All Fluctuations Are Not Created Equal: the Differential Roles of Transitory Versus Persistent Changes in Driving Historical Monetary Policy

Richard Ashley
Department of Economics
Virginia Tech
Blacksburg, VA 24061-0316
Email: ashleyr@vt.edu

Kwok Ping Tsang
Department of Economics
Virginia Tech
Blacksburg, VA 24061-0316
Email: byront@vt.edu

Randal J. Verbrugge
Research Department
Federal Reserve Bank of Cleveland
1455 E. 6th St.
Cleveland, OH 44114
Email: randal.verbrugge@clev.frb.org

April 5, 2018

JEL Classification Codes: E52, C22, C32

Keywords: Taylor Rule, Great Inflation, Intermediate Target, Natural Rate, Persistence Dependence

Acknowledgement: The authors would like to thank, without implicating, Todd Clark, Ed Knotek, Jeff Fuhrer, Kurt Lunsford, Christian Mathes, Filippo Occhino, Giovanni Olivei, Ellis Tallman, and participants at the Cleveland Fed Brownbag. We have also benefitted from previous conversations with Marianne Baxter, Luca Benati, John Cochrane, Rob Engle, John Geweke, and Mark Watson.The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or of the Federal Reserve System.

Transitory Versus Persistent Fluctuations in Historical Monetary Policy

Abstract

We ask whether the FOMC treats persistent fluctuations in the unemployment rate, or in the inflation rate, differently from transient fluctuations when it makes its fed funds rate decisions. We do this through the lens of estimated policy rules which are extended to allow for (but not impose) such distinctions. Our estimation method uses real-time data in these rates – as did the FOMC – and requires no *a priori* assumptions on the nature of the processes generating either the data or the natural rate of unemployment. Also, unlike other approaches, our estimation method allows for possible feedback in the relationship. We find that, while its ability to do so improved over time, the FOMC has always distinguished persistent movements in the unemployment rate from other movements; implicitly such movements were treated as an intermediate target, and the FOMC has responded to the implied unemployment gap. In contrast, we find that the FOMC behavior vis-a-vis inflation changed starting with Volcker: its response to the inflation rate became much stronger, and it focused more intensely on very persistent movements in this variable. Our results shed light on the "Great Inflation" experience of the 1970s, casting doubt on the proposition that bad inflation forecasts or poor estimates of slack explain FOMC behavior over that time period.

1 Introduction

1.1 Background

Following its articulation in 1992 by John Taylor (Taylor 1993), a large literature has developed around the scrutiny of simple monetary policy rules. Taylor's rule took the form

$$i_t = \alpha + \beta u_t + \varphi \pi_t + e_t \tag{1}$$

where i_t is the federal funds rate, π_t is the annualized inflation rate from period $t-1$ to period t, u_t is a trendless measure of real activity (output gap or unemployment rate) in period t, and e_t is a stationary exogenous monetary shock. There are many variants of Equation (1). Theory often suggests forward-looking versions (e.g. Clarida, Gali and Gertler (2000)); real-time lags in data collection motivate the use of lagged inflation and real activity in a backward-looking monetary policy rule (e.g. McCallum (1997)), as we formulate it here; and interest rate smoothing considerations, as well as the statistical properties of i_t, motivate adding lags of i_t to the right-hand side of (1).[1] A number of studies have used variants of Equation (1) to examine and explain the "Great Inflation" of the 1970s (e.g., Athanasios 2002) or to examine FOMC behavior over time, generally finding that the central bank's policy changed markedly starting with Volcker (see, e.g. Judd and Rudebusch (1998) or Clarida, Gali and Gertler (2000); more recent studies reaching this conclusion include Debortoli and Lakdawala (2016) and Chen, Kirsanova and Leith (2017)).[2] Some studies extend the policy rule to include other variables such as bond yields (e.g., Fuhrer and Tootell 2008 or Roskelley 2017) or stock and house prices (e.g., Aastveit, Furlanetto and Loria 2017).

Despite the intensity of this research focus, econometric studies of simple monetary policy rules like (1) almost invariably impose strong restrictions on FOMC behavior that don't make sense –

[1] Consolo and Favero (2009) argues, in the forward-looking context, that the inertia is an artifact of a weak-instrument problem for expected inflation. (As will be evident below, our methods sidestep problems with instruments and identification in forward-looking policy rules (see, e.g., Jondeau, Bihan and Galles (2004)), though information from forecasts are incorporated gracefully.) Rudebusch (2002) disagrees with the interest rate smoothing interpretation; using evidence from the term structure, he shows that monetary policy inertia is more likely due to persistent shocks that the central bank faces.

[2] Using the the time-varying parameter framework, Cogley and Sargent (2005), Kim and Nelson (2006) and McCulloch (2007) come to a similar conclusion. Not all studies reach this conclusion; e.g., Sims (2001) and Sims and Zha (2006) find that there is less evidence for significant changes in the reaction coefficients ϕ if one allows for time-varying variance in the monetary policy shock.

an imposition that threatens the resulting inference, and may render conclusions about such things as the origins of the Great Inflation suspect. What are these restrictions?

First, these specifications presume that FOMC members would respond to extremely persistent movements in the unemployment rate – likely driven by supply-side factors – in the same way that they would respond to business-cycle-driven fluctuations in that variable. As this issue is well-known, a fairly common means of addressing it is to replace u_t in Equation (1) with an "unemployment gap" $(u_t - u_t^*)$, as (for example) in McCulloch (2007).[3] To take this approach, one needs an explicit estimate of u_t^*, the time-varying "long-run normal" level of unemployment. But the selection of any particular formulation for estimation of u_t^* may itself compromise inference. For instance, if one's estimate of u_t^* derives from a two-sided filter – such as an HP filter or a symmetric moving average filter – this will distort the coefficient estimate ϕ_u (see Ashley and Verbrugge, 2009b). Further, estimates of u_t^* are inherently problematic in that they are estimated very imprecisely, are subject to large revisions, and typically hinge on untested (and perhaps untestable) auxiliary assumptions about the natural-rate data generating process (such as an explicit formulation of its persistence), any or all which may well be substantially incorrect. As might be expected, then, the u_t^* estimates can vary across concepts and methods widely (see Tasci and Verbrugge, 2014). Use of an output gap instead of an unemployment gap in Equation (1), also common, does not improve matters.

In any case, selection of a particular u_t^* measure implies an assumption that FOMC members were in broad agreement about that particular measure. At least during the 1970s, that seems unlikely. The natural rate of unemployment as a concept came into popular parlance following Friedman (1968), but it was not clear how to measure it.[4] Soon afterwards, the stagflation in the early 1970s was thought to indicate shortcomings in the concept, and prompted the introduction in 1975 of a related but distinct concept, the NIRU (Modigliani and Papademos, 1975), which was later re-termed the NAIRU.[5] Textual analyses of FOMC minutes and transcripts, such as

[3]See Orphanides (2008) and Knotek et al. (2016) for a more general discussion. Other notable studies, such as Ball and Tchaidze (2002) and Detmeister and Babb (2017), do not use an unemployment gap specification.

[4]Hetzel (2017), an author who writes extensively about the historical evolution of central banking, asserts that during the 1970s there was a "general belief that 4 percent unemployment represented full employment" (p.13). Having said this, as our results indicate, such a belief would not necessarily preclude a differential FOMC response to persistent unemployment rate fluctuations.

[5]See Espinosa-Vega and Russell (1997).

Chappelle, McGregor and Vermilyea (2004) or Weise (2012), indicate that the 1970s FOMC was concerned about high and rising unemployment and was grappling with the notion of what level was sustainable, but there was nothing like a settled concept or estimate guiding deliberations. Given this fact, a better approach would be to determine from the data whether and how the FOMC was actually responding to movements in the unemployment rate. Were they, for example, responding to more persistent movements in the unemployment rate in a way that was systematically different than their response to business cycle fluctuations in that variable? Did this aspect of their behavior change starting with the chairmanship of Volcker? In this study, we apply methods that let us credibly answer this question.

Second, but equally fundamentally, these specifications insist that the FOMC responds to highly transient movements in the inflation rate in exactly the same way that they respond to persistent fluctuations. A crucial issue in inflation measurement is noise: in the short run, measured inflation is often subject to large transitory influences, which can affect an aggregate price index for long periods. Policymakers have argued – e.g. Mishkin (2007) – that the central bank should therefore not respond to transitory fluctuations in inflation. If the assumed policy rule does not take this differential response into account, Equation (1) will again be seriously misspecified. Some analysts have attempted to address this issue by making use of so-called "core" inflation measures in Equation (1). This expedient is valid, however, only if all movements in the core inflation measure are identically persistent, which is emphatically not the case; for example, see Bryan and Meyer (2002) and Dolmas and Wynne (2008).[6] Note that the issue we raise here is distinct from whether π_t should instead enter Equation (1) in terms of an "inflation gap," where this gap is the difference between π_t and an inflation target, π^*. As in the case of u_t, ideally one would want to let the data inform us as to how the FOMC was actually responding to movements in the inflation rate – e.g., whether it was ignoring transient movements in this variable, or focusing on the most persistent movements – and whether this behavior changed with the chairmanship of Volcker.

[6] Indeed, in March 2017, a large price drop in cellular communications resulted in a notable drop in the core PCE index. This caused a twelve-month drop in year-over-year core inflation. Yet FOMC communications make it clear that this transient movement was going to be ignored. Such analysis is not new; for example, the FOMC Greenbooks in early 1969 explicitly highlighted transitory inflation fluctuations deriving from Federal pay increases – clearly the FOMC regarded these as transitory.

1.2 The Approach Used Here and Our Broad Findings

In this paper, we relax the aforementioned restrictions implicit in (1), and allow the data to speak as to how the FOMC historically has responded to fluctuations with various persistence levels in the unemployment rate and in the inflation rate. Our approach delivers new insights into FOMC behavior and into how it changed starting with the chairmanship of Volcker.

We estimate the central bank's assumed monetary policy rule using a new method proposed by Ashley and Verbrugge (2009b).[7] This method allows us to test whether, and quantify how, the FOMC in fact distinguished between fluctuations with differing persistence levels. Essentially, we use a one-sided filtering technique to partition the real-time unemployment rate and inflation rate into components with differing levels of persistence. It is then straightforward to determine whether the FOMC responded (in real time) to the persistent part in the same manner that it responded to the less persistent part, or to the high-frequency part. The technique is briefly described as follows (for more details, refer to the Technical Appendix here, or to Ashley and Verbrugge 2009b).

We use moving windows to filter the real-time data at each time t, partitioning the time t observation into various frequency (or persistence) components. One-sided filtering is necessary for two reasons: first, two-sided filtering cannot be conducted using real-time data; and second, two-sided filtering results in distorted coefficient estimates and generally destroys the ability to make causal statements.[8] The precision of the partitioning is substantially enhanced by extending the data in each window with forecasts or "projections" (see also Dagum (1978), Mise, Kim and Newbold (2005) or Clark and Kozicki (2005)). Thus, for example, to partition the unemployment rate at time t, we apply the Ashley-Verbrugge filter to a rolling window of 96 observations ending in t, augmented with a set of 1 through 24-month forecasts, starting at time t+1. The filter is then applied to this 120-month window, and the time-t decomposition is saved. (The inflation rate is partitioned in the same manner, except in that case, a 42-month window is augmented by 18 months of forecasts. We argue below why distinct window lengths for the two variables are

[7] The Technical Appendix describes the latest version of this method.

[8] Two-sided filters (such as the HP filter) applied to both the explanatory and dependent variables, as in classical RBC studies or in recent New Keynesian DSGE modeling, distort relationships amongst variables which are contemporaneously cross-correlated or in a feedback relationship, because two-sided filtering inherently mixes up future and past values of the time series; see Ashley and Verbrugge 2009b for a more detailed exposition of this point. For the same reason, and because such calculations are incompatible with the use of real-time data, two-sided cross-spectral estimates or filtering with wavelets are similarly ruled out.

appropriate.) Figures 1 and 2 display the resulting components of each of these two variables, in each case partitioning the data by persistence level into three persistence components.[9]

Obviously, we are not arguing that the FOMC made use of techniques from the future in order to guide their decisionmaking.[10] Instead, our analysis asks whether the FOMC, using the tools it had on hand at various points in time (such as forecasts, judgment, various trend estimates, or detailed studies of particular unemployment rate or inflation movements), in fact differently reacted to persistent versus transient movements in u_t and π_t in its decisionmaking. Indeed, our method uses real-time filtering of the data to assist in identifying persistent versus transient movements in roughly the same manner that an informed observer might do using judgment.[11] In particular, we are using new tools to better understand and describe the behavior of the central bank, allowing the data to inform us as to the manner in which the central bank has actually responded to fluctuations in these macroeconomic variables.

While the original Taylor-type monetary policy response function is attractive in its simplicity, our extension of it broadens its generality and descriptive power, yielding novel results. By allowing the data a chance to tell us if the FOMC treated apparently persistent fluctuations differently from transient fluctuations, we obtain a clearer picture of how the FOMC's actual behavior has evolved over time – e.g., how it differs for the Martin-Burns-Miller (MBM) period (roughly, March 1960 to August 1979) versus the Volcker-Greenspan-Bernanke (VGB) period (here taken to run from September 1979 to August 2008).

Our results are surprising along some dimensions, but in other respects accord well with intuition and some other accounts. There are three major conclusions. First, estimates of (1) as it stands imply that the FOMC was unresponsive to unemployment rate fluctuations in the VGB period; as we explain below, our result underscores the desirability of allowing for a policy response distinction

[9] A fourth component distinguishing fluctuations varying within a quarter proved unnecessary.

[10] While economists were already making use of the conventional tools of spectral analysis by the early 1970s, given the early cogent critique of Granger (1969) and the relative paucity of that sort of research, it is safe to say that such methods were not at the forefront of FOMC decisionmaking in the 1970s – and indeed, still aren't.

