# A Cautionary Tale of Fat Tails

Chetan Dave* Scott J. Dressler† Samreen Malik‡

March 4, 2022

**Abstract**

Several macroeconomic time series exhibit excess kurtosis or "Fat Tails" possibly due to rare but large shocks (i.e., tail events). We document the extent to which tail events are attributable to long-run growth shocks. We show that excess kurtosis is not a uniform characteristic of postwar US data, but attributable to episodes containing well-documented growth shocks. A general equilibrium model captures these observations assuming Gaussian business-cycle shocks and a single growth shock from various sources. The model matches the data best with a growth shock to labor productivity while investment-specific technology shocks drive cycles.

JEL codes: E0, E3

*Keywords:* fat tails, growth shocks, real business cycles.

*Email: cdave@ualberta.ca. University of Alberta.
†Email: scott.dressler@villanova.edu. Villanova University.
‡Email: samreen.malik@nyu.edu. New York University Abu Dhabi.

# 1 Introduction

Our analysis is motivated by two observations. First, many macroeconomic time series are not normally distributed and instead exhibit a kurtosis exceeding 3 (i.e., they exhibit *fat tails*). Table 1 reports the kurtosis of real, per-capita output (GDP), consumption, investment, and labor hours for the post-war US (1948:Q1-2019:Q4). When considering HP-detrended logged series, the table reports that both output and investment significantly exhibit fat tails. When considering annualized growth rates, the list of variables significantly exhibiting fat tails increases to all four primary business-cycle variables. Second, this excess kurtosis is primarily attributable to large economic shocks that differ from regular business-cycle fluctuations (i.e., *tail events*). Figure 1 follows Acemoglu et al. (2017) and illustrates two sets of quantile-quantile (Q-Q) plots for the same output observations presented in Table 1 against the standard normal distribution. The top row of the figure considers the full sample and suggests that a normal distribution significantly underestimates the frequency of large economic shocks. The bottom row of the figure shows that a normal distribution does a much better job at capturing the data when restricting attention to the interdecile range. Namely, the interior 90 percent of observations after removing those exceeding 1.645 standard deviations away from their respective means.

Table 1: US Data (1948:QI-2019:QIV)

| | Output | Consumption | Investment | Hours |
|---|---|---|---|---|
| HP filter | 3.86* | 3.47 | 4.53* | 2.82 |
| Growth rates | 4.60* | 8.43* | 5.82* | 4.71* |

Note: * indicates rejection of $H_0 : kurtosis = 3$ with 90% confidence

These two observations as well as their implications have been documented in previous studies. Christiano (2007) finds substantial evidence against the use of a Gaussian likelihood function in empirical VAR analyses, suggesting the potential for specification bias. Mishkin (2011) warns against the use of Gaussian shocks in quantitative studies of optimal monetary policy, suggesting that "the distribution of shocks hitting the economy is more complex." This warning not only

Figure 1: Quantile-Quantile Plots of Postwar (1948:QI-2019:QIV) US GDP

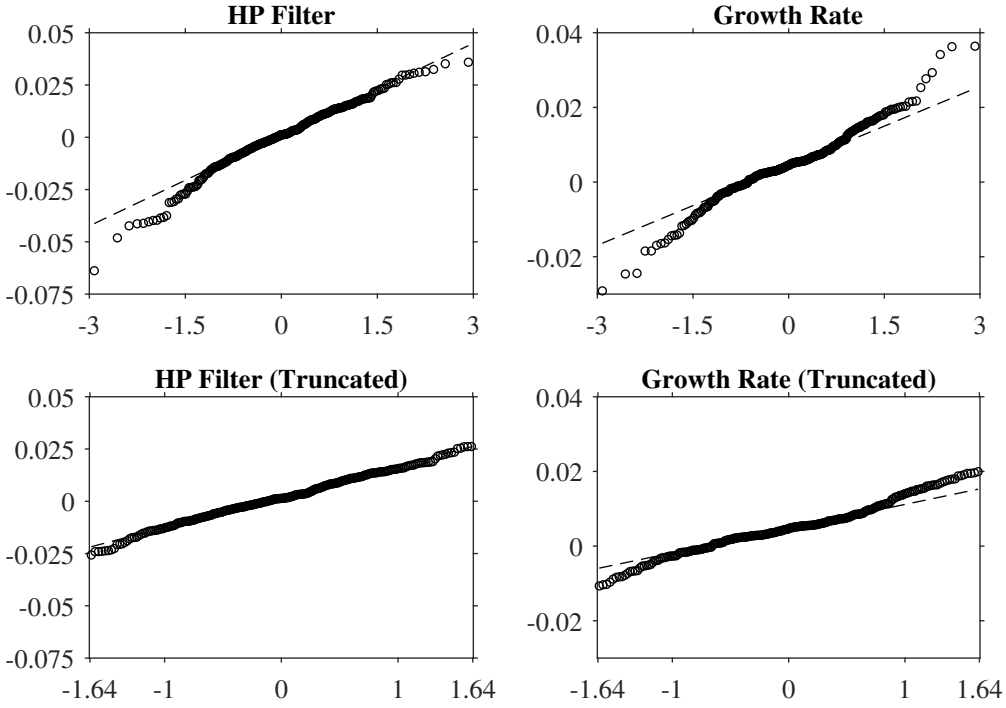


has implications for the examination of monetary policy, but for any proposed macroprudential policy or welfare analysis.[1]

We make two novel contributions to the topic of fat tails in business-cycle aggregates being attributable to tail events. Empirically, we attempt to locate these tail events in US output by examining the distributional properties of many subsets of the full sample. Our *rolling-windows* analysis (discussed in detail below) suggests that these tail events are clustered into one or two episodes, and these episodes are historically associated with shifts in long-run trends. In other words, these tail events can potentially be associated with *growth shocks*. We focus on one US episode (late 1970s to early 1980s) where tail events are identified in both HP-filtered GDP and GDP growth, and employ a standard real business cycle model to determine the extent to which a tail event can be attributable to a growth shock. Our model is capable of considering multiple sources of growth shocks and business-cycle fluctuations (i.e., preferences, total factor produc-

[1]These observations as well as additional implications have been documented by Kim and Nelson (1999), Blanchard and Simon (2001), Stock and Watson (2002), Fagiolo et al. (2008), and several others.

tivity, and investment-specific technology), and we are able to consider different combinations of short and long-run shocks to see which fit the totality of the data best.

We show in detail below that the model finds a growth shock to be a very plausible tail event. By plausible, we mean: (i) *one* growth shock amidst regular (Gaussian) business-cycle fluctuations is sufficient to deliver excess kurtosis in a full simulated sample; (ii) the size of this growth shock is not excessively large and always within 1.5 standard deviations of the business-cycle shock process; and (iii) all sources of growth shocks considered are able to match the full sample kurtosis. While all pairwise combinations of shocks mentioned have the ability to match the fat tails observed in US output, we find the shock combinations allowing the model to best match the data across other business-cycle variables have investment-specific technology driving the business cycle while a growth shock impacts the labor market (on either the supply or demand-side).[2]

We consider this paper a *cautionary* tale of fat-tail behavior because if fat tails in business cycle variables are due to growth shocks, then this suggests an instance where *long-run* shocks are distorting *short-run* dynamics.[3] If fat tails were in fact a convolution of trend and cycle, then applying a filter that reportedly does a more efficient job at removing trend from cycle should deliver a short-run series that does not exhibit fat tails. We reexamine our empirical and simulated-model results using the filter proposed by Hamilton (2018) and find that to be the case. In particular, we show that US output, consumption, investment, and labor hours show no evidence of fat tails in the full US sample when detrended using Hamilton's filter. The same result holds for our simulated data using Hamilton's filter, even though the same simulated data predict significant fat tails in either growth rates or when detrended using the HP filter.

Our results make several contributions to the extant literature regarding an explanation of the fat tails of macroeconomic aggregates. One branch of the literature focuses on exogenous sources of tail events stemming from either non-Gaussian shock distributions, stochastic volatility of the

[2]Our conclusion of investment-specific technology being the likely driver of business cycles is consistent with the conclusions of Fisher (2006).

[3]An extensive literature (e.g., Kahn and Rich (2007)) reports that short-run shocks can distort the impact of long-run shocks and propose identifying restrictions to best extract the long-run properties of the data. Our results suggest that the opposite result, that long-run shocks distort the impact of short-run shocks, can be true as well.

exogenous shock process, or uncertainty.[4] Ascari et al. (2015), for example, find that both real business cycle and New Keynesian environments are unable to deliver fat tails in endogenous variables when being subjected to exogenous, Gaussian shocks. Our model can deliver fat tails by considering a combination of Gaussian business-cycle shocks and a single growth shock which can be considered a draw from a second Gaussian distribution.[5] Another branch of the literature focuses on endogenous sources of tail events such as exogenous shock amplification due to fiscal multipliers, constant gain adaptive learning, indeterminacies and sunspots, or disaggragated-sectoral shocks.[6] While the growth shocks in our model are general and not indicative of any particular endogenous amplification mechanism, our model does indicate that a labor market disruption is the most likely candidate. Moreover, our empirical contribution suggesting that tail events arise in specific episodes of US data can be used to better discipline subsequent analyses. For example, any endogenous explanation of tail events should account for when the tail events actually occurred.

The paper proceeds as follows. Section 2 details our empirical analysis and lends support for our conjecture that the tail events observed in the data might in fact be due to growth shocks. Section 3 lays out our model, while Section 4 reports the quantitative analysis of the model. Section 5 concludes.

[4]Literature considering non-Gaussian shock distributions include Barro (2009), Cúrdia et al. (2014), and Chib and Ramamurthy (2014). Literature considering stochastic volatility of the exogenous shock process include Cogley and Sargent (2001, 2005), Lux and Sornette (2002), and Justiniano and Primiceri (2008). Literature considering uncertainty shocks include Bloom (2009), Mumtaz and Theodoridis (2017), and Ludvigson et al. (2021).

[5]One can imagine a single growth shock being drawn from a Gaussian distribution when only considering the rare, non-zero shocks as part of the support. We acknowledge that one can argue for a non-Gaussian distribution of growth shocks because all of the zero (no growth shock) observations would result in any non-zero growth shock being several standard deviations away from its mean.

[6]Literature considering fiscal multipliers include Auerbach and Gorodnichenko (2012). Literature considering learning include Benhabib and Dave (2014), Dave and Tsang (2014), and Dave and Malik (2017). Papers considering sunspots include Ascari et al. (2019) and Dave and Sorge (2020, 2021). Literature considering sectoral shocks include Leduc and Liu (2016) and Acemoglu et al. (2017).

## 2 The Data

Given the evidence documented above that tail events are attributable to fat tails in short-run macroeconomic aggregates, this section lays out an analysis to get an idea of exactly *when* these tail events took place.

