Dave, C. and K. P. Tsang (2014): "Recursive preferences, learning and large deviations," *Economics letters*, 124, 329–334.

Fagiolo, G., M. Napoletano, and A. Roventini (2008): "Are output growth-rate distributions fat-tailed? Some evidence from OECD countries," *Journal of Applied Econometrics*, 23, 639–669.

Fisher, J. D. (2006): "The dynamic effects of neutral and investment-specific technology shocks," *Journal of political Economy*, 114, 413–451.

Ghez, G. and G. S. Becker (1975): *The Allocation of Time and Goods over the Life Cycle*, no. ghez75-1 in NBER Books, National Bureau of Economic Research, Inc.

Grant, A. L. and J. C. Chan (2017): "Reconciling output gaps: Unobserved components model and Hodrick–Prescott filter," *Journal of Economic Dynamics and Control*, 75, 114–121.

Grazzini, J. and D. Massaro (2022): "Great volatility, great moderation and great moderation again," *Review of Economic Dynamics*, 44, 269–283.

Hansen, G. D. (1985): "Indivisible labor and the business cycle," *Journal of Monetary Economics*, 16, 309–327.

Kahn, J. and R. Rich (2007): "Tracking the new economy: Using growth theory to detect changes in trend productivity," *Journal of Monetary Economics*, 54, 1670–1701.

Kim, C.-J. and C. R. Nelson (1999): "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle," *The review of economics and statistics*, 81, 608–616.

Liu, X. (2019): "On tail fatness of macroeconomic dynamics," *Journal of Macroeconomics*, 62.

Mishkin, F. S. (2011): "Monetary policy strategy: lessons from the crisis," Tech. rep., National Bureau of Economic Research.

Morley, J. C., C. R. Nelson, and E. Zivot (2003): "Why Are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP So Different?" *The Review of Economics and Statistics*, 85, 235–243.

Pakko, M. R. (2002): "What happens when the technology growth trend changes? Transition dynamics, capital growth, and the "new economy"," *Review of Economic Dynamics*, 5, 376–407.

Rogerson, R. (1988): "Indivisible labor, lotteries and equilibrium," *Journal of Monetary Economics*, 21, 3–16.

Schmid, F. and M. Trede (2003): "Simple tests for peakedness, fat tails and leptokurtosis based on quantiles," *Computational Statistics Data Analysis*, 43, 1–12.

Stock, J. H. and M. W. Watson (2002): "Has the Business Cycle Changed and Why?" *NBER macroeconomics annual*, 17, 159–218.

A1. Appendix

A1.1. Output Calculations

All data variables were taken from from FRED for the period 1948:QI-2019:Q4. Real GDP per capita (in dollars) is defined as

$$Y_t = \frac{\frac{GDPC1_t}{4}}{CNP160V_t} \times 1000000.$$

This measure of output is then transformed into annualized growth rates as well as logged and detrended using the HP and Hamilton filters.

A1.2. Extracting Trend Shocks

Two of the three experiments performed above utilize shocks to the long-run growth rate of technology at the quarterly frequency. Since the model assumes exogenous growth, trend shocks to output and technology are identical, and this section details how output trend shocks are extracted from the data.

The Real GDP per capita series defined in the previous section is logged and detrended via the HP filter. Define $y_t = log(Y_t)$ and let the cycle and trend components of output be defined by $y_t = y_t^{cycle} + y_t^{trend}$. The trend component is the smoothed, moving-average series extracted from the unfiltered series and illustrated by the dashed line in the upper panel of Figure 1. The cycle component is the difference between the unfiltered series and the trend component, is stationary by definition, and illustrated by the solid line in the lower panel of Figure 1.

The logged-trend series is first-differenced, resulting in quarterly growth rates along the lines of Equation 9 in the text.

$$g_t = y_t^{trend} - y_{t-1}^{trend}$$

The quarterly growth rate series for output is used as a proxy for the growth rate series for technology and used to estimate the coefficients for the AR4 regression detailed in Equation 9.

The resulting coefficients are

$$\tau_1 = 3.71$$

$$\tau_2 = -5.23$$

$$\tau_3 = 3.33$$

$$\tau_4 = -0.80,$$

where all estimated coefficients are significantly different from zero at the 95-percent confidence level and sum to 0.9994. The resulting *trend* shocks from the regression (ν_t) are illustrated by the dashed line in the bottom panel of Figure 1 and is exactly what is fed into the model for the relevant experiments.