[11] This appears to be consistent with the viewpoint of Meyer, Venkatu and Zaman (2013), who state: "By specifying the inflation threshold in terms of its forecasted values, the FOMC will still be able to 'look through' transitory price changes, like they did, for example, when energy prices spiked in 2008. At that time, the year-over-year growth rate in the Consumer Price Index (CPI) jumped up above 5.0 percent but subsequently plummeted below zero a year later when the bottom fell out on energy prices. At the time, the Committee maintained the federal funds rate target at 2.0 percent, choosing not to react to the energy price spike."

between persistent movements in the unemployment rate versus other fluctuations. Second, by allowing the data to inform us about FOMC responses to persistent fluctuations versus higher-frequency fluctuations, our striking finding is that the FOMC did *not* seem to be struggling with mismeasurement of u_t^* in the 1970s, as some analysts have argued. Instead, the FOMC has *always* ignored extremely persistent fluctuations in the unemployment rate – which one might largely identify with natural rate fluctuations, although an "intermediate target" interpretation is possible, as we outline below – a response that stands in sharp contrast to its strong response to business-cycle fluctuations in the unemployment rate. This suggests that some previous work, which has focused on particular estimates of an output gap (rather than on an unemployment gap), may have come to erroneous conclusions about FOMC behavior in the 1970s. Our third major finding is that, in contrast to its continuity in response to the unemployment rate over both periods, the FOMC response to inflation movements changed in not just one, but two ways, starting in the VGB period. First, the VGB FOMC became more aggressive in its response to inflation – something other authors have noted (e.g., Clarida, Gali and Gertler 2000). But second, it also became much more focused on fighting the *persistent* fluctuations in inflation, and ignoring more transient movements.[12]

Our results thus shed some light on the origins of the Great Inflation. They imply that the "ideas" hypothesis – that FOMC errors in the 1970s were due to erroneous beliefs about the structure of the economy, including an unrealistically low estimate of the natural rate of unemployment (e.g., Orphanides 2002, Romer and Romer 2002, Romer 2005) – cannot fully explain FOMC policy during that time period. Our findings point in this direction because we show that FOMC behavior with respect to the fed funds rate response to the unemployment rate has remained rather stable: irrespective of any belief it may have had about the natural rate of unemployment, it has always ignored the most persistent unemployment rate fluctuations. Implicitly, the FOMC has always implicitly formed an unemployment gap in the same manner, with the lowest-frequency fluctuations serving the role of an intermediate target (in the words of Fuhrer).[13] So what can explain the high

[12] While the coefficient estimates themselves are consistent with a differential-inflation-persistence-response interpretation in the MBM period, formal statistical tests fail to reject the null hypothesis that the FOMC responded to all inflation fluctuations in the same way during this period. We interpret this as indicating that the FOMC during the VGB period became either more focused or more accurate at identifying persistent inflation fluctuations. Broadly speaking, this accords with the analysis of Goodfriend and King (2013); in discussing this paper, Svensson (2013) states "...that a major explanation for the Great Inflation could be a small weight on inflation stabilization and drifting inflation target does not seem so far-fetched." (p. 213)

[13] Jeffrey Fuhrer suggested this language in his discussion of Fuhrer and Olivei (2017) during the Boston Fed's annual central banking conference in October 2017.

inflation of that period? On the face of it, our results suggest that a change in policy preferences – possibly driven by political pressures (see, e.g., Chappell, Havrilesky and McGregor (1993), Meltzer (2009), Weise (2012), or Levin and Taylor (2013)) – is a far more likely explanation. What about the conjecture that FOMC behavior in the 1970s resulted from inaccurate inflation forecasts (see, e.g., Orphanides 2002, Levin and Taylor 2013 or Fuhrer and Olivei 2017)? This mechanism could certainly operate in conjunction. However, our results indicate that the actual behavior of the FOMC in the 1970s is more consistent with its using real-time CPI inflation data, where identification of the persistent part of inflation fluctuations uses information from simple-but-reliable inflation forecasts.

2 Empirical Results

2.1 Data Description and Plan of This Section

We use real-time data on the civilian unemployment rate (u_t) and the 12-month growth rate in the (real-time) Consumer Price Index inflation rate (π_t),[14] taken from St. Louis Federal Reserve Bank ALFRED data set. (In our investigation of the experience 1970s, for comparison purposes we also make use of real-time quarterly GNP deflator data, and inflation forecasts deriving from FOMC Greenbooks, both available from the Real-Time Data Research Center at the Federal Reserve Bank of Philadelphia.) For our projections of the unemployment rate, a case where univariate models can be slow to detect turning points in the data, we also make use of the Survey of Professional Forecasters projections of this variable, also available from the Federal Reserve Bank of Philadelphia.

As the backward-looking filtering employed here uses up a good deal of sample data as a 'start-up' period, our regressions begin in January 1965. Our sample period ends in August 2008, just prior to the point when the sample variation in i_t becomes minimal. The data we are analyzing correspond closely to those which were available to the FOMC at the time it set the federal funds

[14]Specifically, we use the inflation rate defined as the 12-month growth rate in percentage terms – i.e., $100ln(CPI_t/CPI_{t-12})$ – where CPI_t is the non-seasonally adjusted Consumer Price Index for urban wage earners and clerical workers until February 1978 and the non-seasonally adjusted Consumer Price Index for all urban consumers thereafter. The value for CPI_{t-12} in π_t is that which was available when CPI_t was released.

rate (i_t).[15]

Our treatment of the federal funds rate, and the information available for the FOMC during its meeting each month, warrants some discussion. Over most of our sample, the federal funds rate experienced notable day-to-day changes, and was also subject to end-of-month effects as banks addressed regulatory constraints. Our goal is to estimate the FOMC reaction function, which models the FOMC fed funds rate decision in their meeting, in reaction to or based upon the information that the FOMC had available when it made the decision. However, FOMC meeting dates – and conference phone conversations at which decisions could also be taken – did not occur on the same day each month.[16] Hence, for each month we estimate the monthly fed funds rate by taking a trimmed-mean estimate of the fed funds rate over the six business days following an FOMC meeting, trimming the highest and lowest daily rates of that period, and taking the average over the remaining four observations. If there was no meeting that month, we implicitly assume an FOMC meeting on the 20th day of the month, at which the rate was left unchanged. For the monthly unemployment rate and inflation rate series, we then utilized the information that was actually available immediately prior to the FOMC deliberations that month. Most commonly, the FOMC had available the unemployment rate from the previous month, and the CPI inflation rate from two months prior. However, for meetings that occured late in the month, CPI inflation data from the previous month were often available. (We describe below a robustness exercise which used GNP deflator information aligned to the same meeting dates.)

Following Clarida, Gali, and Gertler (2000), we consider two sub-sample periods. The first of these is January 1965 to August 1979, which roughly corresponds to the Martin-Burns-Miller period and is here denoted 'MBM'. The second sub-sample runs from September 1979 to August 2008; it covers Volcker's, Greenspan's and part of Bernanke's tenures; it is here denoted 'VGB'. Most of the VGB period is also referred to as the 'Great Moderation' – see McConnell and Perez-Quiros (2000) and Kim and Nelson (1999) – as it is characterized by low variance in most macroeconomic variables. Since the onset of the Great Recession, of course, most macroeconomic variables have

[15] Source: http://research.stlouisfed.org/fred2/. See Orphanides (2001) for evidence that estimated monetary policy rules are likely not robust to the vintage of the data.

[16] In fact, occasionally, there were two or meetings in a single month. In those cases, we selected one of the meetings, preferentially the meeting or conference call at which a rate change took place. Earlier in the sample, when Greenbook forecast information was more variable, we also gave consideration to meetings at which richer forecast information was included in the Greenbook.

become more volatile.

Each monthly observation on u_t and π_t is here decomposed into frequency (persistence) components, as noted above and described in the Technical Appendix, using moving windows. In specifying the length of the moving windows used in decomposing the real-time unemployment rate, we note that the natural rate of unemployment is thought to be quite slowly-varying. For instance, Boivin (2006) uses a five-year moving average of the unemployment rate as a proxy for the natural rate. A shorter window – of length, say, 36 months – would risk including business cycle effects within the lowest-frequency unemployment rate component, and likely not properly distinguish very persistent supply-side pressures on the unemployment rate from less persistent business-cycle-related influences on the unemployment rate. Put differently, this short window length choice would impose the restriction that the central bank responds in much the same way to an unemployment rate fluctuation with a reversion period of 36 months as it does to a fluctuations with substantially larger reversion periods – e.g., of 5 years, or even 10 years. Hence, we conjectured that a ten-year window would likely result in a satisfactory decomposition for the unemployment rate. As will be evident below, the outcome of the empirical analysis presented here will inform us as to whether or not a shorter window would have been adequate. In specifying the length of the moving windows used in decomposing the real-time inflation rate, in contrast, we considered that a central bank might very well react differently to a fluctuation in the real-time inflation rate which has persisted just 12 months – as compared to one which has persisted ca. 36 or 48 months – but that it seems unlikely *a priori* that it would attend to a 60-month-long fluctuation substantially differently than to one which has persisted substantially longer than this.[17] We thus judged that a moving 60-month window would suffice for the inflation rate.[18]

[17] A (centered) 36-month moving average is often used as an estimate of trend inflation; see, for example, Cecchetti (1997) or Brischetto and Richards (2007); Giannone and Matheson (2007) and Higgins and Verbrugge (2015) argue for using even shorter moving averages.

[18] Our empirical results are not particularly sensitive to specifying somewhat shorter (or longer) moving windows for use in decomposing u_t and π_t. As noted in the Technical Appendix, this filtering provides a usefully accurate decomposition for the last sample observation in a moving window only if the sample observations in the window are augmented by a number of projected observations. We use 18 months of projections for the inflation rate partitioning, and 24 months of projections for the unemployment rate partitioning. These choices imply that roughly 80% of each window consists of sample data. Based on minimizing BIC as a selection criterion, the π_t projection forecasting model chosen for use here combines a random walk and an ARMA(6,2) process; the u_t forecasting model combines forecasts from the *Survey of Professional Forecasters* with an AR(2) process. As with the window lengths themselves, we find that our empirical results not sensitive to minor changes in the number of projections or the form of the projection forecasting model used, but it *is* necessary to use at least 12 projection months and some kind reasonable projection model for filling out each window. RATS code implementing this decomposition methodology

In previous versions of the paper, we have explored more disaggregated partitions of the data – e.g., splitting u_t and π_t each into ten components with differing persistence levels and estimating a distinct policy rule coefficient for each of these components. We also explored more parsimonious specifications in which the parameter variation across the ten components was constrained by a quadratic or cubic polynomial, as in Almon (1965). While the latter is a bit more satisfactory (and the former method still yields reasonable results), it is abundantly clear that a tripartite partition suffices, yielding the same economic inferences in a much more transparent manner. Accordingly, here we only quote results from a tripartite partition of each variable.

In particular, here we isolate the most persistent fluctuations that are resolvable given the length of the filtering window – thus the most persistent fluctuations in the unemployment rate consist of all fluctuations that take longer than 120 months to complete, while the most persistent fluctuations in the inflation rate consist of all fluctuations that take longer than 60 months to complete. We will denote these here as the "very persistent" fluctuations.[19] Next, we split out the "transient" fluctuations of each variable, where these consist of all fluctuations that take 12 months or less to complete. We will be able to investigate whether the FOMC effectively ignored these fluctuations. Below we refer to the remaining component of u_t or π_t as the "moderately persistent" fluctuations.

is available from the authors. For non-RATS users, we also have available a ready-to-use Windows-based executable which implements the decomposition using simple AR(p) projection; in this case all that the user needs to do is specify the order p, the number of projections, and the window length – the program then produces the maximal number of frequency (persistence) components allowed by the window length; the user can then aggregate these components into frequency (or persistence) bands.

[19] As we explain below, the extremely persistent component will include all frequency-0 fluctuations, in addition to all fluctuations corresponding to reversion periods longer than the filtering window. Analysts may choose to include fluctuations with lesser persistence, although here we do not do so. In terms of the frequency domain analysis that underlies these decompositions, a 120-month reversion period corresponds to the sinusoidal frequencies proportional to 1/120; see the Technical Appendix for details.