Our analysis begins with considering a macroeconomic aggregate of interest, e.g., post-war US GDP per capita. Our data ranges from the first quarter of 1948 to the last quarter of 2019, resulting in a total of 288 observations.[7] This data series possesses a full-sample kurtosis of 3.86 when HP-detrended and logged, and 4.60 when considered in growth rates (see Table 1). In order to determine if fat tails are either consistent throughout the sample or attributable to particular subsets of data, one can split the detrended sample into subsets and calculate the kurtosis of each subset independently.[8] In order to thoroughly look at all possible subsamples of data, we performed the following *rolling windows* analysis. First, we set a subsample size of $n = 60$ and construct subsamples of the full sample where the first subsample contains observations 1 through 60, the second contains observations 2 through 61, and so on. This results in 229 $(288 - n + 1)$ overlapping subsets.[9] Second, we calculate the kurtosis of each subset and test if this value is significantly different than 3 via a t-test as we did with the full sample.

The results of our rolling window analysis are illustrated in Figure 2. The solid line indicates the kurtosis of each data window indicated at its starting date, and the vertical bars indicate NBER dated recessions. The dashed horizontal line at 4.04 indicates the critical (t-test) value for kurtosis being significantly different from 3 with 90 percent confidence given 60 observations, and a black dot indicates windows with kurtosis values exceeding this critical value. The top panel of the figure considers HP-detrended logged US GDP per capita, and although the full sample kurtosis was significantly different than 3, significant kurtosis values only stem from windows with a

[7] We intentionally excluded the data corresponding to the COVID pandemic.

[8] Note that we detrend the entire sample once prior to splitting the sample. This ensures that our results are not in any way attributable to the performance of the detrending filter at the tails of the split samples.

[9] We also consider alternative window sizes of 40 and 80 quarters consisting of 249 and 209 overlapping subsets of data, respectively. We present the resulting Kurtosis in each subset in Figure A1-A2, and find that our main stylized fact is robust to alternative window selection.

start date beginning from 1980:Q2 to 1983:Q1. Focusing on this episode allows us to get a rough idea of when a tail event occurred. The first significant rolling window in 1980:Q2 has a kurtosis of 4.11 and includes observations between 1980:Q2 and 1995:Q2 (15 years). This means that if a tail event was driving the kurtosis value, then it could have occurred at anytime throughout the sample. The last significant rolling window in 1983:Q1 has a kurtosis of 4.8 and the rolling window immediately after (1983:Q2) has a kurtosis of only 2.79. This sharp drop in kurtosis from one window to the next is quite surprising when considering that they share many of the same observations and differ in only one observation at each end.[10] Since the window beginning in 1983:Q1 displays significant excess kurtosis but the window beginning one quarter later does not, this leads us to conclude that 1983:Q1 is a likely date for a tail event.

It should be stressed that while our rolling-window analysis is informative, it is not a precise tool. For example, if a significant tail event occurred in 1983:Q1, then why do we not see excess kurtosis in every window containing this date (i.e., windows beginning in 1968:Q1 and so on)? One of the many possible reasons is the HP filter itself, this is why the middle panel considers the same analysis using GDP growth. While the full sample kurtosis of GDP growth was 4.60, this panel indicates that three episodes show rolling windows with kurtosis significantly different than 3. The middle episode roughly lines up with the episode delivered by the HP filter and does show more potential dates for a tail event within this episode. While the growth results indicate two additional episodes (mid 1950s and post 1994), we focus on the middle episode solely as support for the HP filtered results.[11]

The empirical result thus far is that there is evidence of a tail event occurring somewhere within the late 1970s to early 1980s when considering either HP-detrended log GDP or GDP growth. The interesting feature of this particular episode is that it is full of what one may potentially consider to be growth shocks. For example, the literature often discusses events such as the two oil crises, the Great Inflation / Volker monetary policy shocks, the Reagan-era double

[10] Recall that the full sample was detrended once before creating the windows, meaning that this result is not an artifact of the HP filter.

[11] While not to detract from the scope of the current analysis, we find the rolling window results from later in the sample to be less reliable given the presence of nonstationary growth due to the Global Financial Crisis.

dip recessions, changes in immigration laws, and other significant events that could be argued as having a long-run impact on the US economy. While making no particular case for any single incident giving rise to a growth shock during this episode, it appears reasonable to believe that at least one long-run deviation from trend occurred during this episode.

If it is true that the fat tails observed in the full sample and rolling windows analyses are attributable to a growth shock, then it implies that long-run shocks are in fact distorting the short-run dynamics. Further along the lines of this argument, if a particular filter did a better job at disentangling the short-run (cyclical) and long-run (trend) dynamics then those used thus far, then the short-run dynamics obtained from this filter should possess less excess kurtosis. We investigate this line of reasoning by considering a relatively new filter described in Hamilton (2018). This filter is reported to achieve all of the objectives of the HP filter with none of the drawbacks.[12] Repeating the empirical analysis using data that was logged and detrended using Hamilton's filter delivers two results. First, the full-sample analysis fails to deliver any kurtosis measures significantly different than 3. In particular, the full-sample kurtosis for output, consumption, investment, and labor hours using this filter are respectively 3.30, 3.31, 3.39, and 3.51. Second, the bottom panel of Figure 2 reports the rolling window results using the Hamilton filter. The figure indicates that a large amount of the episodic fat tail behavior observed in the previous two data transformations is significantly reduced. In particular, the Hamilton filter indicates a short (two-quarter) episode ending in 1983:Q1 inline with the HP filter, and another episode towards the end of the sample inline with the data in growth. We view these results as indicating that at least some of the full sample excess kurtosis can be attributed to shifts in long-run trends confounding the short-run dynamics. The fact that the rolling window analysis cannot completely remove the 1982:Q4-1983:Q1 episode only indicates further that something significant occurred at that time that can be considered a tail event.

[12] Hamilton (2018) states the drawbacks of the HP filter as being: (i) it produces series with spurious dynamic relations that have no basis in the underlying data-generating process; (ii) filtered values at the end of the sample are very different from those in the middle; and (iii) A statistical formalization of the problem typically produces values for the smoothing parameter vastly at odds with common practice, e.g., a value for $\lambda$ far below 1600 for quarterly data. The filter proposed by Hamilton (2018) is an intuitive regression-based filter, and the interested reader is directed there for details.

This section concludes with briefly reporting the rolling-window analysis for the other business-cycle variables using the three data transformations previously discussed. Figure 3 reports the rolling window kurtosis results for HP-detrended consumption (top panel), Investment (middle panel), and labor hours (bottom panel). Along the lines of investment being the only variable with kurtosis significantly different than 3 in the full sample, it is also the only variable having significant kurtosis in any of the rolling windows. There appear to be two episodes where the windows exhibit fat tails that are not shared with GDP, but the middle episode is inline with GDP and even shares the drop in kurtosis between the 1983:Q1 and Q2 windows. Figure 4 reports the rolling window analysis considering the same variables in growth rates. As in the full-sample results, there appears to be more excess kurtosis discovered when using growth rates as opposed to the HP filter. In addition, the middle episode of the rolling-window analysis for consumption growth appears to line up with output growth, while the episode for growth in investment and labor hours appear to occur a bit earlier. Finally, Figure 5 considers the same analysis using Hamilton's filter. Interestingly, the episode shared across all four business-cycle variables is around 1995 and not in 1983 as the output results indicate. Otherwise, these results are similar to those obtained using the HP filter, that there are a couple of episodes of excess kurtosis in investment but fewer episodes in consumption or labor hours.