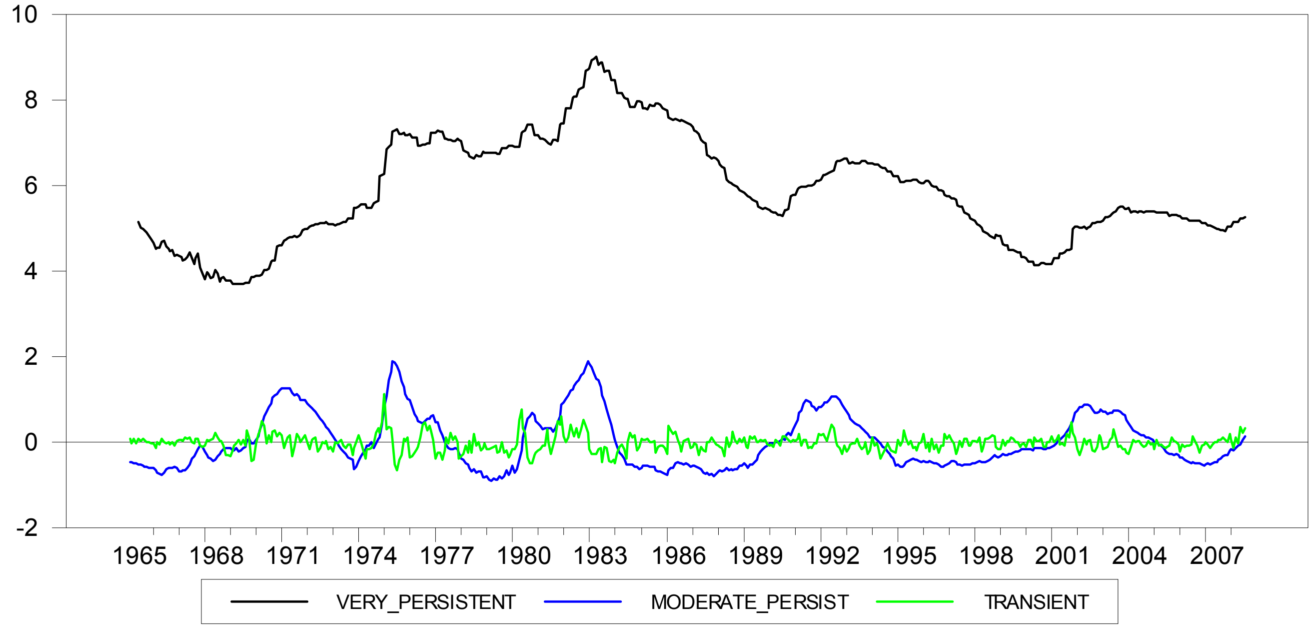

Figure 1: Components of u_t By Persistence

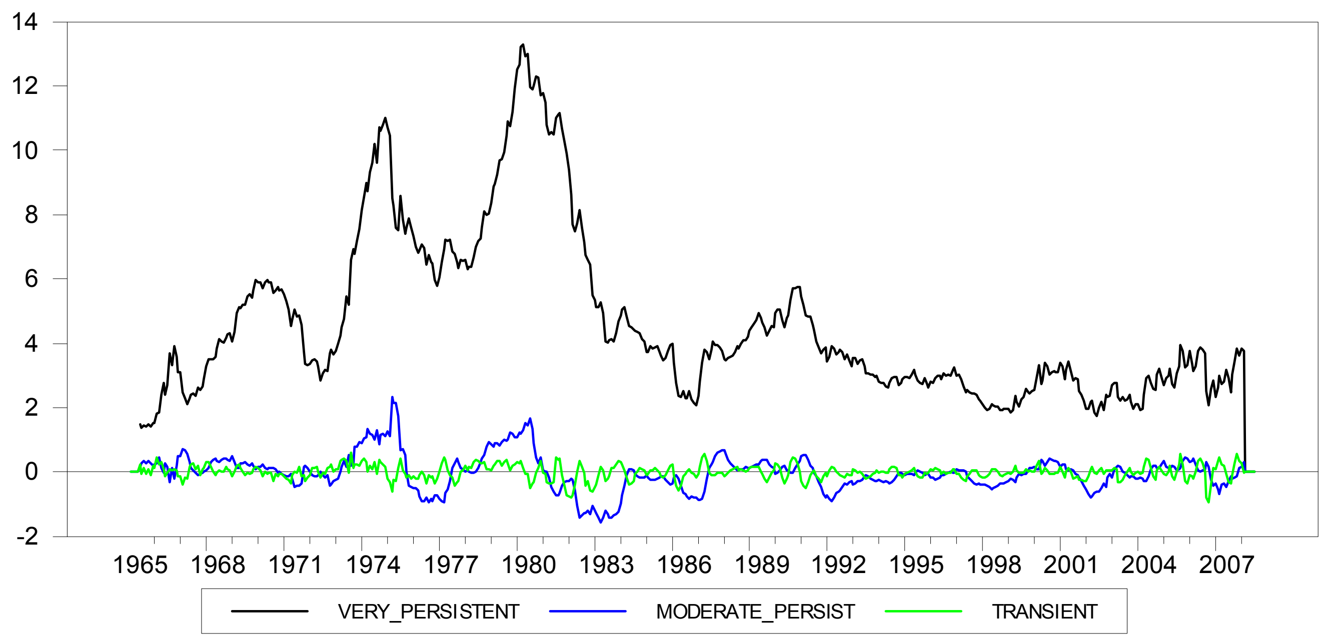

Figure 2: Components of π_t By Persistence

Time plots of "very persistent" components ("persistent_un" or "persistent_pi"), "moderately persistent" components ("moderate_persist"), and "transient" components ("transient_un" or "transient_pi) are displayed in Figure 1 for the u_t components and in Figure 2 for the π_t components.

A few words of explanation may be in order. The very persistent component is not as smooth as one would, from the nomenclature, expect; and there appears to be a modest correlation between the very persistent and the moderately persistent components. The first issue comes about because both the real time data and the projections used in each window can and do shift each period, leading to not-particularly-smooth fluctuations in what otherwise would have been a very smooth "very persistent"component. The second issue arises for similar reasons. Had the entire data sample been subsumed into one long window, these components would be precisely uncorrelated. But in real-time, there is always a time-t innovation that represents a departure from the time $t-1$ forecast of the variable at time t. The movement at time t must be decomposed into various persistence components – and without observations of the variable at times $t+1, t+2$, etc., it is impossible to parse the time-t change into its persistence components without error. A given innovation in real-time u_t or π_t will unavoidably be somewhat misattributed across the components, yielding both a modest level of correlation between the components.[20]

In Section 2.2 the usual simple monetary policy rule model discussed as Equation (1) in Section 1.1 is generalized. First, we take account of persistence in i_t by including lags of i_t in the equation. This generalization is econometrically necessary, so as to yield model errors free from serial correlation; and it is economically interesting, in that such dynamics correspond to the FOMC acting so as to smooth the time path of the federal funds rate. Second, the model is further extended to allow for our estimated FOMC reaction functions to differentiate between highly persistent fluctuations in each variable, moderately persistent fluctuations in each variable, and transient fluctuations in

[20] Most of the resulting inter-component correlations we observe in Figures 1 and 2 are actually quite small, typically in the 0.1 range. The most persistent part of each series has a correlation with the moderately persistent part on the order of 0.2 (u_t) to 0.35 (π_t), however.

each variable.

2.2 Specification and Results

2.2.1 Specification

Our empirical specification was guided by prior work and by theoretical considerations, in conjunction with a desire to allow the data to speak as clearly as possible. First, the federal funds rate is highly persistent. Empirical estimates of central bank policy functions from around the world generally indicate the presence of substantial inertia (see, e.g., Goodhart 1999 or Coibion and Gorodnichenko 2012), consistent with partial adjustment of the interest rate toward its target. (In their study of the U.S. Great Inflation episode in the 1970s, Humpage and Mukherjee (2015) find that inertia was deliberate.) Incorporating such inertia is also found to improve outcomes in many theoretical models (see, e.g., Woodford 2003, Taylor and Williams 2011). In the U.S. at least, there is some evidence for "double inertia "(see Carlstrom and Fuerst, 2014): two lags of the federal funds rate are necessary to yield serially uncorrelated fitting errors. Ignoring such persistence can lead to severe parameter estimate distortions (see Ashley and Verbrugge, 2009a); these distortions are consequently avoided here by the inclusion of lags in i_t.

Second, as discussed in Section 1.1, it is common to see an estimate of a "natural rate of unemployment" and/or an inflation target in a central bank policy reaction function specification. Such inclusions are inherently awkward because these quantities are time-varying (albeit slowly so), but their time variation cannot be identified without making strong (and untestable) assumptions about their time-evolution. Such variables are not included in our base model specification here because the estimation methodology we use in our immediately subsequent model specification – i.e., Equation (3) – will allow for the coefficients on u_t and π_t to differ for slowly-varying ("highly persistent") fluctuations in these variables.

Thus, our base model specification is given by:

$$i_t = \delta_1 i_{t-1} + \delta_2 i_{t-2} + (1 - \delta_1 - \delta_2)(\alpha + \beta u_t + \varphi \pi_t) + e_t. \tag{2}$$

As is standard in the literature, lagged values of π_t and u_t are not included. Because of the parameters δ_1 and δ_2, equation (2) is estimated via non-linear least squares. The term $(\alpha+\beta u_t+\varphi\pi_t)$ is often interpreted as a target interest rate, with the central bank eliminating a fraction $(1-\delta_1-\delta_2)$ of the gap between the current target and the current federal funds rate each month.[21][22]

To investigate whether the FOMC reaction function differentiated between highly persistent movements of each variable, moderately persistent movements of each variable, and higher-frequency movements of each variable, we re-specify Equation (2) to allow for the possibility that the coefficients on π_t and u_t depend on the persistence levels of the fluctuations in these variates. This yields the model:

$$i_t = \delta_1 i_{t-1} + \delta_2 i_{t-2} + (1-\delta_1-\delta_2)\left(\alpha + \sum_{k=1}^{3}\beta_k u_t^k + \sum_{j=1}^{3}\varphi_j \pi_t^j\right) + \epsilon_t. \tag{3}$$

In Equation (3), u_t^1 is the "very persistent" component of u_t; u_t^2 is the "moderately persistent" component of of u_t; and u_t^3 is the "transient" component of of u_t, as described in Section 2.1 above. This model specification allows us to estimate distinct policy-response coefficients (β_1, β_2, and β_3) for these components of unemployment rate fluctuations, as perceived by the FOMC based upon the real-time data available to it at the time a decision was made. The variables $\left(\pi_t^1, \pi_t^2, \text{and } \pi_t^3\right)$ and coefficients (φ_1, φ_2, and φ_3) are analogously defined, corresponding to the "highly persistent,"

[21] Several of the terms in Equation (2) are extremely persistent. While the presence of lagged interest rate terms ensures that all of the coefficients will be estimated consistently, Ashley and Verbrugge (2009a) have shown that a degree of finite-sample bias can be expected with persistent regressors. The sample lengths here – of $T = 176$ in the MBM period and $T = 348$ in the VGB period – are not small, but we note that those simulation results warn that these biases are not necessarily completely negligible when there is substantial regressor persistence, even with several hundred observations. It is worth pointing out, on the other hand, that unit root hypotheses are all strongly rejected in each of the MBM and VGB subperiods when these periods are considered separately. This is in keeping with what one might expect, given the arguments in Clarida, Gali and Gertler (1999) to the effect that stationarity for these variables is implied by theoretical models in which simple monetary policy functions play a role. Also, partitioning the sample to some degree alleviates the problem of time-varying error-term variances mentioned in Sims and Zha (2001, 2006), but we use Eicker-White standard error estimates throughout nonetheless. In a robustness exercise, we also conduct inference using the wild bootstrap.

[22] Implicitly, (2) imposes a fixed equilibrium or natural real interest rate over time. This assumption is almost certainly a poor description of reality over the last decade, but is arguably innocuous up until that time. Prior to the 1980s, estimates often did not reject constancy of the realized real interest rate (see, e.g., Carlson 1977) and such constancy was routinely assumed in academic work during that period (see, e.g., Fama 1975) and even later. During the VGB period, it is possible that the FOMC may have been cognizant of a time-varying natural rate of interest. As robustness checks, during the VGB period we included in (2) the real-time r_t^* estimate of Fuhrer and Olivei (2017) or the one-sided r_t^* estimate of Laubach and Williams (2003, 2015). (Unfortunately, such estimates – regardless of method – "are very imprecise and subject to considerable real time mismeasurement" (LW 2003, Clark and Kozicki 2003).) Our results were qualitatively unchanged.

"moderately persistent," and "transient" components of inflation rate data, as perceived by the FOMC based upon the real-time data available at the time a decision was made.

2.2.2 Results

The first two columns of Table 1 display NLS estimates of δ_1, δ_2, β, and φ separately over the MBM and VGB subperiods, corresponding to Equation (2). As noted above, the inclusion of two lags in i_t in this model suffices to yield serially uncorrelated model fitting errors. Eicker-White standard errors are quoted for all coefficient estimates, here and below, to account, at least asymptotically, for any heteroskedasticity in ϵ_t.[23] The coefficients δ_1 and δ_2 can be taken to quantify the "double-inertia" interest rate smoothing behavior by the FOMC alluded to in Section 2.2.1,, so it is noteworthy that the null hypothesis that both of these coefficients are zero can be always be rejected with $P < 0.0005$.[24]

Our base model specification – Equation (2) – imposes the assumption that the FOMC did not distinguish the most persistent movements in u_t and π_t from other fluctuations. While β and φ enter with conventional signs and statistical significance during the MBM period, during the VGB period the estimates would appear to indicate that the FOMC ignored fluctuations in u_t in this period, since $H_0 : \beta = 0$ cannot be rejected at conventional levels of significance. Taking these coefficient estimates at face value also indicates that the FOMC's response to inflation rate fluctuations was notably smaller in the MBM than in the VGB period: on average the FOMC increased the federal

[23] Possible parameter estimation distortion due to three outlying observations in the fitting errors – for July 1973, May 1980, and Feb 1981 – was addressed using dummy variables to shift the intercept. The estimated coefficients on these dummy variables were always highly significant – and (negative, negative, positive) in signs, respectively – but their exclusion did not substantively affect the inference results reported below. Consequently, the listing of these coefficient estimates – and the model intercept term (α) – is, for simplicity, suppressed in Table 1.

[24] Fed funds rate dynamics were significantly more volatile prior to 1983 than afterwards. What might explain these changes? One potential explanation is regime shifts. First, the Federal Reserve formally adopted monetary targets in 1970. However, by the mid-1970s the fed funds rate had essentially become the operational target. Then, in October 1979, ostensibly to control high inflation, the Federal Reserve switched to targeting monetary aggregates. During this period, there were several episodes during which large fluctuations in the federal funds rate prompted no FOMC response (see Gilbert, 1994). Finally, in October 1982, the Federal Reserve effectively began targeting the fed funds rate again (although arguably this may not have always been the chief focus of attention; for instance, during a conference call on January 9, 1991, Richard Sternlight, who was manager of the trading desk at the Federal Reserve Bank of New York, remarked, "Somewhat by default we might place more guidance on, of all things, the federal funds rate..."). Regime switches often lead to uncertainty, experimentation, and apparent policy reversals, and the associated diminishment of any beneficial effects arising from stable expectations.

funds rate by only 0.63% for every 1% increase in the inflation rate in MGM period, whereas in the VGB period the estimated response is 1.82%. The FOMC's short-run response to a 1% increase in the unemployment rate is only economically meaningful during the MBM period, and even then it is of only modest economic significance (-0.84%).[25]

But these base model results are artifactual. Allowing for (but not imposing an assumption that) the FOMC could distinguish the most persistent movements in u_t and π_t from other fluctuations yields interestingly different results. Indeed, the imposition that the FOMC did not distinguish fluctuations by persistence level leads to seriously misleading conclusions as to the FOMC's past behavior. In particular, we turn now to the estimation results based on Equation (3), which explicitly allows for "persistence dependence" in the FOMC response to unemployment and inflation rate fluctuations.

[25] If one re-estimates Equation (2) using a conventional unemployment rate gap, one finds – in keeping with most of the literature – evidence for a response to unemployment rate fluctuations in the VGB period. Evidently, the apparent non-response to unemployment rate fluctuations in the VGB period indicated in Table 1 for Equation (2) arises from the fact that the FOMC does not respond to extremely persistent movements in the unemployment rate, and these comprise a large part of the variance.