Figure 2: GDP Per Capita

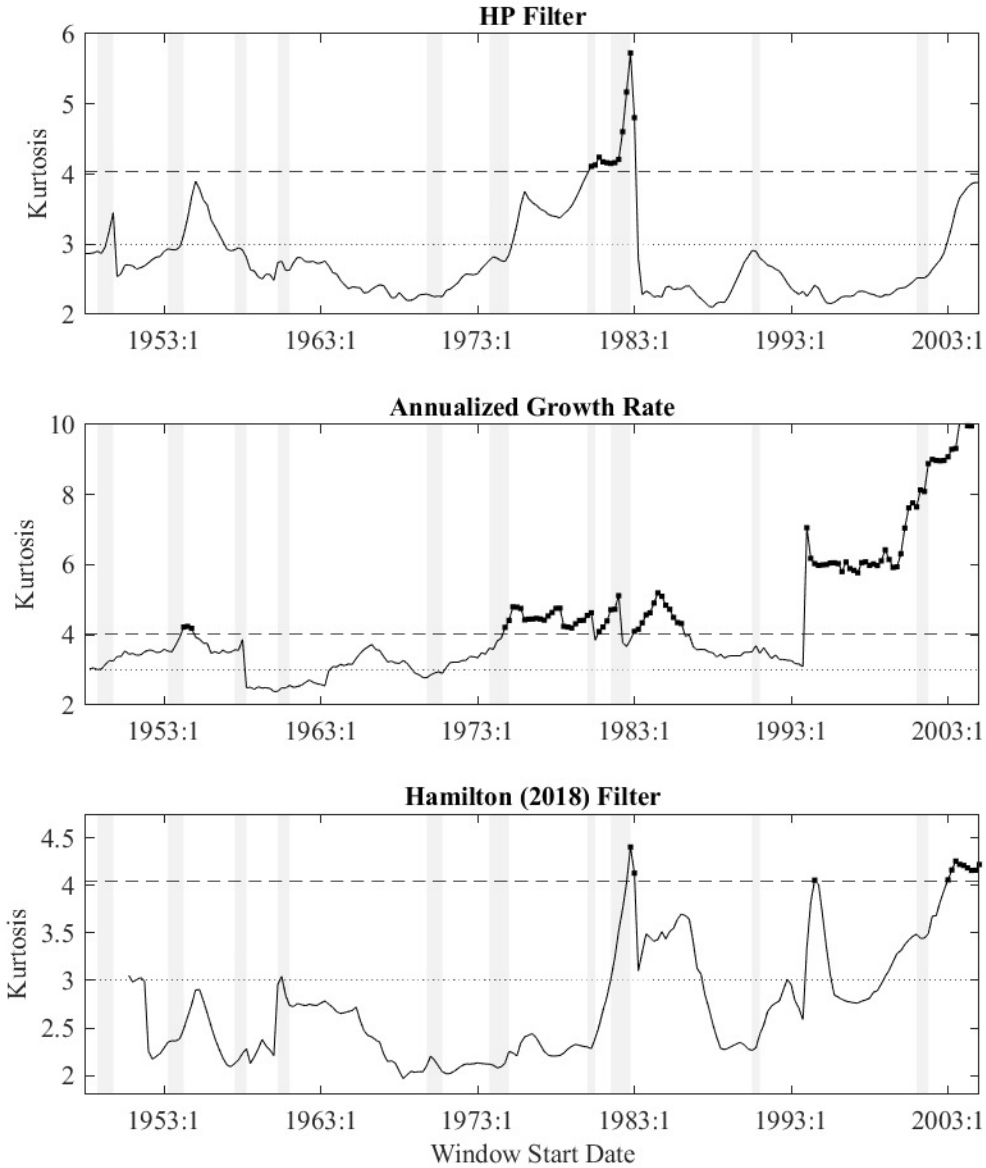

Figure 3: HP Filtered Data

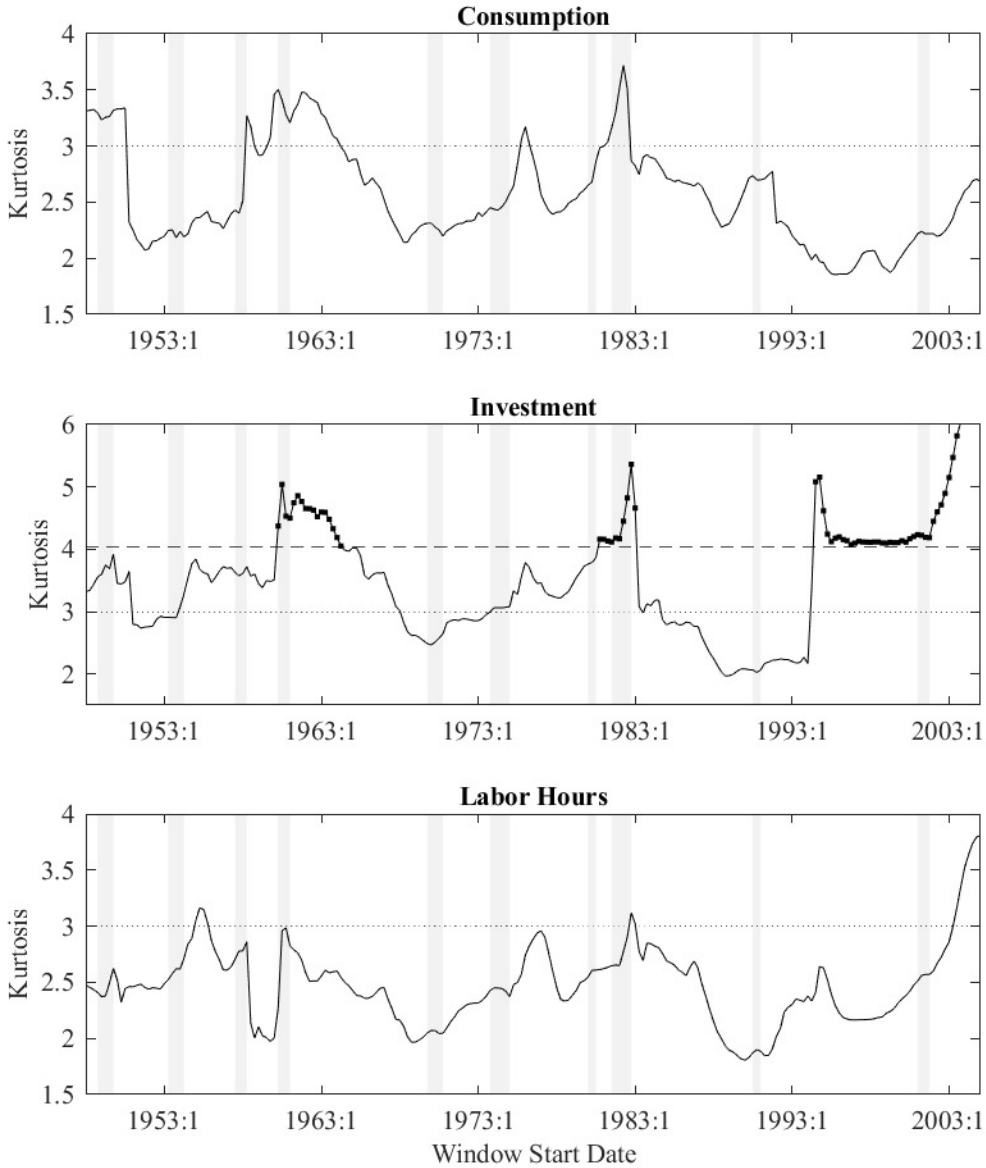

Figure 4: Growth Data

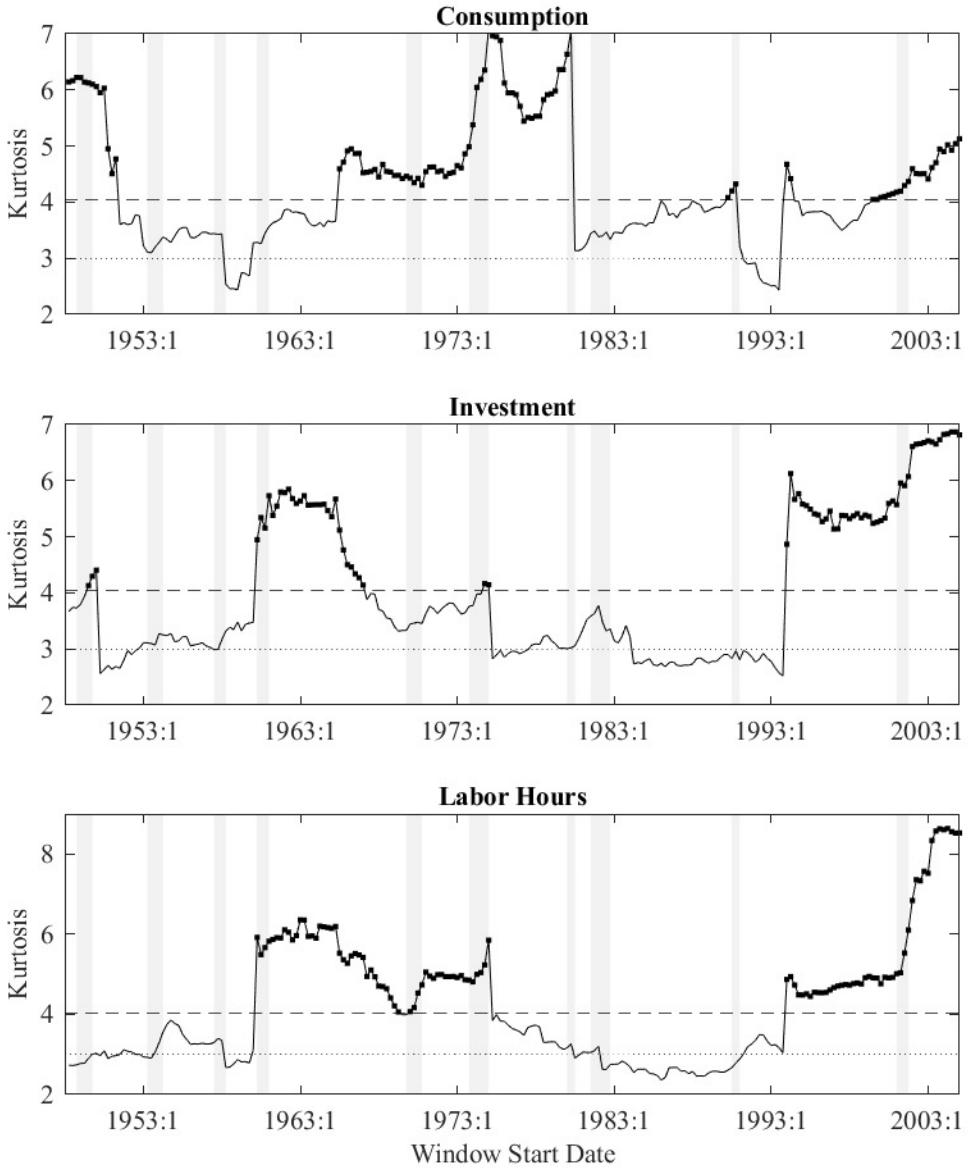

Figure 5: Hamilton (2018) Filtered Data

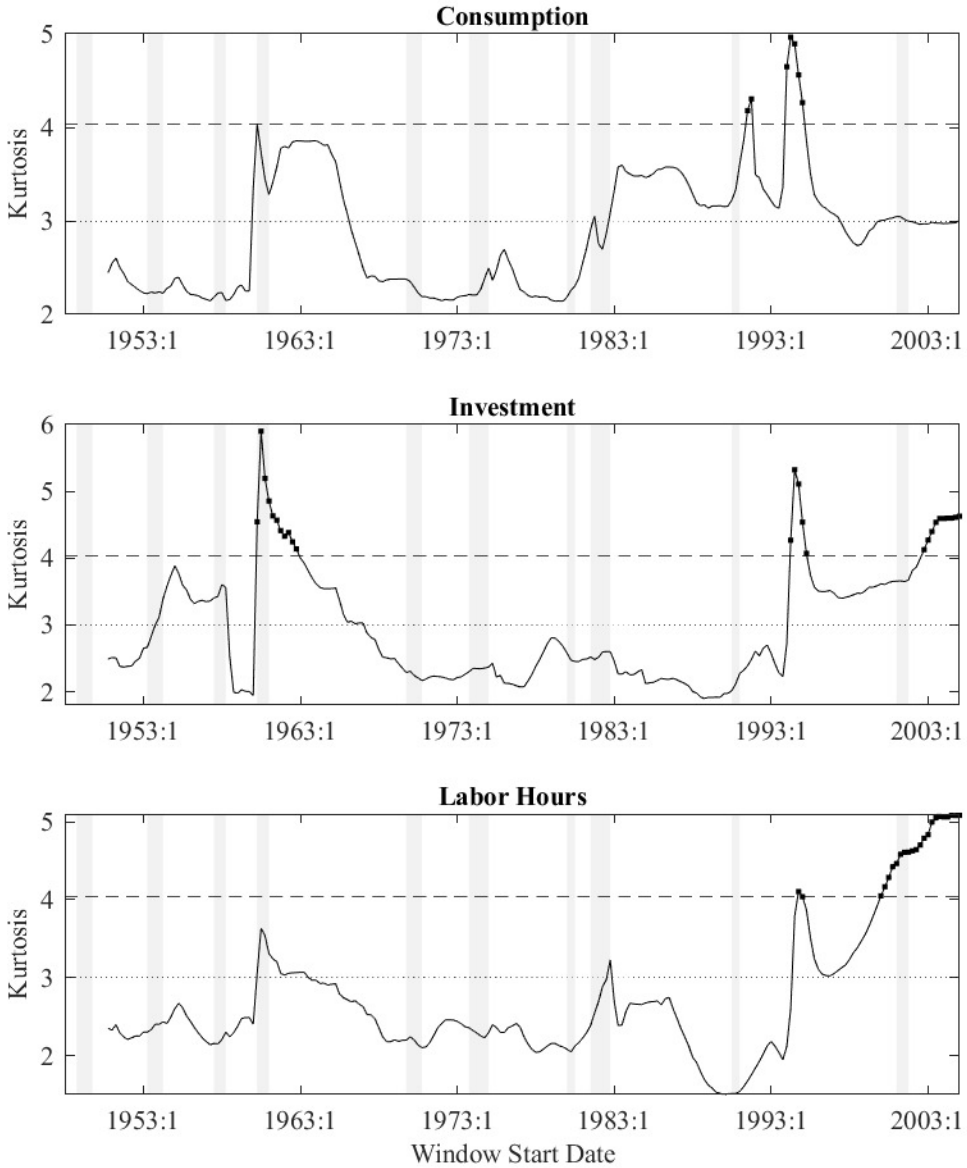

# 3 Model

## 3.1 Overview

The model begins with the neoclassical framework of Fisher (2006) that distinguishes between shocks to labor-augmenting total-factor productivity (TFP) and investment-specific technology (IST), and adds preference shocks as in Ireland and Schuh (2008). Following Pakko (2002), each of our three shock processes contain a short-run component that is stationary in levels and a separate long-run component that is stationary in growth rates.