Table 1: Estimation Results

	Equation 2			Equation 3	
	MBM Period	VGB Period		MBM Period	VGB Period
δ_1	$+1.46^{***}$ (0.09)	$+1.37^{***}$ (0.08)	δ_1	$+1.33^{***}$ (0.09)	$+1.28^{***}$ (0.08)
δ_2	-0.55^{***} (0.09)	-0.42^{***} (0.08)	δ_2	-0.42^{***} (0.09)	-0.37^{***} (0.07)
			β_1	-0.46 (0.32)	$+0.38$ (0.35)
β (u_t)	-0.84^{***} (0.26)	$+0.08$ (0.36)	β_2	-1.57^{***} (0.50)	-2.63^{***} (0.51)
			β_3	-6.58^{***} (2.67)	-5.14^{***} (2.05)
			φ_1	$+0.60^{***}$ (0.17)	$+1.63^{***}$ (0.30)
φ (π_t)	$+0.63^{***}$ (0.16)	$+1.82^{***}$ (0.48)	φ_2	$+0.09$ (0.60)	-0.28 (1.11)
			φ_3	$+1.16$ (1.82)	-5.05^{***} (1.80)
			F-test: $\beta_i = 0\ \forall i$	0.004	0.000
			F-test: $\beta_i = \beta_j\ \forall i, j$	0.023	0.000
			F-test: $\beta_2 = \beta_3$	0.76	0.24
			F-test: $\varphi_i = 0\ \forall i$	0.001	0.000
			F-test: $\varphi_i = \varphi_j\ \forall i, j$	0.66	0.013
			F-test: $\varphi_2 = \varphi_3$	0.81	0.016

With respect to the persistence-decomposed u_t fluctuations, we note that the null hypothesis $\beta_1 = \beta_2 = \beta_3$ can be rejected in both periods; thus the Equation (2) base-model specification artificially imposing equality of these coefficients is not supported by the data, so that inferences based on that model are likely distorted. The source of that rejection is clear: β_1, the coefficient on the very persistent component of u_t is statistically insignificant in both periods. Thus, in contrast to the distorted results from the base model, we can conclude that during the MBM period, *just like in the VGB period*, the FOMC did not respond to – i.e., was ignoring – the most persistent unemployment rate fluctuations. Now turn to the β_2 estimates, i.e. the coefficients on the "moderately persistent" components of u_t for the two periods. We note that during the VGB period, *even more than in the MBM period*, the FOMC aggressively responded to those unemployment rate fluctuations that are, perhaps, the most responsive to policy changes. Thus, we find that the FOMC has *in both periods* ignored very persistent movements in u_t, and has significantly focused on responding to less-persistent fluctuations. Our method allows us to determine this without making use of ancillary assumptions, or imposing any concepts that were arguably not guiding FOMC discussions during this period.

This continuity of behavior that we find vis-a-vis fluctuations in the unemployment rate is evidence against the "views" explanation of the Great Inflation. Arguably, u_t^* conceptual or measurement struggles do *not* explain FOMC behavior in the 1970s, as some analysts have asserted. Instead, as noted above, the FOMC has *always* ignored very persistent fluctuations in the unemployment rate, a response that stands in sharp contrast to its strong response to moderately persistent fluctuations of the unemployment rate. This ignoring-of-persistent-movements-in-u_t is consistent with its viewing such movements as being natural rate fluctuations; but this interpretation is not necessary: one might alternatively interpret this behavior as treating the persistent part as an "intermediate target." Our results thus suggest that much previous work (which has often focused on particular estimates of an output gap, rather than on an unemployment gap) may have come to erroneous conclusions about FOMC behavior in the 1970s. We further note that, as is evident from the plots in Figure 1, the FOMC's implicit intermediate target for u_t is much more volatile that conventional "natural rate" estimates.

With respect to inflation rate (π_t) fluctuations, we find that – in contrast to its continuity of response to the unemployment rate over both periods – the FOMC response to inflation movements changed in not just one, but two ways, starting in the VGB period. First, comparing the φ_1, φ_2, and φ_3 estimates for the two periods, the VGB FOMC became notably more aggressive in its response to inflation – something other authors have noted (e.g., Clarida, Gali and Gertler 2000 or Conrad and Eife 2012). Indeed, in contrast to many studies,[26] we find that – if one takes the coefficient estimates $\widehat{\varphi}_1$ across the different periods at face value, so that the comparison is

[26] See, e.g., Coibion and Gorodnichenko (2011) and discussion therein.

sensible – we can reject the hypothesis that these coefficients are equal across the two periods (Z-score = 2.99).[27] And second, in the VGB period the FOMC also became much more focused on fighting the *persistent* fluctuations in inflation, and ignoring more transient movements.[28] While the φ_k coefficient estimates themselves are consistent with a differential-inflation-persistence-response interpretation in the MBM period, the formal statistical test of coefficient equality fails to reject the null hypothesis that the FOMC responded to all inflation fluctuations in the same way during this period. We interpret this as indicating that the FOMC during the VGB period became either more focused or more accurate at identifying persistent inflation fluctuations.

2.2.3 Comparisons and Robustness

As a robustness check, and because of the potential for distortion in least-squares inference with regard to coefficients on highly persistent regressors in samples of modest length – as noted in Ashley and Verbrugge (2009a) – we also used wild bootstrap simulation to test the various null hypotheses indicated in Table 1. (The wild bootstrap was used in view of strong evidence for heteroscedasticity in the model errors.) With one exception, the resulting inferences were essentially unchanged: the standard error estimates and inference p-values for the hypothesis tests results given in Table 1 were generally quite similar to those obtained using the NLS asymptotic heteroscedasticity-robust standard errors. The only exception to this is the rejection p-value for the test of the joint null hypothesis that the coefficients on all three u_t components were equal in the MBM period, where the wild bootstrap simulations resulted in a rejection p-value of 0.09, whereas the p-value resulting from the NLS asymptotics yields a rejection p-value of 0.023. Each of these p-values is, itself, just an (asymptotically-justified) sample estimate. Consequently we take the inference here to be rejection

[27] To further test this result, we computed the statistic $(\Delta\varphi_1)^2 + (\Delta\varphi_2)^2 + (\Delta\varphi_3)^2$, where $\Delta\varphi$ refers to the difference between the coefficient estimates across the two periods, and performed a wild bootstrap to generate p-values for this statistic. The resulting p-value is 0.056. Conversely, performing the same exercise using β coefficients yields a p-value in excess of 0.70.

[28] On its face, φ_3 is estimated to be significantly negative during the VGB period, ostensibly indicating high-frequency accomodation of inflation rate fluctuations. However, this result is actually an artifact, due to this coefficient being unstable during the VGB period: Between October 1979 and October 1982, rather than targeting the fed funds rate, the FOMC was targeting nonborrowed reserves (as a means to control monetary aggregates, and hence inflation). During this period, the fed funds rate was highly variable, with reversals that roughly line up with high-frequency inflation shifts. This is consistent with the view that high-frequency shifts in inflation occurred along with (and perhaps causing) high-frequency shifts in money demand that, in turn, induced high-frequency shifts in the fed funds rate. When one allows a break in φ_3 after 1982, this coefficient is statistically different from zero only during the pre-1983 period.

at the 2% to 10% level of significance. This significance level is not as compelling as we would like; however, we note that this result is likely an artifact of the nonlinear least squares estimation technique that Equations (3) and (4) require due to the gap expression. In particular, when we instead reformulate these models simply using two lagged dependent variables in the specification and estimate by OLS (with HAC standard errors), the null hypothesis $\beta_i = \beta_j \; \forall i, j$ is resoundingly rejected at the 0.00001 level of significance.

One might ask, is our Equation (3) model for the VGB period significantly better than a conventional alternative? In particular, how do its predictions compare to those of an alternative model which uses a more conventional estimate of u_t^*, such as that of the CBO, and ignores the distinction betweeen persistent and transient movements in inflation? To answer this question, we used the Diebold-Mariano forecast comparison test to compare one-step-ahead forecast errors from two linear models over the VGB period, estimated over rolling ten-year windows.[29] To make the treatment of u_t comparable across models, we specified both models in "gap" form: in the CBO model, the unemployment gap was defined to be $\left(u_t - u_t^{*,CBO}\right)$, while in our model, the unemployment gap was defined to be $\left(u_t - u_t^{persistent}\right)$. In both models, this gap term was lagged one month. Regarding the treatment of π_t, the CBO model simply uses the real-time estimate of π_t, while our model uses the real-time estimate of $\pi_t^{persistent}$. In both models, this term was lagged two months. Then each model included a constant and one lag of the fed funds rate. Our model significantly outperformed the CBO model: the p-value of Diebold-Mariano test statistic was $<$ 0.0005.

As we have noted above, our results are robust to a number of specification choices. One in particular may be noteworthy. It may be argued that during the VGB period, the FOMC was implicitly using a time-varying target natural rate of interest (r_t^*) concept. If we modify (3) to incorporate an r^* in the parenthetical expression, and then use the real-time r_t^* estimate from Fuhrer and Olivei (2017),[30] or the one-sided r_t^* estimate from Laubach and Williams (2015), our results are nearly unchanged.

[29] This is not a true real-time forecast comparison, since the CBO estimate is taken from the future. One might argue, however, that this biases the comparison against our Equation (3) model.

[30] We thank these authors for sharing these data with us.

2.2.4 Regarding the Great Inflation

Our results also shed some light on the origins of the Great Inflation. They imply that the "ideas" hypothesis – that FOMC errors in the 1970s were due to erroneous beliefs about the structure of the economy, including an unrealistically low estimate of the natural rate of unemployment (e.g., Orphanides 2002, Romer and Romer 2002, Romer 2005) – cannot fully explain FOMC policy during that time period. Our findings point in this direction because we show that FOMC behavior with respect to the fed funds rate response to the unemployment rate has remained rather stable: it has always ignored the most persistent fluctuations in u_t. So what can explain the high inflation of that period? On the face of it, our results suggest that a change in policy preferences – possibly driven by political pressures (see, e.g., Chappell, Havrilesky and McGregor (1993), Meltzer (2009), Weise (2012), or Levin and Taylor (2013)) – is a far more likely explanation.[31]

What about the conjecture that FOMC behavior in the 1970s resulted from inaccurate inflation forecasts (see, e.g., Orphanides 2002, Levin and Taylor 2013 or Fuhrer and Olivei 2017)? This mechanism could certainly operate in conjunction. To investigate this, in parallel to our treatment of the real-time monthly CPI inflation rate, we constructed monthly estimates of the highly persistent part of the quarterly GNP deflator. For this exercise, one-sided filtering was conducted for each month using a five-year window of quarterly data based upon 15 quarters of real-time data, augmented with the monthly real-time nowcasts and 4 quarterly forecasts in the Greenbook of that month. (Early in the sample, in order to obtain sufficient number of forecasts, we had to augment Greenbook forecasts. We did this using SPF forecasts when available, or using simple univariate forecasts which treated available Greenbook forecasts (or Greenbook + SPF forecasts) as additional data.) We find that, as one might have expected based upon previous research, the implied estimate of $\pi_t^{persistent}$ is systematically significantly lower than our baseline CPI-based estimate during a large portion of the 1970s. However, we conducted a second out-of-sample forecast comparison test to determine which model – our baseline CPI-based model, or the alternative GNP-deflator-with-Greenbook-forecasts – better predicted or described actual FOMC behavior. As before, we

[31] Many commentators (e.g., Sargent 1999, Hetzel 2018) note that the narrative evidence indicates that the FOMC largely held pessimistic beliefs about the "sacrifice ratio," i.e. that reducing inflation would require very high levels of unemployment. But as Levin and Taylor (2013) argue, such a belief should translate into a more *vigorous* response to incipient increases in inflation.

used the Diebold-Mariano forecast comparison test to compare two linear models over the MBM period, estimated over rolling ten-year windows. Both models treat u_t identically, with this variable entering in gap form as in $\left(u_t - u_t^{persistent}\right)$; in both models, this term was lagged one month. Regarding the treatment of π_t, the GNP deflator model uses the *current* real-time estimate of $\pi_t^{persistent}$, while our model uses our real-time estimate of $\pi_t^{persistent}$, lagged two months. Then each model included a constant and one lag of the fed funds rate. Despite our model's informational disadvantage – the GNP deflator model always uses all available inflation information, including CPI inflation and other inflation-relevant data, while our baseline model ignores other inflation-relevant data and is sometimes one month out of date relative to the GNP deflator data – our baseline model significantly outperformed the GNP deflator model: the p-value of Diebold/Mariano test statistic was 0.037. This indicates that the actual behavior of the FOMC in the 1970s is more consistent with its using real-time CPI inflation data (where identification of the persistent part of inflation fluctuations uses information from simple-but-reliable inflation forecasts). Evidently, poor inflation forecasts do not explain the Great Inflation.[32]

3 Conclusions

Using the lens of simple monetary policy reaction functions, we apply recently developed econometric tools to deepen our understanding of FOMC behavior and how it has (or has not) changed starting with the chairmanship of Volcker. Standard simple monetary policy rules like Equation (2) properly allow for persistence in the fed funds rate but arbitrarily impose an assumption that the FOMC treated persistent fluctuations in the unemployment rate, and persistent fluctuations in the inflation rate, in the same way that they treated transient movements in those variables. In this paper, we relax that restriction and test it. Our results are surprising along some dimensions, but accord well with intuition and some other accounts. Our study reaches the following conclusions.

[32] We have conducted several different investigations into the post-sample forecasting effectiveness of our models, comparing the ability of models incorporating frequency-partitioned explanatory variables relative to analogous models specified without frequency dependence; two of these were reported above. These results were encouraging, in that a disaggregated model is always somewhat superior on this metric, yielding credence to our modeling approach. We note, however, that the present study is about inference rather than forecasting, and we lay no claim here to having developed a new forecasting approach. For a study along those lines, see, e.g., Carlstrom and Zaman (2014).

First, estimates of Equation (2) as it stands imply that the FOMC was unresponsive to unemployment rate fluctuations in the VGB period; as we explain below, this fundamentally misleading result underscores the necessity of allowing for a distinction between the response to a persistent movement in the unemployment rate versus the response to a less persistent fluctuation. By extension, we would argue that there are likely many other macroeconomic relationships in which the relationship between two variables differs by persistence level. In such cases, restricting this relationship to be the same across differing persistence levels will quite likely to miss out on uncovering some interesting features in the data, and may well lead to over-simplified and distorted inferences.