The model is subjected to a series of experiments where only one shock source is assumed active in the short-run and only one shock source in the long-run. Since there are three short-run and long-run shocks each, there are nine total model specifications. For a given specification, the model is simulated with a series of short-run shocks and a single long-run shock corresponding to 1983:Q1 as detailed above. Key parameters governing the short and long-run shocks are determined via a dynamic calibration so the model matches the standard deviation and kurtosis of HP-detrended output observed in the full post-war US data. The remaining predictions of each model specification with respect to the rolling windows analysis, as well as consumption, investment, and labor hours are then considered to assess the plausibility of each model specification.

## 3.2 Preferences and Technologies

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

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[ ln(C_t) - \Psi H_t / \tilde{A}_t \right], \tag{1}$$

where $\beta$ is the discount factor and $C_t$ and $H_t$ respectively denote consumption and labor hours supplied at time $t$. The representative household's utility is logarithmic in consumption and linear in leisure to make the model consistent with balanced growth. While some non-separable utility

specifications consistent with balanced growth could be used, the specification in (1) 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 preference shock $\tilde{A}_t$ in (1) impacts the marginal rate of substitution between consumption and labor supply such that an increase in $\tilde{A}_t$ is associated with an increase in equilibrium hours worked.[13]

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

$$C_t + I_t \leq K_t^{\alpha}(\tilde{Z}_t H_t)^{1-\alpha} \tag{2}$$

$$K_{t+1} \leq (1-\delta)K_t + \tilde{V}_t I_t \tag{3}$$

In the above equations, the Cobb-Douglas share parameter $\alpha$ and the depreciation rate of the capital stock $\delta$ both lie between zero and one.

Equations (2) and (3) each possess a technology shock. The TFP shock $\tilde{Z}_t$ in (2) is labor-augmenting in order for balanced growth to be feasible. In contrast to the preference shock $\tilde{A}_t$, technology shocks of this kind are interpreted as a source of disturbances that directly influence general production possibilities such as changes in taxes, regulation, or market structure. The investment-specific technology (IST) shock $\tilde{V}_t$ in (3) directly influences the transformation of output invested in the current period into capital available for production in the following period, and has been interpreted as a source of financial shocks by Justiniano and Primiceri (2008). Modelling IST shocks in this manner imply that the number of consumption units that must be exchanged to acquire an efficiency unit of the investment good is $1/\tilde{V}_t$. This exogenous, real price

[13]Preference shocks of this kind have been interpreted as a source of non-technological disturbances that can drive aggregate fluctuations at either short or long horizons by Hall (1997), Chang and Schorfheide (2003), Comin and Gertler (2006), Galí et al. (2007), and others.

of an investment good directly shows that while TFP shocks impact production possibilities, only IST shocks can impact the marginal rate of transformation between consumption and investment goods.

The model described above is the simplest environment that can make distinctions between non-technological or demand-side fluctuations ($\tilde{A}_t$), general technological or supply-side fluctuations ($\tilde{Z}_t$), and specific production fluctuations ($\tilde{V}_t$). This distinction is all that is required at the moment to see if short and long-run fluctuations from any combination sources can plausibly account for the stylized facts. Environments with specific consumption and investment producing sectors such as Ireland and Schuh (2008) or multiple capital goods such as Greenwood et al. (1997) would undoubtedly add detail to the analysis, but the complications associated with these more elaborate models are not necessary at this stage of our analysis.

## 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 (1), subject to the constraints imposed by (2) and (3) for all $t$. Letting $\Lambda_t$ denote the non-negative multiplier on the budget constraint (2), and using (3) to remove $I_t$ from the problem, the first-order conditions for this problem can be written as

$$\frac{1}{C_t} = \Lambda_t, \tag{4}$$

$$\frac{\Psi}{\tilde{A}_t} = \Lambda_t \left[ (1-\alpha) K_t^{\alpha} H_t^{-\alpha} \tilde{Z}_t^{1-\alpha} \right], \tag{5}$$

$$\frac{\Lambda_t}{\tilde{V}_t} = \beta E_t \Lambda_{t+1} \left[ \alpha K_{t+1}^{\alpha-1} (\tilde{Z}_{t+1} H_{t+1})^{1-\alpha} + (1-\delta)/\tilde{V}_{t+1} \right], \tag{6}$$

and (2) with equality for all $t$.

Intuitively, while (2) shows how the TFP shock $\tilde{Z}_t$ directly impacts the household's budget constraint, (4) indicates how $\Lambda_t$ measures the marginal utility of consumption for the household, (5) shows how the preference shock $\tilde{A}_t$ impacts the trade-off between the marginal utility of

consumption and the marginal product of labor, and (6) shows how the IST shock $\tilde{V}_t$ impacts the intertemporal trade-off between consumption and investment.

### 3.4 Driving Processes

The model is closed by specifying the stochastic behavior of the three exogenous shocks. As stated previously, each of these shocks possess a short-run component stationary in levels, and a long-run component stationary in growth rates. These components are formally defined for the preference shock as $\tilde{A}_t = e^{a_t} A_t$, where $a_t$ and $A_t$ are respectively the short and long-run components. The short-run component follows an AR(1) process

$$a_t = \rho_a a_{t-1} + \varepsilon_t^a, \tag{7}$$

with $|\rho_a| < 1$ and $\varepsilon_t^a$ representing iid draws from a Normal distribution with zero mean and standard deviation $\sigma_a^l$. The long-run component $A_t$ represents the cumulative product of growth shocks to preferences where

$$A_t = e^{g_t^a} A_{t-1} = \prod_{s=0}^{t} e^{g_s^a}, \tag{8}$$

$$g_t^a = (1 - \tau_a) g^a + \tau_a g_{t-1}^a + \nu_t^a, \tag{9}$$

with $|\tau_a| < 1$ and $\nu_t^a$ representing iid draws from a Normal distribution with zero mean and standard deviation $\sigma_a^g$. The term $g^a$ in (9) denotes the long-run growth rate of preferences.

Expressions for the TFP and IST shocks are specified in a similar fashion, $\tilde{Z}_t = e^{z_t} Z_t$ and

$\tilde{V}_t = e^{v_t} V_t$, where

$$z_t = \rho_z z_{t-1} + \varepsilon_t^z, \tag{10}$$

$$Z_t = e^{g_t^z} Z_{t-1} = \prod_{s=0}^{t} e^{g_s^z}, \tag{11}$$

$$g_t^z = (1-\tau_z) g^z + \tau_z g_{t-1}^z + \nu_t^z, \tag{12}$$

$$v_t = \rho_v v_{t-1} + \varepsilon_t^v, \tag{13}$$

$$V_t = e^{g_t^V} V_{t-1} = \prod_{s=0}^{t} e^{g_s^v}, \tag{14}$$

$$g_t^v = (1-\tau_v) g^v + \tau_v g_{t-1}^v + \nu_t^v, \tag{15}$$

with $|\rho_z|,\ |\rho_v|,\ |\tau_z|,\ |\tau_v| < 1$, and $\varepsilon_t^z$, $\varepsilon_t^v$, $\nu_t^z$, and $\nu_t^v$ representing iid draws from Normal distributions with zero means and standard deviations given by $\sigma_z^l$, $\sigma_v^l$, $\sigma_z^g$, and $\sigma_v^g$, respectively.

In the short-run, shocks to both level and trend will impact the model variables. The purpose of the quantitative exercise below is to determine if a short-run data characteristic like excess kurtosis can be attributable to a confounding of 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. Since these three shocks can have different growth rates, specific transformations are necessary for all model variables to be stationary. In particular, labor hours grow at the same rate as the long-run component of the preference shock so

$$h_t = \frac{H_t}{A_{t-1}}.$$

Output, consumption, and investment must all grow by the same rate according to the budget constraint (2). This rate is given by

$$G_{t-1} = A_{t-1} Z_{t-1} V_{t-1}^{\frac{\alpha}{1-\alpha}},$$

so $c_t = C_t/G_{t-1}$, $i_t = I_t/G_{t-1}$, and $y_t = Y_t/G_{t-1}$ where $Y_t = C_t + I_t$. Finally, (3) states that the capital stock grows faster than output, so $k_t = K_t/(G_{t-1}V_{t-1})$. These transformations can be confirmed to deliver stationary variables along a balanced growth path.

# 4 Quantitative Analysis

## 4.1 Overview

Our strategy for analyzing the model is as follows. All model parameters identifiable at steady state are calibrated such that the model matches several long-run moments of the US economy. With these parameters fixed, a model specification is chosen where one of the three level shocks is designated to be the short-run shock and one of the three trend shocks is designated to be the long-run shock. The model is simulated with multiple shocks to levels and a single trend shock. The size of the short-run shocks are determined by the standard deviation $\sigma^x,\ x \in \{a, z, v\}$, while the size of the single long-run shock is given by $g^X,\ X \in \{A, Z, V\}$. Since neither $\sigma^x$ or $g^X$ are identifiable in steady state, they are determined via a *dynamic* calibration so the model simulations match the kurtosis and standard deviation of the full, post-war US sample. With these parameters determined, the model is then simulated 1000 times for 288 quarters each (the size of the post-war sample), and the simulated data is analyzed in the same manner as the true data detailed above. The calibration and the model simulation results are detailed below.

## 4.2 Calibration

The first step of the 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 $\Psi$ is set to 2.87 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 $\alpha = 0.36$.

All exogenous, steady-state growth rates $(g^A, g^Z, g^X)$ are set to 1-percent annually, and then $\beta = 0.997$ so the capital stock to annual output ratio of the model is 2.5. While the sizes of the exogenous growth rates impact the size of the discount parameter $\beta$, it has been verified that the quantitative results presented below are not sensitive to the particular choices of steady-state growth rates.

The second step of the calibration strategy is to select one source of exogenous shocks to levels and one source of a single exogenous shock to trend. While the shocks to levels will occur once every quarter as in traditional business-cycle analyses, the single shock to trend is chosen to occur in the quarter corresponding to 1983:1 (as argued in the previous section).

In order to place all level shocks on an equal field, the persistence parameter $\rho_x = 0.9,\ x \in \{a, z, v\}$ for all simulations. This value is within the traditional value of TFP persistence of 0.95 as in Prescott (1986) and IST persistence of 0.8 as in Justiniano et al. (2011). We similarly treat trend shocks equally: the persistence parameter $\tau_X = 0.95,\ X \in \{A, Z, V\}$ for all simulations.[14]

With these details established, the dynamic calibration strategy amounts to determining the standard deviation of the shock to levels $\sigma^x$ and a single shock to trend $g^X$ so the predicted values of select moments match those of the full US data sample. The moments selected are the full-sample kurtosis and standard deviation of HP-filtered (logged) US GDP, respectively calculated to be 3.8616 and 0.0161. The resulting values of $\sigma^x$ and $g^X$ are those that minimize the squared distance between the moments in the data and those predicted from the model.

The results from the dynamic calibration strategy performed on all nine specifications of the model are reported in Table 2. While the totality of model predictions for a few of the specifications listed in the table will be detailed below, there are a couple of general results worth noting.