Second, by allowing the data to inform us about FOMC responses to persistent fluctuations versus higher-frequency fluctuations, we notably find that mismeasurement of u_t^* (or the output gap) is *not* a convincing explanation of FOMC behavior in the 1970s, as some analysts have argued. Instead, whether or not such an equilibrium concept was understood formally, or simply incorporated in judgment (whether consciously or unconsciously), the FOMC in *both* periods ignored extremely persistent fluctuations in the unemployment rate (although it arguably got better at this starting with Volcker). We follow Fuhrer and term this extremely persistent value the FOMC's "intermediate target" for the unemployment rate. This lack of response to extremely persistent fluctuations stands in sharp contrast to the FOMC's strong response to *moderately* persistent fluctuations of the unemployment rate. (Clearly, it is "business-cycle" fluctuations in u_t that are generally thought to be responsive to movements in the federal funds rate.) Our finding of this continuity of FOMC behavior between the MBM and VGB periods suggests that much previous work, which has incorporated particular estimates of an output gap (rather than an appropriately modeled unemployment gap), may have come to erroneous conclusions about FOMC behavior in the 1970s.

Third, the FOMC response to inflation movements changed in not just one, but two ways, starting in the VGB period. First, the VGB FOMC became more aggressive in its response to inflation – something other authors have noted (e.g., Clarida, Gali and Gertler 2000). But second, it also became much more focused on fighting the most *persistent* fluctuations in inflation, and began ignoring more transient movements.

Our results thus shed some light on the origins of the Great Inflation. They imply that the "ideas" hypothesis – that FOMC errors in the 1970s were due to erroneous beliefs about the structure of the economy, particularly conceptual or measurement errors regarding the natural rate of unemployment (e.g., Orphanides 2002, Romer and Romer 2002, Romer 2005) – cannot fully explain FOMC policy during that time period. They further imply that the "poor inflation forecasts" hypothesis also fails to explain FOMC behavior. Hence, taken together, our results suggest that a change in policy preferences – possibly driven by political pressures – is a far more likely explanation of the Great Inflation.

Economic theory often suggests that decision-makers distinguish between fluctuations with different degrees of persistence. The relatively straightforward technique used here illustrates how one can easily test this hypothesis – even in settings where (as here) feedback is likely, so that ordinary spectral analysis is inappropriate. This methodological innovation can lead to sharp new insights about the process generating such data, whereas ordinary time-series techniques (in either the time domain or the frequency domain) would both fail to reliably uncover features of this nature in the data generating process, and yield distorted inference results for having glossed over them.

Notably, the "persistence dependence" in empirical model coefficients directly addressed by the new econometric methodology used here is very clearly interpretable in intuitive economic terms. This methodology is therefore particularly well suited toward guiding the construction of deeper and richer structural model specifications for the economic processes underlying the data generating mechanism whose properties have been thus unveiled.

References

[1] **Aastveit, Knut Are, Francesco Furlanetto, and Francesca Loria** (2017) "Has the Fed Responded to House and Stock Prices? A Time-Varying Analysis." Norges Bank Research 1/2017.

[2] **Almon, Shirley** (1965): "The Distributed Lag between Capital Appropriations and Expenditures," *Econometrica*, 30, 178-196.

[3] **Ashley, Richard and Kwok Ping Tsang** (2013): "The Oil Price-Macroeconomy Relationship: Does Persistence Matter? " working paper.

[4] **Ashley, Richard and Guo Li** (2014): "Re-Examining the Impact of Housing Wealth and Stock Wealth on Retail Sales: Does Persistence in Wealth Changes Matter?," *Journal of Housing Economics*, 26, 109-118.

[5] **Ashley, Richard and Randal J. Verbrugge** (2009a): "To Difference or Not to Difference: A Monte Carlo Investigation of Inference in Vector Autoregression Models," *International Journal of Data Analysis Techniques and Strategies* Vol. 1, No. 3, 242-274.

[6] **Ashley, Richard and Randal J. Verbrugge** (2009b): "Frequency Dependence in Regression Model Coefficients: An Alternative Approach for Modeling Nonlinear Dynamic Relationships in Time Series,"*Econometric Reviews*, 28(1-3), 4-20.

[7] **Ashley, Richard and Randal J. Verbrugge** (2018): "The Phillips Curve Includes Two Gaps, Not One" (formerly titled, "Mis-Specification in Phillips Curve Regressions: Quantifying Frequency Dependence in This Relationship While Allowing for Feedback") Working paper.

[8] **Basu, Susanto, Fernald, John G. and Miles S. Kimball** (2006) "Are Technology Improvements Contractionary?," *American Economic Review*, Vol. 96, No. 5, 1418-1448.

[9] **Bernanke, Ben S. and Jean Boivin** (2003) "Monetary policy in a data-rich environment," *Journal of Monetary Economics*, vol. 50, no. 3, p. 525-546.

[10] **Boivin, Jean** (2006) "Has U.S. Monetary Policy Changed? Evidence from Drifting Coefficients and Real-Time Data," *Journal of Money, Credit and Banking*, vol. 38, no. 5, p. 1149-1173.

[11] **Box, George E.P. and Gwilym Jenkins** (1976): *Time Series Analysis, Forecasting and Control*, Holden-Day, San Francisco.

[12] **Brischetto, Andrea, and Anthony Richards** (2007), "The Performance of Trimmed Mean Measures of Underlying Inflation," paper presented at the Conference on Price Measurement for Monetary Policy, May 24-25, 2007, Jointly sponsored by the Federal Reserve Banks of Cleveland and Dallas.

[13] **Bryan, Michael F. and Brent Meyer** (2002): "Are Some Prices in the CPI More Forward Looking Than Others? We Think So", *Economic Commentry*, No. 2010-2, Federal Reserve Bank of Cleveland.

[14] **Carlstrom, Charles T. and Timothy S. Fuerst** (2014): "Adding Double Inertia to Taylor Rules to Improve Accuracy" *Federal Reserve Bank of Cleveland Economic Commentary 2014-08*.

[15] **Carlstrom, Charles T. and Timothy Stehulak** (2015): "The Long-Run Natural Rate of Interest" *Federal Reserve Bank of Cleveland Economic Trends*.

[16] **Cecchetti, Stephen G.** (1997): "Measuring Short-Run Inflation For Central Bankers", *FRB Saint Louis - Review*, , vol. 79, no. 3, p. 143-155.

[17] **Chappell, Henry W. Jr, Thomas M. Havrilesky, and Rob Roy McGregor** (1993): Partisan Monetary Policies: Presidential Influence Through the Power of Appointment. Quarterly Journal of Economics 108: 185-218.

[18] **Chappell, Henry W. Jr, Rob Roy McGregor, and Todd Vermilyea** (2004): Majority Rule, Consensus Building, and the Power of the Chairman: Arthur Burns and the FOMC. Journal of Money, Credit and Banking, Vol. 36, No. 3, Part 1 (Jun., 2004), pp. 407-422

[19] **Chen, Xiaoshan, Tatiana Kirsanova, and Campbell Leith** (2017) "How Optimal is US Monetary Policy?" Journal of Monetary Economics 92, 96-111.

[20] **Christiano, Lawrence J. and Terry J. Fitzgerald**(2003): "The Band Pass Filter," *International Economic Review*, vol. 44, no. 2, p. 435-465.

[21] **Clarida, Richard, Jordi Gali and Mark Gertler** (2000): "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics*, vol. 115, no. 1, p. 147-180.

[22] **Cochrane, John H.** (1989): "The Return of the Liquidity Effect: A Study of the Short-Run Relation between Money Growth and Interest Rates", *Journal of Business and Economic Statistics*, Volume 7, 75-83.

[23] **Cogley, Timothy, Riccardo Colacito and Thomas J. Sargent** (2007): "Benefits from U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas," *Journal of Money, Credit and Banking*, Vol.39, No.1, 67-99.

[24] **Coibion, Olivier, and Yuriy Gorodnichenko** 2012: "Why Are Target Interest Rate Changes So Persistent?" *American Economic Journal: Macroeconomics*, 4(4): 126-62.

[25] **Conrad, Christian, and Thomas A. Eife** (2012) "Explaining Inflation-Gap Persistence by a Time-Varying Taylor Rule." *Journal of Macroeconomics*, 34, 419-428.

[26] **Consolo, Agostino and Carlo A. Favero** (2009): "Monetary Policy Inertia: More a Fiction than a Fact?" *Journal of Monetary Economics*, Volume 56, Issue 6, 900-906.

[27] **Curdia, Vasco, Andrea Ferrero, Ging Cee Ng, and Andrea Tambalotti** (2015): "Has U.S. Monetary Policy Tracked the Efficient Interest Rate?" *Journal of Monetary Economics*, 70, 72-83.

[28] **Dagum, Estela** (1978): "Modelling, Forecasting and Seasonally Adjusting Economic Time Series with the X11 ARIMA Method," *The Statistician*, 27, 3/4, 203-216.

[29] **Debortoli, Davide, and Aeimit Lakdawala** (2016) "How Credible is the Federal Reserve? A Structural Estimation of Policy Re-Optimizations." American Economic Journal: Macroeconomics, 8.3, 42-76.

[30] **Detmeister, Alan, and Nathan Babb** (2017): "Nonlinearities in the Phillips Curve for the United States," Federal Reserve Bank of Cleveland, Working Paper.

[31] **Dolmas, Jim and Mark A.Wynne** (2008): "Measuring Core Inflation: Notes from a 2007 Dallas Fed Conference," Staff Papers No. 4, May 2008, Federal Reserve Bank of Dallas.

[32] **Engle, Robert** (1974): "Band Spectrum Regression," *International Economic Review*, 15, 1-11.

[33] **Espinosa-Vega, Marco A., and Steven Russell** (1997): History and Theory of the NAIRU: A Critical Review. *Federal Reserve Bank of Atlanta Economic Review* Second Quarter 1997, 4- 25.

[34] **Friedman, Milton** (1968): "The Role of Monetary Policy". *American Economic Review* 58: 1-17.

[35] **Fuhrer, Jeffrey, and Geoffrey Tootell** (2008): "Eyes on the Prize: How Did the Fed Respond to the Stock Market?" *Journal of Monetary Economics*, Volume 55, 796-805.

[36] **Fuhrer, Jeffrey, and Giovanni Olivei** (2017): "Rules and discretion: An Empirical Assessment," Working paper, Federal Reserve Bank of Boston. Prepared for the Federal Reserve Bank of Boston's 61^{st} Annual Economic Conference, "Are Rules Made to be Broken? Discretion and Monetary Policy".

[37] **Gali, Jordi** (2009): *Monetary Policy, Inflation, and the Business Cycle*, Princeton University Press.

[38] **Giannone, Domenico, and Troy D. Matheson** (2007): "A New Core Inflation Indicator for New Zealand," *International Journal of Central Banking*, December, p. 145-180.

[39] **Gilbert, R. Anton** (1994): "A Case Study in Monetary Control: 1980-82" *Federal Reserve of Saint Louis Review*, Vol 76, No. 6, 35-55.

[40] **Goodfriend, Marvin, and Robert G. King** (2013) "The Great Inflation Drift." in Michael D Bordo and Athanasios Orphanides, Eds., *The Great Inflation: The Rebirth of Modern Central Banking.* Chicago: University of Chicago Press.

[41] **Goodhart, Charles A.** (1999): "Central Bankers and Uncertainty" Keynes Lecture in Economics *Proceedings of the British Academy*, 101: 229-71.

[42] **Gordon, Robert J.** (1975): "Alternative Responses of Policy to External Supply Shocks," *Brookings Papers on Economic Activity*, 1975(1), p. 183-206.

[43] **Granger, Clive W.J.** (1969): "Investigating Causal Relations by Econometric Models and Cross-Spectral Methods," *Econometrica*, 37, 424-438.

[44] **Hamilton, James D., Ethan S. Harris, Jan Hatzius, and Kenneth D. West.** (2015): "The Equilibrium Real Funds Rate: Past, Present, and Future" (Presented at the US Monetary Policy Forum, New York, February 27, 2015.)

[45] **Hannan, E.** (1963): "Regression for Time Series," in M. Rosenblatt, ed. *Time Series Analysis*, John Wiley: New York, 14.37.

[46] **Hetzel, Robert L.** (2018) "The Evolution of U.S. Monetary Policy." WP 18-01, Federal Reserve Bank of Richmond.

[47] **Higgins, Amy, and Randal Verbrugge** (2015): "Tracking Trend Inflation: Nonseasonally Adjusted Variants of the Median and Trimmed-Mean CPI," Federal Reserve Bank of Cleveland, Working Paper p. 15-27.

[48] **Holston, Kathryn, Thomas Laubach, John C. Williams.** (2016): "Measuring the Natural Rate of Interest: International Trends and Determinants." Federal Reserve Bank of San Francisco Working Paper 2016-11. http://www.frbsf.org/economic-research/publications/workingpapers/ wp2016-11.pdf

[49] **Humpage, Owen F., and Sanchita Mukherjee** (2015): "Even Keel and the Great Inflation" Federal Reserve Bank of Cleveland Working Paper 15-32.

[50] **Johnston, J.** (1972) *Econometric Methods*, McGraw-Hill: New York.

[51] **Jondeau, Eric, Herve Le Bihan and Clementine Galles** (2004): "Assessing Generalized Method-of-Moments Estimates of the Federal Reserve Reaction Function," *Journal of Business and Economics Statistics*, 22(2), 225-239.

[52] **Kim, Chang-Jin and Charles R. Nelson** (2006): "Estimation of a Forward-Looking Monetary Policy Rule: A Time-Varying Parameter Model Using Ex Post Data," *Journal of Monetary Economics*, 53, p. 1949-1966.

[53] **Knotek, Edward S. II, Randal Verbrugge, Christian Garciga, Caitlin Treanor, and Saeed Zaman** (2016): "Federal Funds Rates Based on Seven Simple Monetary Policy Rules." *Federal Reserve Bank of Cleveland Economic Commentary 2016-17.*

[54] **Koopmans, Lambert H.** (1974): *The Spectral Analysis of Time Series*, San Diego, CA: Academic Press.