[14] It has been verified that changing the values of these persistence parameters impacts the values of the parameters determined via our dynamic calibration procedure, but they do not significantly impact the quantitative predictions of the model. In other words, a larger (smaller) value of $\rho_x$ or $\tau_X$ is met with a smaller (larger) value for the dynamically calibrated parameters $\sigma^x$ and $g^X$. The model does an equally good job of matching the targeted moments, and the business-cycle predictions remain unchanged.

First, all nine specifications determined values of $\sigma^x$ and $g^X$ such that the predicted kurtosis and standard deviation match the data up to four decimal places. While this result implies that we cannot rule out any of the model specifications at this stage of the analysis, it does suggest that a combination of short-run shocks and a single long-run shock from a variety of sources is easily capable of matching the excess kurtosis of output observed in the full US sample. Second, the size of the single trend shock (relative to $\sigma^x$) is less than 50 percent larger than the standard deviation of the level shock (in absolute value) in all specifications. This implies that the model does not require an implausibly large trend shock to combine with the Gaussian level shocks in order to produce non-Gaussian characteristics in output.

Table 2: Parameters: $[g^X, \sigma^x]$

| | z | v | a |
|---|---|---|---|
| **Z** | [-0.0126, 0.0086] | **[-0.0127, 0.0102]** | [-0.0126, 0.0084] |
| **V** | [-0.0098, 0.0084] | [-0.0099, 0.0099] | [-0.0098, 0.0082] |
| **A** | [-0.0124, 0.0086] | **[-0.0125, 0.0102]** | [-0.0124, 0.0084] |

*Notes:* Items in bold are specifications that where the model predictions fit the totality of data observations best

## 4.3 Main Model Results

This section reports the model results from the specification where the level shocks stem from IST ($x = v$) and the trend shock stems from TFP ($X = Z$). This specification was selected because it is one of the nine specifications where the model predictions best match the full-sample kurtosis characteristics of business-cycle variables that were not used as targets in the dynamic calibration. After the results from this specification are presented in detail, a following section briefly presents the predictions of other specifications to illustrate how some specifications can come close to fitting the totality of data as well as the preferred specification, while others fail along several dimensions.

The full-sample kurtosis results of the model specification is reported in Table 3. While the model exactly matches the full-sample kurtosis of output observed in the data by design, the

kurtosis predictions of the model with respect to the other business-cycle variables are not that far off from the data. The model predicts a slightly higher kurtosis of consumption and labor hours than what is observed in the data, with the former kurtosis prediction being significantly different than 3 in the model when neither test is rejected in the data. The model also underpredicts the kurtosis of investment, but the prediction still tests significantly different than 3 as in the data. Note that all of these predictions of excess kurtosis are stemming from one single shock to TFP trend. The third row of the table reports the kurtosis predictions of the model when the shock to TFP trend is removed and only the shocks to IST levels remain. Since the model in this instance is only being driven by one traditional source of Gaussian shocks, it is not surprising that the kurtosis predictions are not significantly different than 3. This shows that any significant excess kurtosis observed in the model is exclusively due to the single, plausibly-sized shock to trend.

Table 3: US Data vs Model

| | Output | Consumption | Investment | Hours |
|---|---|---|---|---|
| **Data** | 3.86* | 3.47 | 4.53* | 2.82 |
| **Model** | 3.86* | 3.88* | 3.63* | 3.48 |
| $g^Z = 0$ | 2.97 | 3.13 | 3.26 | 2.98 |

*Notes*: * indicates rejection of $H_0 : kurtosis = 3$ with 90% confidence.

We performed the rolling-window analysis on the model simulations and illustrated the results along with the data in Figure 6. Both lines in the figure represent the kurtosis of a 60-quarter window beginning at the indicated start date, and a dot indicates a kurtosis value significantly different than 3 at the 90-percent level (i.e., a critical value of 4.04). In addition, since the kurtosis of each rolling window for the model is the average of 1000 model trials, the gray band illustrates the *inter-quartile range* of observed values. In other words, the lower (upper) value of the gray band indicates the 25th (75th) percentile kurtosis value observed across all model trials for that particular window.

As the figure illustrates, the simulated results have a great deal in common with the data. First, there is no excess kurtosis in windows not including the trend shock (assumed to take place

in 1983:1). Second, almost all simulated windows containing the trend shock display significant excess kurtosis. Third, both the data and model results share the steep sharp drop in kurtosis observed in 1983:2.

While these similarities between the model and data are rather surprising given that the model is attributing them to one single trend shock, there are some discrepancies. First, there is much more significant excess kurtosis than what is observed in the data. Second, the largest kurtosis in the model is observed in the earlier windows containing the trend shock while the largest kurtosis in the data is observed in the period immediately before the assumed growth shock (1982:4).

There are several possible reasons behind these discrepancies, ranging from the existence of multiple trend shocks in the data to stochastic volatility. While considering multiple trend shocks lacks discipline and is beyond the scope of the current analysis, the impact of stochastic volatility on the model predictions can be assessed. Since kurtosis is the ratio of the fourth-central moment (i.e., $\sum_{i=1}^{n}(x_i - \bar{x})^4$) and the standard deviation raised to the fourth power ($\sigma^4$), the kurtosis results can be decomposed to see how the model performs on these individual components.

Figure 7 illustrates the same model and data comparison of Figure 6 in the top panel, and includes the fourth central moment of each rolling window in the middle panel, and the standard deviation of each rolling window in the bottom panel. The figure shows that the model does a markedly better job at tracking the individual components of kurtosis. Both the standard deviation and fourth central moment display sharp upticks around 1967, the first windows including the 1983:1 shock in the model. Both components also display sharp declines around 1983:2. One discrepancy is that both components in the data begin to increase before 1962, which could indicate trend shocks occurring prior to our selected date. Another discrepancy is the gradual decline in both components between 1968 and 1982 observed in the data, while the predicted increase in volatility in the model is consistent across all windows that include the single trend shock. This comparison shows that while a trend shock provides a source of stochastic volatility, the model dynamics cannot capture the exact degree of stochastic volatility observed throughout the episode.

Figure 8 illustrates the same rolling-window comparison between model and data, of kurtosis and its components for consumption, investment, and labor hours. The top row of panels indicate an inability of the model to explain the stochastic kurtosis observed in the data for these variables. However, the remaining panels again illustrate the sharp rise in the components around 1967 and a sharp decline around 1983. The model ultimately underpredicts the increase in the components for consumption and overpredicts the increase in components for labor hours. While these discrepancies can be attributable to the same stochastic volatility delivering the discrepancies in output, they might also be attributable to the family of preferences necessary for balanced growth.

## 4.4 Alternative Filters

While the previous section considered the level and trend shock combination that led to the model predictions best matching the HP-filtered US data, this section briefly considers the predictions of the model when considering the other transformations of the variables considered in Figure 2. In other words, taking the same model specification considered in the previous section and maintaining the same parameter values, we transform the data into annualized growth rates as well as using the Hamilton (2018) filter. Since the same, true data was subjected to three filters in Figure 2, we thought it interesting to see what the simulated model data predicted when subjected to the same three filters.

The full sample results of the model for all three data transformations are compared with the data in Table 4. Rows one and two are the data and model results where the data was detrended using the HP Filter, rows three and four use annualized growth rates, and rows five and six use the Hamilton (2018) filter. The model results show that the full-sample kurtosis results using growth rates is larger that those with the HP filter and all significantly different than 3, just like the data. Furthermore, the model and data suggest that there is no evidence of excess kurtosis in the full sample when considering the Hamilton filter. This last fact is particularly compelling. If excess kurtosis in the short-run is due to long-run trend shocks and an inability to fully separate the

Figure 6: Rolling-Window Kurtosis, HP-filtered Output

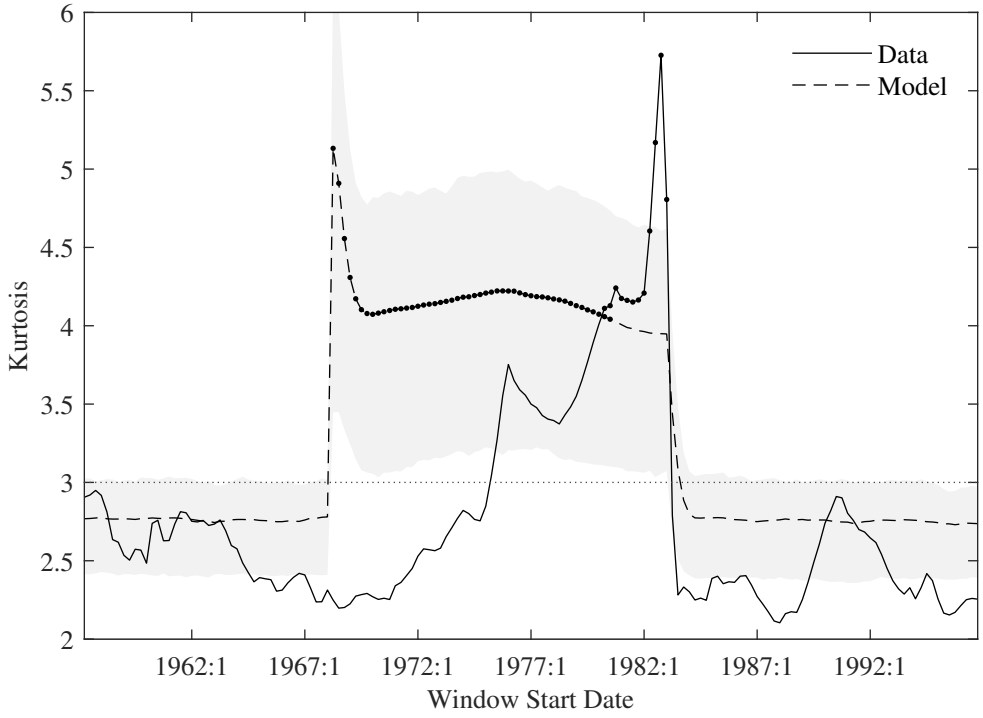


trend from the cycle, then a filter which does a better job at removing the trend should deliver a data series with less kurtosis.

Figures 9 and 10 compare the rolling-windows results of the true and simulated data in annualized growth rates and using the Hamilton filter, respectively. As in the previous illustrated results, the gray bands illustrate the inter-quartile ranges of model simulations while a dot indicates a kurtosis value significantly different than 3 at the 90 percent level. When comparing these simulated results with those using the HP filter, it is easy to see that the predicted kurtosis for the windows containing the long-run shock are significantly larger for growth rates and smaller for the Hamilton filter. This is consistent with the rolling-window analysis of US data (see Figure 2), and it is also consistent with the full sample kurtosis results detailed above.