[55] **Laubach, Thomas, and John C. Williams.** (2003): "Measuring the Natural Rate of Interest" *Review of Economics and Statistics* 85(4), 1063-1070.

[56] **Laubach, Thomas, and John C. Williams.** (2015): "Measuring the Natural Rate of Interest Redux" Federal Reserve Bank of San Francisco Working Papers, 2015-16.

[57] **Levin, Andrew, and John B. Taylor** (2013) "Falling Behind the Curve: A Positive Analysis of the Stop-Start Monetary Policies and the Great Inflation." in Michael D Bordo and Athanasios Orphanides, Eds., *The Great Inflation: The Rebirth of Modern Central Banking.* Chicago: University of Chicago Press.

[58] **McConnell, Margaret M. and Gabriel Perez-Quiros** (2000): "Output Fluctuations in the United States: What Has Changed Since the Early 1980's?" *American Economic Review*, Vol. 90, No. 5, pp. 1464-1476.

[59] **McCulloch, J. Huston** (2007): "Adaptive Least Squares Estimation of the Time-Varying Taylor Rule," working paper.

[60] **Meltzer, Allan H.** (2005): "Origins of the Great Inflation," *Federal Reserve Bank of St. Louis Review*, March/April Part 2, p. 145-176.

[61] **Meyer, Brent, Guhan Venkatu and Saeed Zaman** (2013): "Forecasting Inflation? Target the Middle," *Economic Commentary*, 2013-05, pp. 1-4.

[62] **Mishkin, Frederic S.** (2007): "Inflation Dynamics" in *International Finance*, vol. 10(3), Blackwell Publishing: Oxford, p. 317-334.

[63] **Modigliani, Franco, and Lucas Papademos** (1975): "Targets for Monetary Policy in the Coming Year." Brookings Papers on Economic Activity 1, 141-63.

[64] **Newey, Whitney K. and Kenneth D. West** (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, Vol. 55, No. 3, p. 703-708.

[65] **Orphanides, Athanasios** (2001): "Monetary Policy Rules Based on Real-Time Data," *American Economic Review*, vol. 91, no. 4, p. 964-985.

[66] **Orphanides, Athanasios** (2002): "Monetary-Policy Rules and the Great Inflation," *American Economic Review: Papers and Proceedings*, vol. 92, no. 2, p. 115-120.

[67] **Orphanides, Athanasios** (2008): "Taylor Rules,"in *The New Palgrave Dictionary of Economics*. Second Edition. Steven N. Durlauf and Lawrence E. Blume (eds.) Palgrave Macmillan, 2017. *The New Palgrave Dictionary of Economics Online*. Palgrave Macmillan. (doi: 10.1057/9780230226203.1686)

[68] **Phelps, Edmund S.** (1968). "Money-Wage Dynamics and Labor-Market Equilibrium". *Journal of Political Economy* 76: 678–711.

[69] **Roskelly, Kenneth D.** (2016) "Augmenting the Taylor Rule: Monetary Policy and the Bond Market." Economics Letters 144, 64-67.

[70] **Rudebusch, Glenn D.** (2002): "Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia," *Journal of Monetary Economics*, Volume 49, Issue 6, 1161-1187.

[71] **Sims, Christopher A. and Tao Zha** (2001): "Stability and Instability in US Monetary Policy Behavior," working paper.

[72] **Sims, Christopher A. and Tao Zha** (2006): "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review*, vol. 96, no. 1, p. 54-81.

[73] **Stock, James and Mark Watson** (1999): "Business Cycle Fluctuations in U.S. Macroeconomic Time Series," in (John Taylor and Michael Woodford, eds) *Handbook of Macroeconomics*, Amsterdam: Elsevier, 3-64.

[74] **Svensson, Lars E. O.** (2013) "Comment." in Michael D Bordo and Athanasios Orphanides, Eds., *The Great Inflation: The Rebirth of Modern Central Banking.* Chicago: University of Chicago Press.

[75] **Tan, H. B. and Richard Ashley** (1999a): "On the Inherent Nonlinearity of Frequency Dependent Time Series Relationships," in (P. Rothman, ed.) *Nonlinear Time Series Analysis of Economic and Financial Data*, Kluwer Academic Publishers: Norwell, 129-142.

[76] **Tan, H.B. and Richard Ashley** (1999b): "Detection and Modeling of Regression Parameter Variation Across Frequencies with an Application to Testing the Permanent Income Hypothesis," *Macroeconomic Dynamics*, 3, 69-83.

[77] **Tasci, Murat and Randal Verbrugge** (2014): "How Much Slack is in the Labor Market? That Depends on What You Mean by Slack," *Economic Commentary*, 2014-21, pp. 1-6.

[78] **Taylor, John B.** (1993): "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 93, p. 195-214.

[79] **Taylor, John B. and John C. Williams** (2011): "Simple and Robust Rules for Monetary Policy" in Benjamin M. Friedman and Michael Woodford, eds., *Handbook of Monetary Economics* vol. 3B, North Holland, 829-60.

[80] **Weise, Charles L.** (2012): "Political Pressures on Monetary Policy During the U.S. Great Inflation," American Economic Journal: Macroeconomics, Vol. 4, No. 2, 33-64.

[81] **Woodford, Michael** (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

4 Technical Appendix

4.1 Modeling Frequency Dependence

In this section we discuss the technique used here for modeling frequency dependence in the monetary policy rule, Equation (2) above.[33]

The idea of regression in the frequency domain can be traced back to Hannan (1963) and Engle (1974, 1978), and is further developed in Tan and Ashley (1999a and 1999b), who developed a real-valued reformulation of Engle's (1974) complex-valued framework.

Consider the ordinary regression model:

$$Y = X\beta + e \quad e \sim N(0, \sigma^2 I) \tag{4}$$

where Y and e are each $T \times 1$ and X is $T \times K$. Now define a $T \times T$ matrix A, whose $(s,t)^{th}$ element is given by:

$$a_{s,t} = \begin{array}{l} (\frac{1}{T})^{\frac{1}{2}} \quad \text{for } s = 1; \\ (\frac{2}{T})^{\frac{1}{2}} \cos(\frac{\pi s(t-1)}{T}) \quad \text{for } s = 2, 4, 6, ..., (T-2) \text{ or } (T-1); \\ (\frac{2}{T})^{\frac{1}{2}} \sin(\frac{\pi(s-1)(t-1)}{T}) \quad \text{for } s = 3, 5, 7, ..., (T-1) \text{ or } T; \\ (\frac{1}{T})^{\frac{1}{2}} (-1)^{t+1} \quad \text{for } s = T \text{ when } T \text{ is even.} \end{array} \tag{5}$$

It can be shown that A is an orthonormal matrix, so its transpose is its inverse and Ae is still distributed $N(0, \sigma^2 I)$. Pre-multiplying the regression model (4) by A thus yields,

$$AY = AX\beta + Ae \rightarrow Y^* = X^*\beta + e^*, e^* \sim N(0, \sigma^2 I) \tag{6}$$

where Y^* is defined as AY, X^* is defined as AX, and e^* is defined as Ae. The dimensions of the of Y^*, X^*, and e^* arrays are the same as those of Y, X, and e in Equation (4), but the T components of Y^* and e^* and the rows of X^* now correspond to frequencies instead of time periods.

[33] See Ashley and Verbrugge (2009b) for details; this section provides the most up-to-date exposition, however. Additional descriptions are given in Ashley and Tsang (2013), Ashley and Li (2014), and Ashley and Verbrugge (2018).

To fix ideas, we initially focus on the j^{th} component of X, i.e., column j of the X matrix, corresponding to the $j-1^{st}$ explanatory variable if there is an intercept in the model. The T frequency components are partitioned into M frequency bands, and M $T \times 1$ dimensional dummy variable vectors, $D^{*1}, ..., D^{*M}$, are defined as follows: for elements that fall into the s^{th} frequency band, $D^{*s,j}$ equals $X^*_{\{j\}}$, and the elements are zero otherwise. The regression model can then be generalized as:

$$Y^* = X^*_{\{j\}}\beta_{\{j\}} + \sum_{m=1}^{M} \beta_{j,m} D^{*m,j} + e^* \tag{7}$$

where $X^*_{\{j\}}$ is the X^* matrix with its j^{th} column deleted and $\beta_{\{j\}}$ is the β vector with its j^{th} component deleted.

To test whether the j^{th} component of β is frequency-dependent (i.e., to test whether the effect of the j^{th} variable in X on Y is frequency or persistence dependent) one can then simply test the null hypothesis that $H_0 : \beta_{j,1} = \beta_{j,2} = ... = \beta_{j,M}$.

In the present application, we focus on two columns of X: the real-time inflation rate and the real-time inflation rate; these columns are denoted j and k below. By the same reasoning used above, one can quantify (and test for) frequency dependence in the two model coefficients β_j and β_k corresponding to these two columns by re-writing the regression model (Equation 7) as:

$$Y^* = X^*_{\{j,k\}}\beta_{\{j,k\}} + \sum_{m=1}^{M} \beta_{j,m} D^{*m,j} + \sum_{m=1}^{M} \beta_{k,m} D^{*m,k} + e^*. \tag{8}$$

To make this regression equation a bit more intuitive, one can back-transform Equation (8) back into the time domain by pre-multiplying both sides of this equation with the inverse of A, which (because A is an orthonormal matrix) is just its transpose:

$$A'Y^* = A'X^*_{\{j,k\}}\beta_{\{j,k\}} + A' \sum_{m=1}^{M} \beta_{j,m} D^{*m,j} + A' \sum_{m=1}^{M} \beta_{k,m} D^{*m,k} + A'e^*. \tag{9}$$

This yields the time-domain specification:

$$Y = X_{\{j,k\}}\beta_{\{j,k\}} + \sum_{m=1}^{M} \beta_{j,m} D^{m,j} + \sum_{m=1}^{M} \beta_{k,m} D^{m,k} + e. \tag{10}$$

where $X_{\{j,k\}}$ is the original X matrix, omitting columns j and k and $\beta_{\{j,k\}}$ is the original β vector, omitting components j and k.

Note that now the dependent variable is the same time series (Y) as in the original model, Equation (4). Similarly, all of the explanatory variables – except for the j^{th} and k^{th} – are the same as in the original model. Indeed, the only difference is that these two explanatory variables have each been replaced by M new variables: i.e., the explanatory variable X_j has been replaced by $D^{1,j}...D^{M,j}$ and the the explanatory variable X_k has been replaced by $D^{1,k}...D^{M,k}$. Each of these M variables can be viewed as a bandpass-filtered version of the original data (the j^{th} or k^{th}column of the X matrix), with the nice property that the M frequency component variables corresponding to column j of the X matrix add up precisely to the j^{th} column of X and the M frequency component variables corresponding to column k of the X matrix add up precisely to the k^{th} column of X.

In other words, the j^{th} column of X – for example – is now partitioned into M parts. Reference to definition of the A matrix in Equation (5) shows that the first (lowest-frequency) component utilizing the first row ($s = 1$) of the A matrix and corresponding to $m = 1$, if the first band has only a single component, is proportional to the sample average of the data for this explanatory variable. Similarly, the last component (corresponding to $m = M$, if the last band has only a single component – as it will whenever M is an even number) utilizes the last row ($s = T$) of the A matrix[34] is essentially a sequence of changes in the data; hence this is the highest-frequency component that can be extracted from the data on this variable. To test for frequency dependence in the regression coefficient on this j^{th} regressor, then, all that one need do is test the joint null hypothesis that $\beta_{j,1} = \beta_{j,2} = ... = \beta_{j,M}$. Similarly, X_k is replaced by $D^{1,k}...D^{M,k}$ and one tests the null hypothesis that $\beta_{k,1} = \beta_{k,2} = ... = \beta_{k,M}$.

However, because the A transformation mixes up past and future values (as in any Fourier-based bandpass filter, or any two-sided filter for that matter), it can be shown that these M frequency components are correlated with the model error term e if there is feedback between Y and either of these two explanatory variables, leading to inconsistent estimation of the parameters $\beta_{j,1}, \beta_{j,2}, ...\beta_{j,M}$ and $\beta_{k,1}, \beta_{k,2}, ...\beta_{k,M}$ in that case. Feedback between the federal funds rate and

[34] When T is even; it uses the last two rows when T is odd.

inflation or unemployment rates is certainly likely, so this is an important issue here.

To avoid this problem in general, Ashley and Verbrugge (2009b) suggest modifying the procedure described above in order to obtain a *one-sided* filter for partitioning a variable into its frequency components. In particular, they provisionally[35] suggest decomposing X_j, the j^{th} explanatory variable data vector, into frequency components by applying the transformation described above within a moving window of length T. (Since the discussion from this point forward entirely focuses on extracting frequency component from such moving windows, the symbol T_{full} will henceforth be used to denote the length of the full sample, as this is now distinct from dimensionality (T) of the A matrix for each window. T is a constant – that is, the windows do not increase in length as they move through the sample.) Letting the window for decomposing X_j at time period t consist of the data from time $t - T + 1$ to time t, then for time t only the time t frequency component values (of the T values that are calculated from this window) are retained.

The filtering of X_k is similar to that for X_j, but we note that there is no need to constrain the window length (T) to be identical for both X_j and X_k: this is a modeling decision which should be based on relevant economic theory and on the data themselves. Indeed, we discuss this issue above (in Section 2.1) and do set distinct window lengths for use in obtaining the frequency components for each of the two explanatory variables in the Taylor Rule formulations considered here. In particular, we set a window length of $T_j \equiv 60$ months for decomposing the inflation rate data ("X_j") into its frequency/persistence components, and we set a window length of $T_k \equiv 120$ months for similarly decomposing the unemployment rate data. Thus, the transformation matrix A defined in Equation (5) is of dimension $T_j = 60 \times 60$ months for the inflation rate data and is of dimension $T_k = 120 \times 120$ months for the unemployment rate data. For expositional clarity, however, we focus almost entirely on the case $T_j = 60$ in the remainder of the present sub-section, and below.

Decompositions by frequency must be performed on trendless data, and accordingly, frequency-domain filters typically detrend the sample data prior to performing the decomposition, and add the trend back afterwards. The procedure here is no different in this respect, except that detrending

[35] The word "provisional" is used here because – to fix ideas– this description initially ignores the use of projections for each window; the need for these projections is discussed below.

must occur within each sample window. We consequently instead separately de-trend the data in each window prior to filtering using a linear time trend regression. Upon detrending the data within the window, one implication is that the first (low-frequency) component $D^{1,j}$ is now zero (at machine precision). Of course, these trend estimates are not dropped: after converting the data back to the time domain, Ashley and Verbrugge (2009b) add the time-t estimate of the (within-window) trend to the lowest-frequency component $D^{1,j}$. That is, included in the most persistent component of X_j at time t is the trend estimate of this variable pertaining to time t (i.e., the trend is estimated for the time-t window, and the time-t portion of that trend is added to the most persistent component at t).

It would then, at the outset, appear that the sample average over the T_j observations – corresponding to the zero-frequency component ($m = 1$), and obtained using the first row of the A matrix – thus becomes a moving average of order T_j once the filtering is now being applied to a moving window, so that this first component is now extracting a backward-looking nonlinear trend from the full sample of data, using a moving average of order T_j. It is well-known, however, that moving average trend estimates of this type have very poor properties. In particular, they are known to induce pronounced phase shifts in the estimated trend time series that substantially distort lag length selection and the apparent turning points in the time series.

There is a second weakness in this provisional procedure. When decomposing X_j using a window, one must confront the usual problem of "edge effects" near the window endpoints.

Both the turning point problem and the edge effect problem are addressed by augmenting the data within a window with projections. In particular, as in Dagum (1978), Stock and Watson (1999), Mise, Kim and Newbold (2005) and Clark and Kozicki (2005), we augment the window sample data with projected data, here postpending to the 42 sample data points (ending with period t) in our 60-month π_t windows 18 months of projected values. (And we similarly postpend 24 months of u_t projections to the 96 sample data points (ending in period t) in our 120-month u_t windows). Thus, for example, a 60-month window incorporating the real-time data on X_j as of period t includes 42 past values of X_j – as known at time t – plus projections (forecasts) of its values for months $t + 1$ to $t + 18$. This window of data is then used, as described above, to compute the corresponding $M = 31$ components of X_j – i.e., the vectors $D^{1,j} ... D^{31,j}$. The 42^{nd} element of each of these 60-vectors is then used as the period-t filtered value of X_j for this particular frequency.

In our experience, the estimated values of the coefficient (β_k) on $D^{k,j}$ and its estimated standard error are generally not very sensitive to the number of projection periods chosen (as long as at least 12 months of projections are used), nor to the details of how the projections (forecasts) are produced.[36] Since it is well known that the FOMC makes extensive use of forecasts in its decision-making, utilizing projections of this nature in our windowed bandpass filters seems particularly appropriate. As in the procedure described above – "provisional" because we there omitted discussion of these projections – the estimated trend (now at time t) is included in the zero-frequency component $D^{1,j}$. Because of this window-specific de-trending, where the last 18 months of the data in for example, a 60-month window are projections, the resulting backward-looking non-linear trend time series (which then includes both the zero-frequency component at time t, and also all frequency components with frequency $< \frac{\pi}{30}$, i.e. with reversion periods longer than 60 months) is now a component that embodies the very-low-frequency fluctuations. These complications ensure that this most-persistent component does not exhibit substantive phase and turning point distortions.

This moving-window approach is used in the present paper for an additional, and here crucial, reason. The moving window makes it possible – and, indeed, easy – to use real-time data for the values of X_j and X_k: the data used in each window is simply that available at the time period which is the window's endpoint. The frequency decomposition is in this way gracefully consistent in each period with the data which were available to the policymakers at the time. Conversely, any two-sided approach – whether a conventional analysis based upon gain and transfer analysis (further discussed below), or an approach based upon wavelets – inherently rules out the use of real-time data.

Presuming that T_j and T_k (the window lengths, for decomposing X_j and X_k, respectively) are both even numbers, then reference to Equation (5), and to Table 2 below (for which T is set to 60), makes it evident that there will be $M_j = \frac{T_j}{2} + 1$ distinct frequency components – and hence $\frac{T_j}{2} + 1$ coefficients to estimate in the leftmost sum within Equation (10); similarly there will be $M_k = \frac{T_k}{2} + 1$ distinct frequency components – and hence $\frac{T_k}{2} + 1$ coefficients to estimate in rightmost sum. In the present application, T_j was set to 60 – for the windows used in partitioning the sample variation

[36] See footnote 18 in Section 2.1 for details on the particular projection models used here.

in $\pi(t)$ – yielding 31 distinct frequency components, and T_k was set to 120 – for the windows used in partitioning the sample variation in $u(t)$ – yielding 61 distinct frequency components. While not outright infeasible, it is clearly not desirable to be estimating all 92 of these coefficients. In economic applications, one will generally be interested in more highly aggregated frequency bands, so as to facilitate exposition of one's results on economic or intuitive grounds, as in Ashley and Tsang (2013) and Ashley and Li (2014), where the frequency components are aggregated into just three bands.[37] In Section 2.1 we do the same, defining (for the inflation rate variable, π_t) a low frequency ("very persistent", or "persistent_pi(t)") component corresponding to fluctuations with reversion periods greater than 60 months, a medium frequency ("moderately persistent" or "moderate_persist_pi(t)") component corresponding to fluctuations with reversion periods greater than or equal to 12 months and less than or equal to 60 months, and a high frequency ("transient" or "transient_pi(t)") component corresponding to fluctuations with reversion periods of less than 12 months.[38] Referring to Table 2 (with regard to the 60-period windows for the π_t data), the "persistent_pi(t)" component thus consists of all frequency components with reversion periods greater than the length of the window, 60 months; the "moderate_persist_pi(t)" component is the sum of the five components with intermediate frequencies (components 2-6, corresponding to reversion periods of 12, 15, 20, 30, and 60 months); and the "transient_pi(t)" component comprises the sum of the remaining 25 (highest frequency) components, which all revert more quickly. One might consider this particular *a priori* partitioning of the 31 components (of π_t) or the 61 components (of u_t) into just three bands to be a bit arbitrary. On the other hand, these three aggregated bands are economically interpretable in terms of roughly the same calendar that the FOMC's policymakers live on.

This is a good point at which to contrast the frequency decomposition used here with superficially similar procedures in the existing literature. For example, in contrast to trend-cycle decomposition methods (e.g., Beveridge-Nelson), our approach does not decompose an explanatory variable like X_j into just two components: an arbitrarily-persistent $I(1)$ or $I(1)$-like trend

[37] The A matrix allocates two rows for every non-zero frequency except, for T even, the highest one. The reversion period corresponding to each frequency is proportional to the reciprocal of its frequency, so the number of components with the very highest frequencies are most numerous. For example, in the 120-month windows, 21 of the 60 non-zero frequencies all correspond to reversion periods of a quarter or less.

[38] The frequency components for u_t are analogously named. For this variate, however, the window length is set to 120 months, so the "persistent_un(t)" component includes all frequency components corresponding to reversion periods greater than 120 months.

Frequency Component	Frequency	Reversion Period[a]	Row Number(s) in A[b]
1	0	>60	1
2	$\pi/30$	$60/1 = 60.00$	2,3
3	$2\pi/30$	$60/2 = 30.00$	4,5
4	$3\pi/30$	$60/3 = 20.00$	6,7
5	$4\pi/30$	$60/4 = 15.00$	8,9
6	$5\pi/30$	$60/5 = 12.00$	10,11
7	$6\pi/30$	$60/6 = 10.00$	12,13
8	$7\pi/30$	$60/7 = 8.57$	14,15
9	$8\pi/30$	$60/8 = 7.50$	16,17
10	$9\pi/30$	$60/9 = 6.67$	18,19
11	$10\pi/30$	$60/10 = 6.00$	20,21
12	$11\pi/30$	$60/11 = 5.45$	22,23
13	$12\pi/30$	$60/12 = 5.00$	24,25
14	$13\pi/30$	$60/13 = 4.62$	26,27
15	$14\pi/30$	$60/14 = 4.29$	28,29
16	$15\pi/30$	$60/15 = 4.00$	30,31
17	$16\pi/30$	$60/16 = 3.75$	32,33
18	$17\pi/30$	$60/17 = 3.53$	34,35
19	$18\pi/30$	$60/18 = 3.33$	36,37
20	$19\pi/30$	$60/19 = 3.16$	38,39
21	$10\pi/30$	$60/20 = 3.00$	40,41
22	$11\pi/30$	$60/21 = 2.86$	42,43
23	$12\pi/30$	$60/22 = 2.73$	44,45
24	$13\pi/30$	$60/23 = 2.61$	46,47
25	$14\pi/30$	$60/24 = 2.50$	48,49
26	$15\pi/30$	$60/25 = 2.40$	50,51
27	$16\pi/30$	$60/26 = 2.31$	52,53
28	$17\pi/30$	$60/27 = 2.22$	54,55
29	$18\pi/30$	$60/28 = 2.14$	56,57
30	$19\pi/30$	$60/29 = 2.07$	58,59
31	$30\pi/30$	$60/30 = 2.00$	60

Table 2: **Frequencies and Reversion Periods for a 60-Month Window** [a]In months, calculated as 2π divided by the frequency. The sinusoids comprising the elements of the row(s) of the A matrix corresponding to this reversion period complete a full cycle in this many months. Thus, the scalar product of such a row with a time-series vector whose fluctuations self-reverse substantially slower than this will be very small. [b]The A matrix is defined in Equation (3).

and a stationary $I(0)$ fluctuation. Our decomposition instead produces M components (adding up to X_j) which span a complete range of persistence levels. And it allows the data itself – via regression analysis applied to Equation (10), or the closely related Equation (3) in Section 2.2.1 – to quantify how the coefficients $\beta_{j,1}, ..., \beta_{j,M}$ vary across all of these persistence levels. Further, our decomposition still yields consistent parameter estimation where (as is typically the case with economic relationships) one cannot rule out feedback (or bi-directional causality); this is in contrast to the earlier spectral regression models cited at the outset of this section, which employ two-sided filtering and hence yield inconsistent parameter estimation in the presence of feedback. Finally, our approach is uniquely appropriate to the present analysis of the FOMC's Taylor Rule behavior, because the central bank surely bases its actual policy decisions on real-time data. In particular, the current real-time history of each of the relevant explanatory variables (the inflation and unemployment rates) corresponds exactly to the data which we use in each window for the decomposition of the current value of each variable into its frequency/persistence components.

We note, in this context, that an analogous kind of analysis based on the gain and phase of a transfer function model for the federal funds rate – as in Box and Jenkins (1976, Part III) – would be problematic because such models characteristically involve lagged values of the dependent and explanatory variables. For one thing, models containing lagged variables are inherently awkward when using real-time data because it is not clear whether the period-t datum to be used for X_j lagged, say, two periods should be the value of for that period as known currently (i.e., in period t) or at the time (i.e., in period $t-2$). In addition, transfer function gain and phase plots are substantially more challenging to interpret than our $\beta_{j,1}, ..., \beta_{j,M}$ coefficients, especially where (as here) bi-directional causality is likely. For example, Granger (1969) notes, "in many realistic economic situations, however, one suspects that feedback is occurring. In these situations the coherence and phase diagrams become difficult or impossible to interpret, particularly the phase diagram."

4.2 The Appeal of this Frequency-based Approach to Disaggregation by Persistence Level

The focus of this paper is to investigate, in a data-driven way, the degree and manner to which the FOMC has responded to persistent innovations in the unemployment rate and the inflation rate differently than it has to more transitory fluctuations in those variables. Thus the objective of partitioning two of the explanatory variable time series – X_j and X_k in Section 4.1 above, which are the real-time unemployment and inflation rates in the present application – is not the bandpass filtering *per se*. Rather, we decompose the unemployment and inflation rates into frequency components so that we can separately estimate the historical impact of fluctuations of distinctly different persistence levels in these two variables on the federal funds rate and make inferences concerning these differential impacts; this allows a richer consideration of how FOMC policy has varied over the several time periods considered.

No representation is made here that the bandpass filtering described in Section 4.1 above is asymptotically optimal – e.g., as in Koopmans (1974) or Christiano and Fitzgerald (2003) – although the relevance of asymptotic optimality in filtering windows of data which are only of length 60 or 120 months, and which might need to be considerably smaller in other applications, is debatable.[39] On the other hand, our method of decomposing a time series into M frequency components has several very nice characteristics, which make this decomposition approach overwhelmingly well-suited to the present application:

1) The M frequency components that are generated from an explanatory variable (i.e., from a column of X) by construction partition it. That is, these M components add up precisely to the original observed data on this column of X. This makes estimation and inference with regard to persistence dependence (or its inverse, frequency dependence) in the corresponding regression coefficient particularly straightforward: we simply replace this explanatory variable in the regression model by a linear form in the M components and analyze the resulting M coefficient estimates.

2) Due to the moving windows used, this particular way of partitioning the data on an explanatory variable into these M frequency components by construction utilizes backward-looking (i.e., one-sided) filters. As demonstrated in Ashley and Verbrugge (2009b), this feature is crucial to consistent OLS coefficient estimation where there is bi-directional Granger-causality (i.e., feedback)

[39] In this context we note that it feasible – albeit somewhat awkward – to iteratively employ a Christiano-Fitzgerald (2003) low-pass filter to partition the data in such a way that the frequency components still add up to the original data. This procedure involves applying the filter repeatedly, at each iteration varying the frequency threshold and applying the filter to the residuals from the previous iteration. This procedure is, of course, no longer even asymptotically optimal, but it does yield frequency components which still add up to the original data – as the ones in our present formulation do automatically. Experiments with decompositions along these lines did not yield substantially different frequency dependence results in the present application, but we are investigating this topic in future work.

between the dependent variable and the explanatory variable being decomposed by frequency. The dependent variable in the present context is the federal funds rate, which is quite likely to be in a feedback relationship with the unemployment and inflation rates.

3) Finally, this way of partitioning the data on an explanatory variable into frequency/persistence components is not just mathematically valid and straightforward, it is also intuitively appealing. In particular – in contrast to many analyses in the frequency domain – our decompositions are not a 'black box.' The next section illustrates this point with a simple example.