Figure 7: Rolling-Window Kurtosis, HP-filtered Output

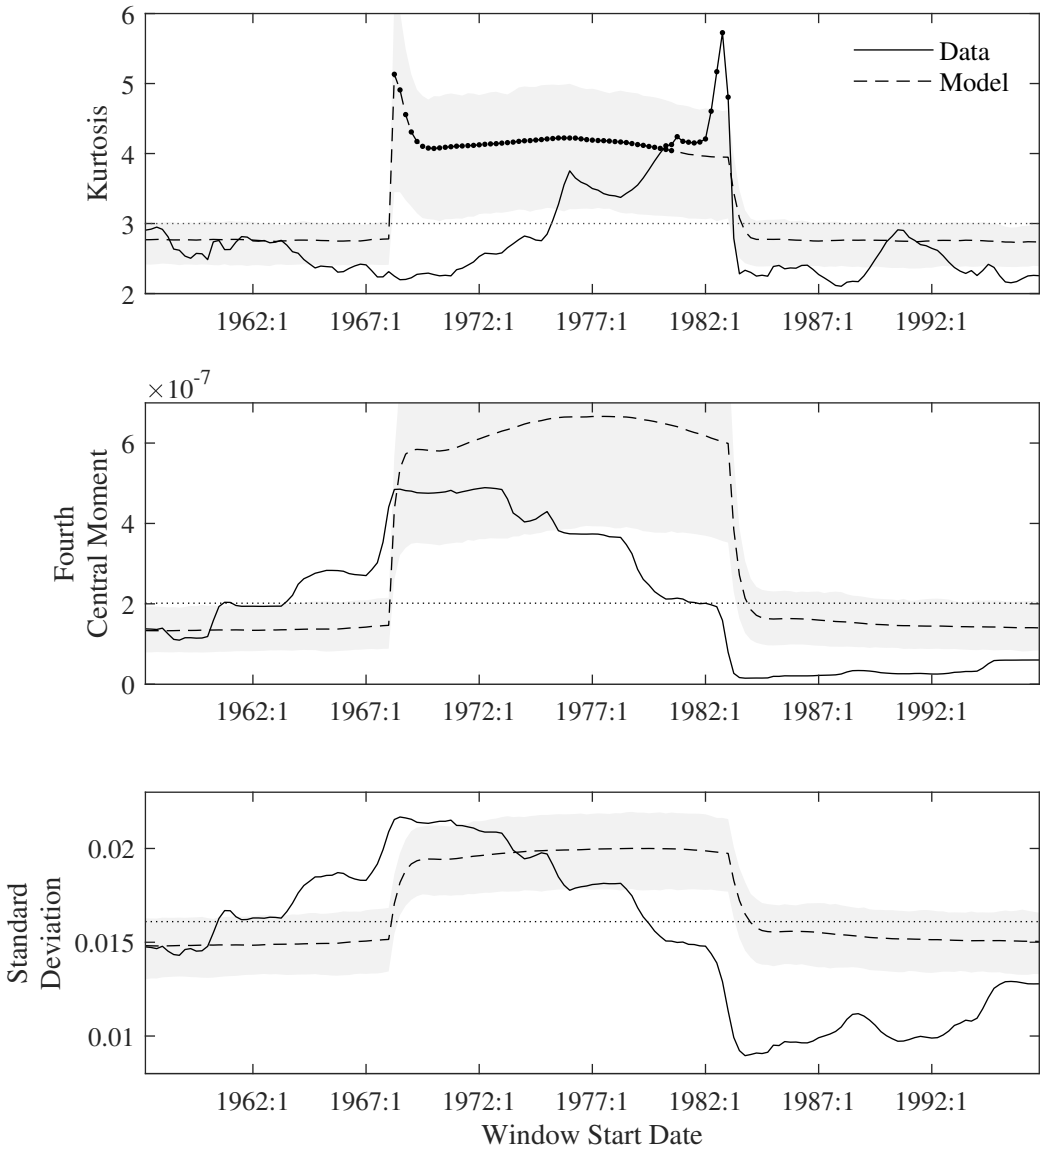

Table 4: US Data vs Model; Alternative Filters

| | | Output | Consumption | Investment | Hours |
|---|---|---|---|---|---|
| **HP** | Data | 3.86* | 3.47 | 4.53* | 2.82 |
| | Model | 3.86* | 3.88* | 3.63* | 3.48 |
| **Growth** | Data | 4.60* | 8.43* | 5.82* | 4.71* |
| | Model | 5.17* | 9.40* | 6.61* | 6.54 |
| **Hamilton** | Data | 3.31 | 3.31 | 3.39 | 3.35 |
| | Model | 3.49 | 3.07 | 3.09 | 3.14 |

*Notes*: * indicates rejection of $H_0 : kurtosis = 3$ with 90% confidence.

## 4.5 Alternative Specifications

Table 5 reports the full-sample predictions of other model specifications considering alternative level and trend shock combinations. The third row of the table reports results of the specification where the level shocks stem from IST ($x = v$) and the trend shock stems from preferences ($X = A$). This specification matches the targeted kurtosis of output by design and predicts a slightly lower kurtosis on consumption compared to our first specification, but it also predicts a kurtosis of investment further below the data and a kurtosis of labor hours further above the data resulting in all four measures having kurtosis significantly different from 3. These discrepancies are not that far apart from those of the preferred specification, and is a clear second when ranking all nine specifications in terms of matching the data. These two specifications taken together deliver two implications. First, IST shocks remain a leading driver of common business cycles as first established by Fisher (2006) and others. Second, since a trend shock from either labor-augmented TFP or labor preferences deliver similar model predictions, it provides further support for the notion that the trend disruption impacted the labor market. However, the model cannot determine if the long-run disruption was to labor supply or demand.

The remaining rows of the table consider model specifications that do not perform as well as others when compared to the data, particularly with the kurtosis of consumption and labor hours. The fourth row reports results from the specification where the level shocks and the trend shock stems from TFP, while the fifth row considers similar level shocks with the trend shock stemming

from IST. Both specifications again match output kurtosis by design, but overpredict the kurtosis for the remaining business cycle variable. In some cases, this overprediction is severe enough to eliminate this potential shock combination from consideration. While the remaining specifications were not listed for sake of brevity, we note that they provided full-sample predictions that were closer to these poor-performing specifications.

Table 5: US Data vs Model

| | Output | Consumption | Investment | Hours |
|---|---|---|---|---|
| **Data** | 3.86* | 3.47 | 4.53* | 2.82 |
| $[g^Z, \sigma^v]$ | 3.86* | 3.88* | 3.63* | 3.48 |
| $[g^A, \sigma^v]$ | 3.86* | 3.79* | 3.61* | 4.06* |
| $[g^Z, \sigma^z]$ | 3.86* | 11.86* | 5.16* | 7.04* |
| $[g^V, \sigma^z]$ | 3.86* | 34.19* | 6.34* | 9.98* |

*Notes*: * indicates rejection of $H_0 : kurtosis = 3$ with 90% confidence.

## 4.6 Sensitivity Analysis

In addition to the robustness of model results we found over several of the parameter values stated above, we found the model results to be robust to several different model specifications. In particular, the model was extended to include internal-habit persistence in consumption, capital-adjustment costs, and real-wage rigidity. For each of these model frictions, we found that increasing the degree of the friction resulted in a larger growth shock determined in the dynamic calibration. These larger growth shocks combined with market frictions delivered quantitatively similar results to those reported above from our baseline model without frictions.[15]

[15] Ascari et al. (2015) compared results from models with and without frictions but similarly-sized shocks, and found that frictions suppressed the ability of the model to deliver excess kurtosis. This feature is shared by our analysis, which is why the dynamic calibration for our models with frictions required larger shock sizes.

# 5 Conclusion

Starting with the well-documented result that many macroeconomic variables exhibit fat tails in postwar US data and that these fat tails can be attributed to large economic shocks or tail events, we provide compelling evidence that tail events are attributable to long-run growth shocks. Empirically, we examine an exhaustive number of subsets of data and provide an approximate timing of the tail events for US data. These episodes where tail events are located are well-known to contain long-run shifts in US dynamics (i.e., growth shocks). Using a standard business cycle model, we provide several results supporting the notion that a growth shock is a reasonable candidate for a tail event. First, we show that a single growth shock amidst Gaussian business-cycle fluctuations is able to deliver fat tails in simulated, macroeconomic time series. Second, the size of the growth shock is plausible insofar as it is within 1.5 standard deviations of the the business-cycle shock. Finally, while we find that growth shocks from multiple sources are able to replicate the fat tails of output we observe in the data, the model suggests that growth shocks impacting the labor market (through either the supply or demand side) combined with investment-specific technology driving the business cycle best match the totality of data.

Our results tell a cautionary tale for any subsequent research on fat tails arising in the business cycle. Our tale is that any exogenous or endogenous explanation behind fat tails should: (i) be consistent with tail events and the episodic excess kurtosis identified in the data; and (ii) be consistent with changes in predicted results across different methods of filtering. Our model results are consistent with these two guidelines, and suggest that a tail event can be associated with a significant long-run disruption stemming from the labor market. An endogenous explanation for this disruption is left for future research.

Figure 8: Rolling-Window Kurtosis, HP-filtered Variables

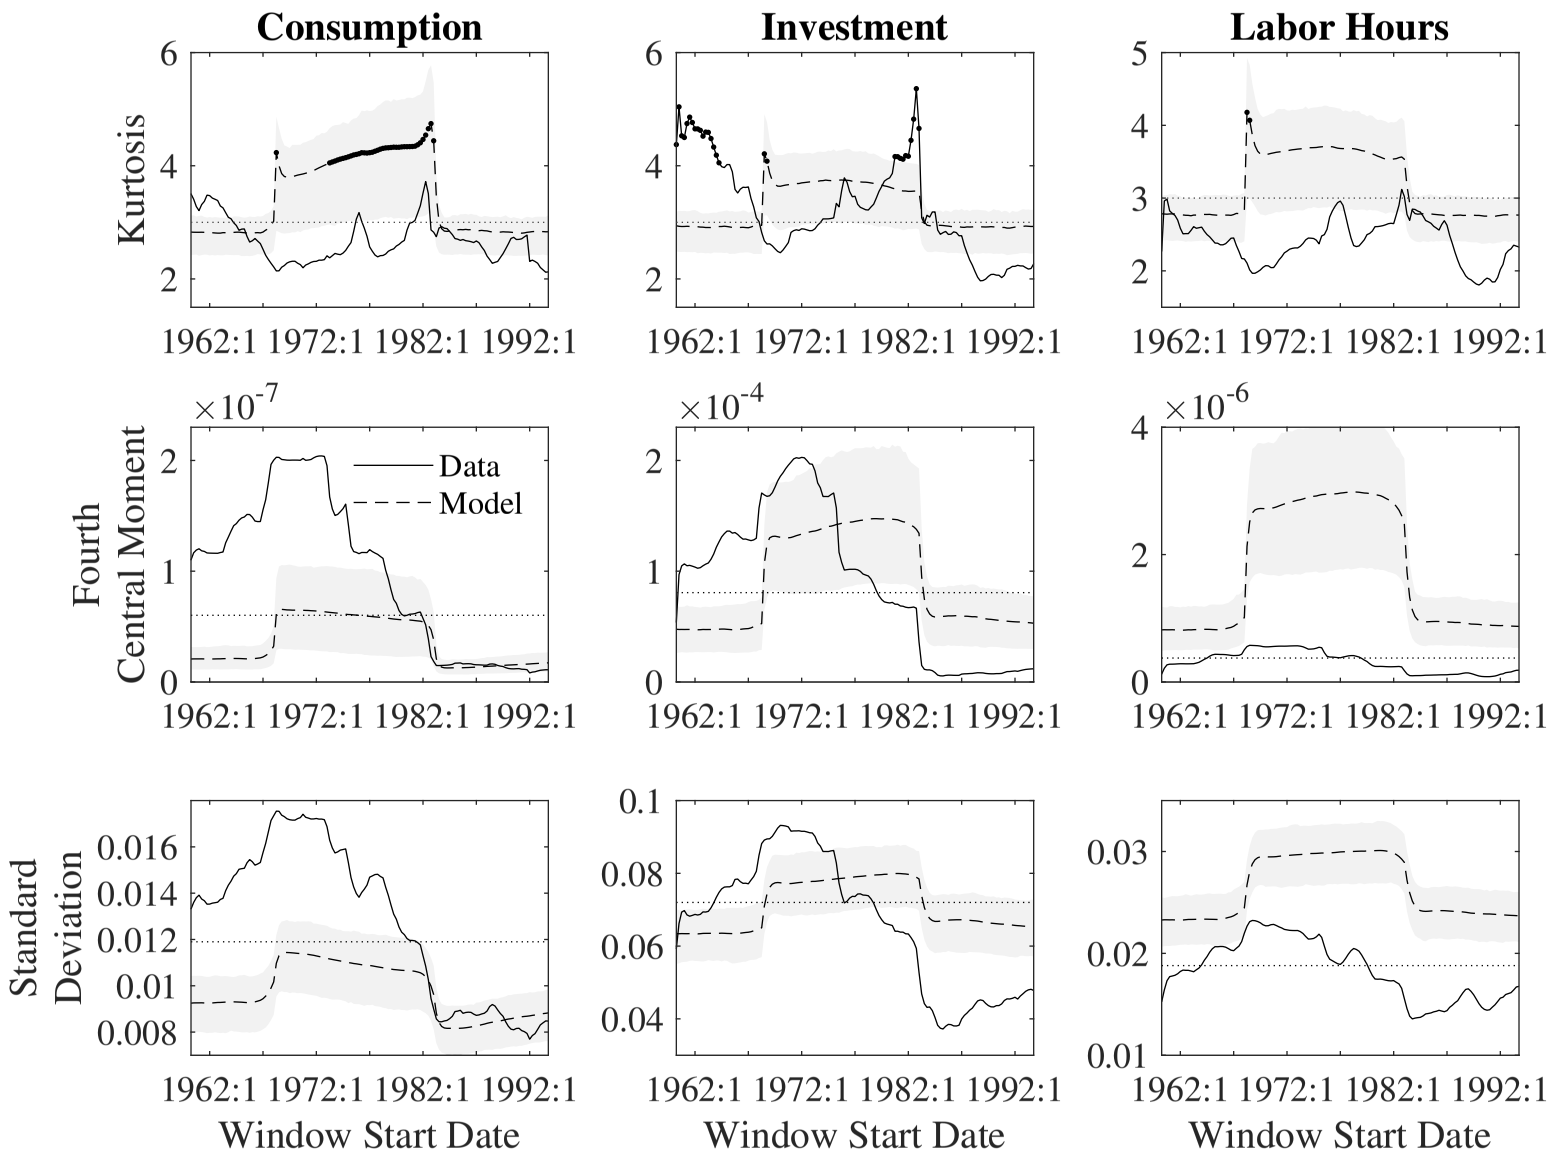

Figure 9: Rolling-Window Kurtosis, Annualized Growth Rates

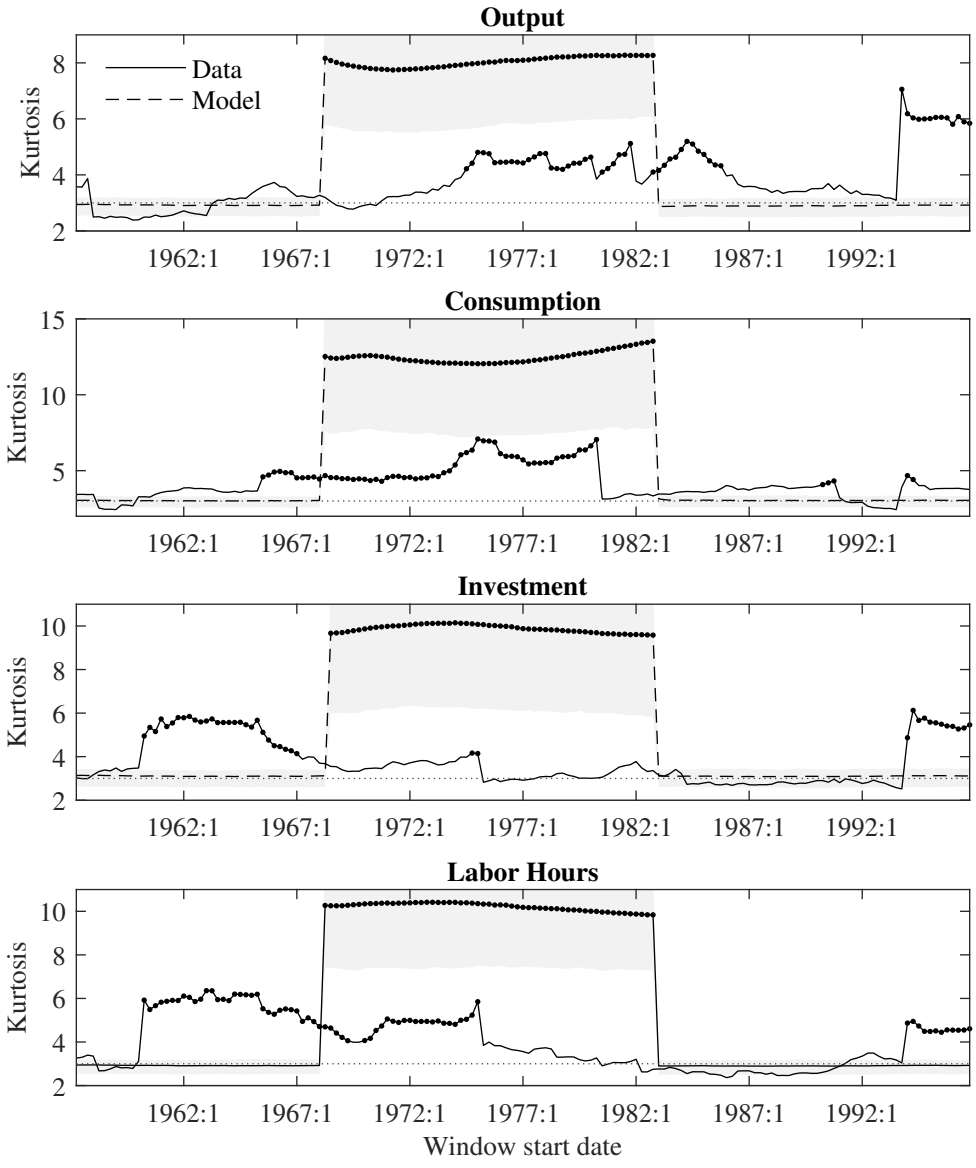

Figure 10: Rolling-Window Kurtosis, Hamilton (2018) Filter

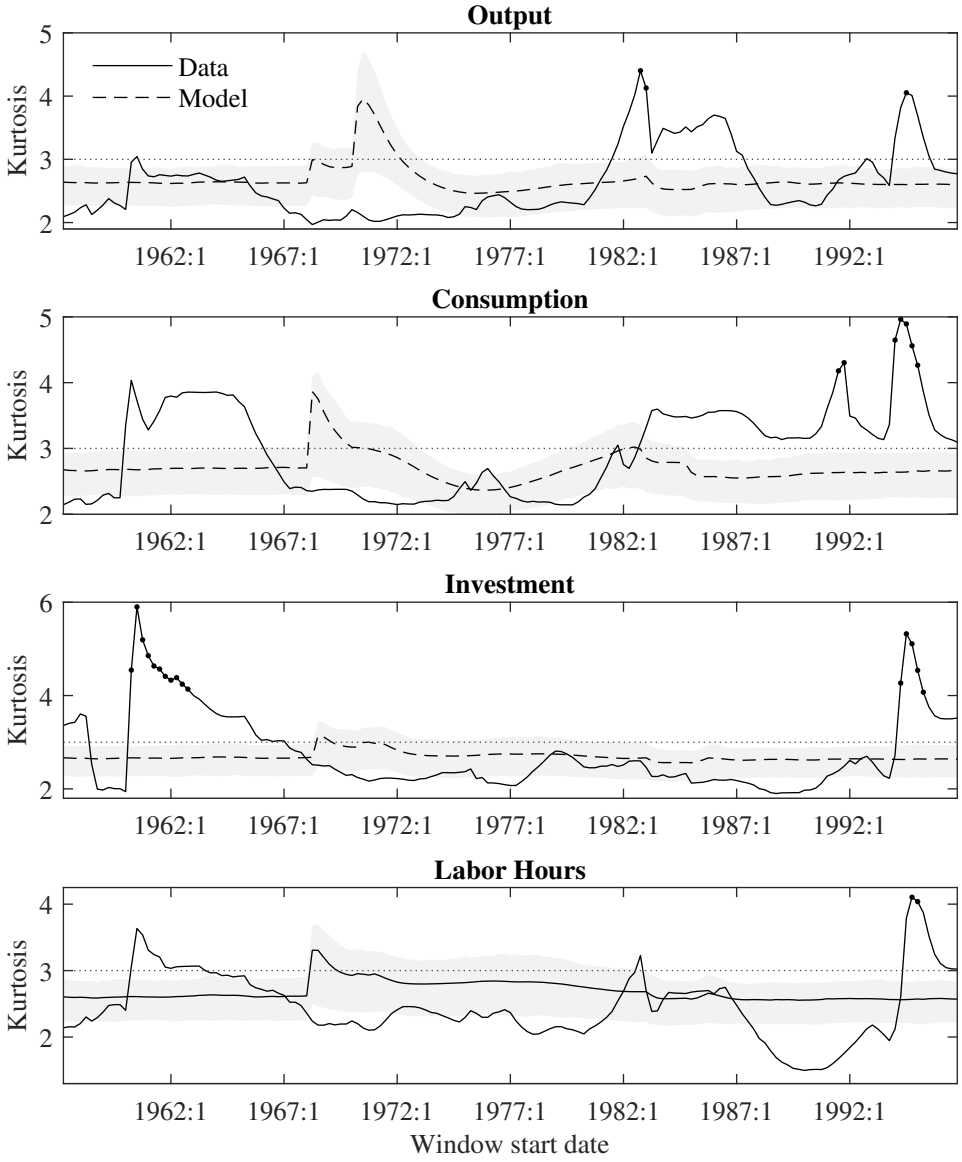

# References


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.

Auerbach, A. J. and Y. Gorodnichenko (2012): *Fiscal Multipliers in Recession and Expansion*, University of Chicago Press, 63–98.

Barro, R. J. (2009): "Rare Disasters, Asset Prices, and Welfare Costs," *The American economic review*, 99, 243–264.

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.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

Chang, Y. and F. Schorfheide (2003): "Labor-supply shifts and economic fluctuations," *Journal of Monetary Economics*, 50, 1751–1768.