4.3 An Illustrative Example with a Very Short Window

An example with a window ten periods in length illustrates the sense in which the frequency components defined above are extracting components of, say, X_j of differing levels of persistence. This window length is sufficient large as to illustrate the point, while sufficiently small as to yield an expositionally manageable example.[40] In particular, Table 3 displays the multiplication of the matrix A – whose elements are defined in Equation (5) – by the ten-component sub-vector of X_j corresponding to a window beginning in the particular period 21 and ending in period 30. For this illustrative example, we will assume that there are no projections, so that this window corresponds to time period 30. We will also assume that the data are trendless over these 10 observations.[41]

The first row of the A matrix is just a constant. The operation of this row of A on this particular ten-dimensional sub-vector of X_j is just calculating the sample mean over these ten observations. The operation of this row of A on this particular ten-dimensional sub-vector of X_j is just calculating the sample mean over these ten observations. Thus, as this window progresses through the entire sample of data X_j, the first component of the vector formed by multiplying each ten-dimensional sub-vector of X_j on the left by A represents a one-sided, real-time, nonlinear trend estimate based (in this example) on a 10-period moving average. This is the "zero-frequency" component of the full

[40] As described above, the empirical implementation in this paper uses a window 60 months in length (for π_t) and 120 months in length (for u_t). See Table 2 for an explicit listing of the component frequencies, the corresponding reversion periods, and the corresponding A matrix rows for a 60-dimensional A matrix.

[41] As noted in Section 4.1 above, in actual practice a portion of every moving window consists of projections, and data within every moving window are detrended. These complications are suppressed in the present subsection so as to focus attention on the A matrix in a setting so simple as to elucidate how the application of this matrix is extracting components for which the terms "frequency" and "reversion period" are intuitively meaningful verbal constructs.

Period	Matrix A										Data
> 10	0.32	0.32	0.32	0.32	0.32	0.32	0.32	0.32	0.32	0.32	$X_j(21)$
10	0.45	0.36	0.14	-0.14	-0.36	-0.45	-0.36	-0.14	0.14	0.36	$X_j(22)$
10	0.00	0.26	0.43	0.43	0.26	0.00	-0.26	-0.43	-0.43	-0.26	$X_j(23)$
5	0.45	0.14	-0.36	-0.36	0.14	0.45	0.14	-0.36	-0.36	0.14	$X_j(24)$
5	0.00	0.43	0.26	-0.26	-0.43	0.00	0.43	0.26	-0.26	-0.43	$X_j(25)$
3.3	0.45	-0.14	-0.36	0.36	0.14	-0.45	0.14	0.36	-0.36	-0.14	$X_j(26)$
3.3	0.00	0.43	-0.26	-0.26	0.43	0.00	-0.43	0.26	0.26	-0.43	$X_j(27)$
2.5	0.45	-0.36	0.14	0.14	-0.36	0.45	-0.36	0.14	0.14	-0.36	$X_j(28)$
2.5	0.00	0.26	-0.43	0.43	-0.26	0.00	0.26	-0.43	0.43	-0.26	$X_j(29)$
2	0.32	-0.32	0.32	-0.32	0.32	-0.32	0.32	-0.32	0.32	-0.32	$X_j(30)$

Table 3: **An Example With a Window of Length Ten Periods**. The first row of A time the data vector simply yields $1/\sqrt{T}$ times the sample mean of the data in this ten-period window. As the window moves through the data set, this operation extracts any, possibly nonlinear, trend as a moving average. Rows two and three take a weighted average of the window data, using smoothly-varying weights which take a full ten periods to reverse, so any fluctuation in window data that reverses in a couple of periods yields a small value. The product of row ten and the window data is essentially calculating five changes in the data which occur during the window period. A long, smooth variation in the window data yields a small value for this frequency component.

X_j vector, corresponding to a sinusoidal reversion period unbounded in length. This component of X_j includes all of its variation at frequencies so low (i.e., reversion periods so large) that they are essentially invisible in a window which is only ten periods in length.

Higher-frequency (lower persistence) components of X_j are, conversely, distinguishable using this window. The "Period" column in Table 3 is the number of observations over which the sine or cosine used in the corresponding row of the A matrix completes one full cycle. This is ten observations for rows two and three of this A matrix, $\frac{10}{2} = 5$ observations for rows four and five, $\frac{10}{3} = 3\frac{1}{3}$ observations for rows six and seven, $\frac{10}{4} = 2\frac{1}{2}$ observations for rows eight and nine, and $\frac{10}{5} = 2$ observations for row ten.[42] In the most common convention, the frequency is defined as $\frac{\pi}{2}$ times the inverse of cycle length (period) of the corresponding sine or cosine for that row of the A matrix, in which case the frequencies run from zero (for row one) to π for row ten.

To see intuitively why multiplication of the X_j vector by, for example, rows two and three extract only slowly-varying fluctuations in X_j, notice that these two rows are smoothly varying weights that will be applied to the ten components of X_j in forming its dot (or scalar) products with these two rows. Slowly-varying fluctuations in X_j will thus have a large impact on these two

[42] The number of observations in the sub-vector is an even integer – ten – in this example, implying that the sine and cosine terms are multiples of one another for what becomes a singleton last (tenth) row of the A matrix.

dot products, whereas rapidly-reverting variations in X_j will have little effect on the values of these two dot products. Hence, components two and three of the matrix product AX_j will 'contain' only those parts of X_j which are slowly varying.

Conversely, it is evident upon inspection of the last row of the A matrix that only high-frequency (low persistence) fluctuations – i.e., fluctuations which reverse in just two months or so – will contribute significantly to the tenth component of AX_j.

Thus, the first rows of the A matrix are distinguishing and extracting what are sensibly the "low-frequency" or "large period"or "highly persistent" or "relatively permanent" components of this ten-month X_j sub-vector as the window moves through the sample. Conversely, the last rows of the A matrix are distinguishing and extracting what are sensibly the "high-frequency" or "small period"or "low persistence" or "relatively temporary" components of this X_j sub-vector.

4.4 Application to a Simple DGP

In this section, we illustrate the decomposition used here via its application to an artificially generated time series simulated using a particularly simple data generating process.[43]

The data generating process is comprised of the sum of four distinct sine waves, plus a white noise error. More specifically,

$$y_t = 2x_{1,t} + x_{2,t} + x_{3,t} + 0.5x_{4,t} + u_t$$

with $u_t \backsim N(0,1)$, where

Component	*Period*
$x_{1,t} = \sin(0.05t)$	125.7
$x_{2,t} = \sin(0.14(t+25))$	44.9
$x_{3,t} = \sin(4t)$	1.57
$x_{4,t} = \sin(22t)$	0.29

[43] We are not suggesting that u_t or π_t is generated in this way; the intent here is solely to present an illustrative example.

We apply our one-sided filtering method, as described above,with moving windows of length 60, of which 24 are projections. Here the projections are obtained from the average of forecasts from a univariate AR(4) model in levels and forecasts from a univariate AR(4) model in first-differences. Here the "very persistent" component is defined to include all fluctuations whose period is longer than 58 time units and simply denoted "persistent" in Figure 3. The cutoff for the transient component is set to a reversion period of 12 time units. Thus, "persistent" should reflect $x_{1,t}$ and "moderate" in Figure 4 should reflect $x_{2,t}$.

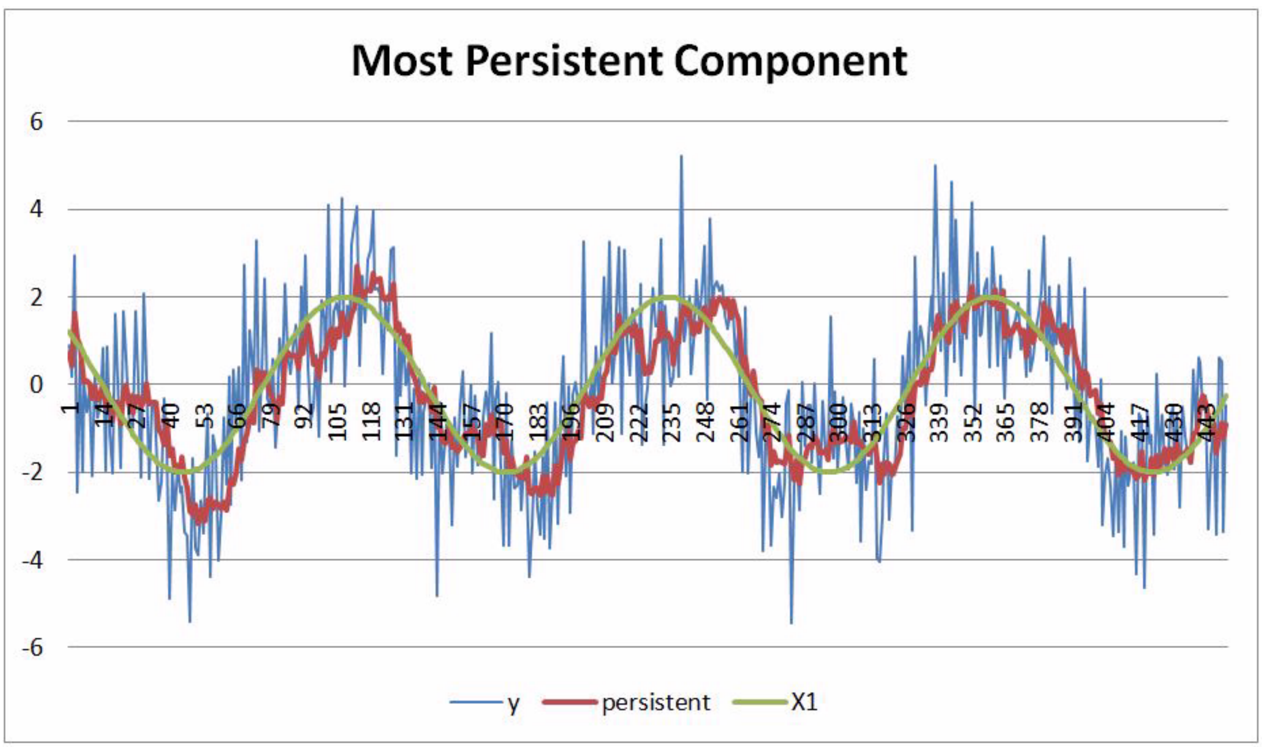

Figure 3: For the simple DGP, our procedure accurately estimates the most persistent component

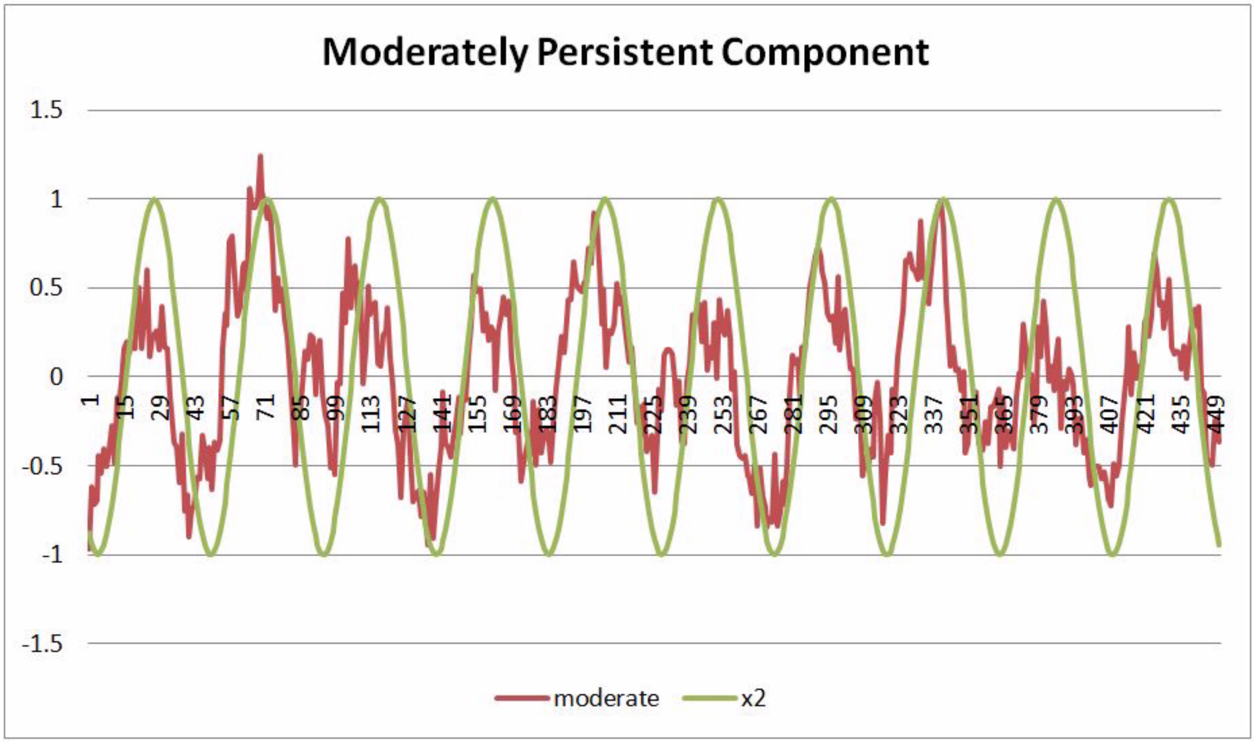

Figure 4: For the simple DGP, our procedure accurately estimates the moderately persistent component

As described above, the most persistent component (here simply labeled "persistent") is effectively estimated using a nonlinear adaptive trend. Its behavior is mainly driven by the window length, which determines the frequency cutoff, and by the length and quality of the projections within that window.

As is evident in Figure 3, our procedure does a reasonable job of estimating $x_{1,t}$, though it is modestly fooled by moderately persistent movements deriving mainly from $x_{2,t}$ And as is evident in Figure 4, our procedure likewise does a reasonable job of estimating $x_{2,t}$.[44]

[44] A linear regression of the most persistent component on $x_1, ..., x_4$ yields a coefficient of 0.93 on x_1 and 0.53 on x_2, with small loadings on x_3 and x_4. A linear regression of the moderately persistent component on $x_1, ..., x_4$ yields a coefficient of 0.00 on x_1 and 0.37 on x_2, with very small loadings on x_3 and x_4. Finally, a linear regression of the transient component on $x_1, ..., x_4$ yields a coefficient of 0.05 on x_1 and 0.06 on x_2, with loadings of 0.90 and 0.91 on x_3 and x_4.