Chib, S. and S. Ramamurthy (2014): "DSGE Models with Student-t Errors," *Econometric reviews*, 33, 152–171.

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.

Cogley, T. and T. J. Sargent (2001): "Evolving Post World War II U.S. Inflation Dynamics," *NBER macroeconomics annual*, 16, 331–373.

——— (2005): "Drifts and volatilities: monetary policies and outcomes in the post WWII US," *Review of economic dynamics*, 8, 262–302.

Comin, D. and M. Gertler (2006): "Medium-Term Business Cycles," *American Economic Review*, 96, 523–551.

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.

Galí, J., M. Gertler, and J. D. López-Salido (2007): "Markups, Gaps, and the Welfare Costs of Business Fluctuations," *The Review of Economics and Statistics*, 89, 44–59.

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.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997): "Long-run implications of investment-specific technological change," *The American economic review*, 342–362.

Hall, R. E. (1997): "Macroeconomic Fluctuations and the Allocation of Time," *Journal of Labor Economics*, 15, S223–S250.

Hamilton, J. D. (2018): "Why you should never use the Hodrick-Prescott filter," *Review of Economics and Statistics*, 100, 831–843.

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

Ireland, P. N. and S. Schuh (2008): "Productivity and US macroeconomic performance: Interpreting the past and predicting the future with a two-sector real business cycle model," *Review of Economic Dynamics*, 11, 473–492.

Justiniano, A., G. Primiceri, and A. Tambalotti (2011): "Investment Shocks and the Relative Price of Investment," *Review of Economic Dynamics*, 14, 101–121.

Justiniano, A. and G. E. Primiceri (2008): "The Time-Varying Volatility of Macroeconomic Fluctuations," *The American economic review*, 98, 604–641.

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.

Leduc, S. and Z. Liu (2016): "Uncertainty shocks are aggregate demand shocks," *Journal of Monetary Economics*, 82, 20–35.

Ludvigson, S. C., S. Ma, and S. Ng (2021): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" *American Economic Journal: Macroeconomics*, 13, 369–410.

Lux, T. and D. Sornette (2002): "On rational bubbles and fat tails," *Journal of Money, Credit and Banking*, 589–610.

Mishkin, F. S. (2011): "Monetary Policy Strategy: Lessons from the Crisis," .

Mumtaz, H. and K. Theodoridis (2017): "Common and country specific economic uncertainty," *Journal of International Economics*, 105, 205–216.

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.

Prescott, E. C. (1986): "Theory ahead of business cycle measurement," Tech. rep.

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

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 Data Calculations

We compute the following from FRED for the period 1984QI-2019QIV:

$$
\begin{aligned}
Y_t &= \frac{\frac{GDPC1_t}{4}}{CNP160V_t} \times 1000000, C_t \frac{\frac{PCECC96_t}{4}}{CNP160V_t} \times 1000000, \\
I_t &= \frac{\frac{GPDIC1_t}{4}}{CNP160V_t} \times 1000000, H_t = \frac{HOANBS_t}{CNP160V_t} \times 6000, \\
g_{Yt} &= \frac{Y_t - Y_{t-1}}{Y_{t-1}}, g_{Ct} = \frac{C_t - C_{t-1}}{C_{t-1}}, g_{It} = \frac{I_t - I_{t-1}}{I_{t-1}}, g_{Ht} = \frac{H_t - H_{t-1}}{H_{t-1}}, \\
y_t &= ln(Y_t), c_t = ln(C_t), i_t = ln(I_t), h_t = ln(H_t),
\end{aligned}
$$

and de-trend $\{y_t, c_t, i_t, h_t\}$ using various methods to report cyclical data statistics.

## A1.2 Alternative Data Figures

Figure A1: GDP Per Capita; 40-Quarter Rolling Windows

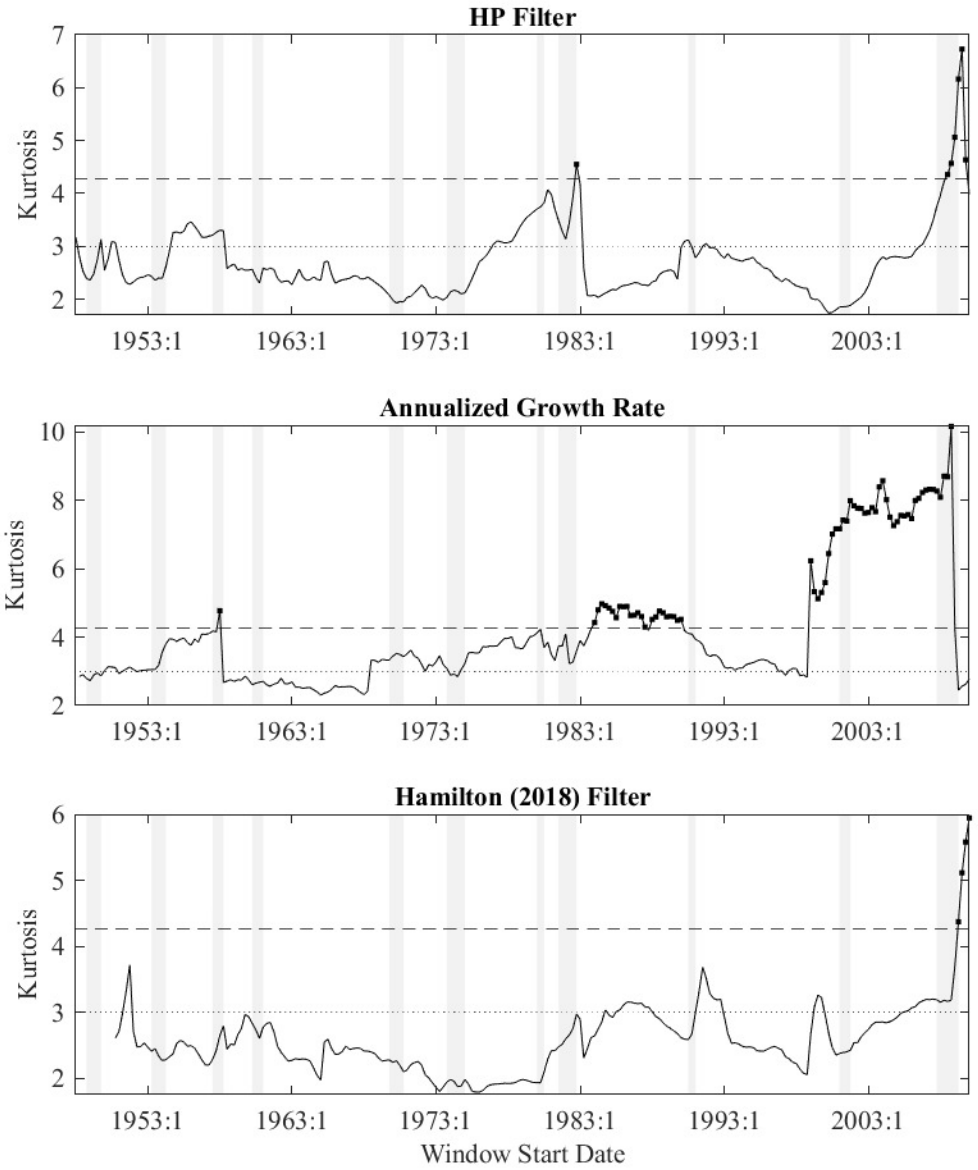
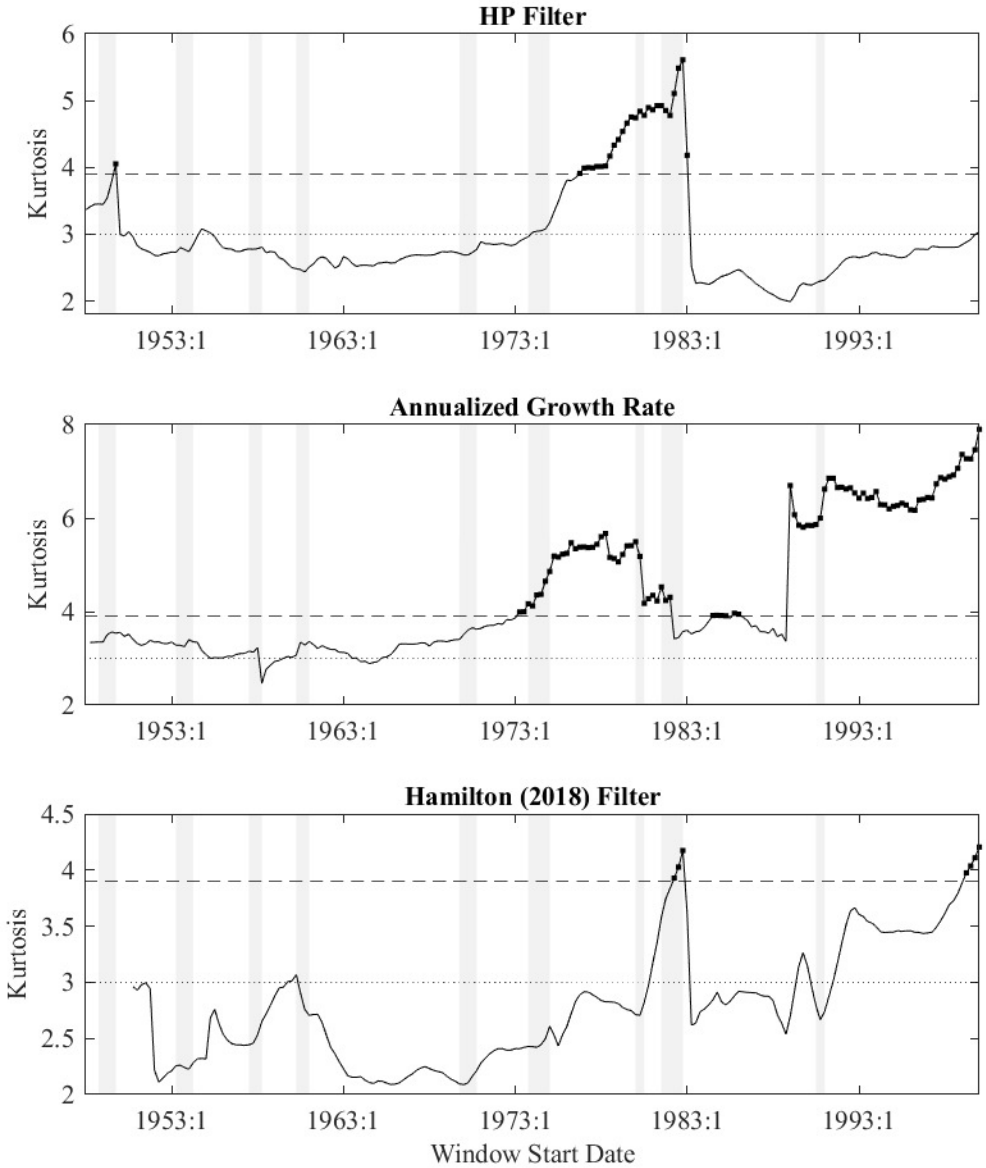
Figure A2: GDP Per Capita; 80-Quarter Rolling Windows