What Can We Learn from 60 Years of PCE Inflation Data?*

Raphael Schoenle
Brandeis University

Dominic A. Smith
U.S. Bureau of Labor Statistics

September 22, 2022

Abstract

Analyzing the distribution of disaggregated PCE category inflation rates, we detect systematic changes in the distribution between 1960 and 2021. Pre-1990, extreme positive tails characterize the distribution, but they moderate post-1990 while more negative tails appear. The distribution is granular, with an increasing importance of granularity over time. The ranking of mean inflation versus robust measures of inflation—medians and trimmed means—inverts several times. The covariance of disaggregated inflation rates decreases more than the variance over time. Our empirical findings point to the use of multi-sector models when appropriately analyzing the stabilization properties of monetary policy. In an application to oil price shocks, we show how the choice of policy regime interacts with the distribution of inflation rates and the measure of aggregate inflation.

Keywords: Inflation, PCE, Granularity

JEL-Codes: C43, E31, E37

*We appreciate feedback from Santiago Alvarez, Raphael Auer, Brian Adams, and David Argente as well as members of the Price Index Number Research Division at the BLS and seminar participants at the AEA meetings, Midwest Economic Association Meetings, and the Bureau of Economic Analysis.

1 Introduction

Understanding inflation is important for macroeconomists and policy-makers. Following the Great Recession inflation was characterized by puzzlingly low rates while double-digit rates were a dominant feature of the inflation experience during the 1970s and 1980s. More recently, inflation appears to be back at elevated rates in the aftermath of the COVID-19 pandemic. This increase in rates has put inflation at the forefront of public interest.

In this paper, we provide new empirical and theoretical insights into inflation as we analyze the *distribution* of its underlying disaggregated inflation rates. Much is hidden behind aggregate measures of inflation. We find that the cross-sectional distribution of disaggregated inflation rates in detailed Personal Consumption Expenditure (PCE) has systematically changed between 1960 and 2021: First, extreme increases in inflation have become rarer, and extreme decreases have appeared. Second, inflation is granular, and importance of granularity has increased over time. Third, the ranking of headline inflation versus robust measures of inflation inverts several times in our data. Fourth, the covariance of inflation rates decreases more than the variance over time. These findings suggest that the inflationary process post-1990 has been driven by granular, idiosyncratic shocks rather than aggregate shocks.

On the theory side, we show that a heterogeneous production model with idiosyncratic shocks is needed to match these facts while providing new insights into the stabilization properties of monetary policy. In such a framework, the choice of policy regime interacts with the distribution of inflation rates and the choice of aggregate inflation measures. To demonstrate the importance of this interaction, we analyze the interaction of average inflation targeting (AIT), as well as a Taylor-type monetary policy rule, with a particular measure of inflation—core inflation—in a calibrated version of the model subject to an idiosyncratic shock in the oil producing sector. We find that in such a scenario targeting core inflation under a Taylor-type rule rather than headline inflation achieves much of the inflation stabilization gained from AIT. However, focusing on core inflation with AIT yields additional benefits. We show that such gains can generally arise for an appropriately defined monetary policy regime in the face of any idiosyncratic shock.

Our analysis begins by presenting new facts about the distribution of highly disaggregated inflation series that underlie the aggregate inflation process. As a basis for this analysis, we first construct a consistent set of disaggregated monthly PCE inflation rates. These data go back to 1959 and cover 98 percent of the aggregate consumption basket, while allowing us to replicate the official PCE headline inflation series. The facts we present derive both from the cross section and the time series.

First, our analysis shows that changes in the cross-sectional distribution appear quite evidently at the extremes: Extreme increases in inflation rates have become rarer, and extreme decreases have appeared over time. Before 1990 there were large positive shocks to inflation. The 99th and 90th percentiles experienced periods of very high inflation. The 1st and 10th percentile of inflation were fairly constant and rarely very negative. After about 1990 this pattern changed with the larger percentiles of the inflation distribution becoming smaller, particularly the 90th percentile. The 1st percentile of inflation saw large declines into the negative. More generally, changes in the shape of the distribution are reflected in changes of other higher moments, such as skewness which has declined over time, and kurtosis which has increased over time.[1]

Second, a few series have a disproportionately large impact on inflation: inflation is granular in the sense of Gabaix (2011). The top 10 out of 183 personal consumption expenditures (PCE) categories account for 41 percent of PCE inflation in 2019. The distribution of inflation rates exhibits fat tails on the positive side before 1990 and on the negative side after 1990, in line with the swings observed for the extremes of the distribution. On top of this change, the quantitative impact of granularity on headline inflation has also increased. We show this trend using a time-series decomposition into an equal-weighted component and a granular residual as in Foerster et al. (2011). In the 1970s and 1980s, the equal-weighted component—which is large when aggregate shocks are important—was the dominant contributor to headline inflation. During the same period, the granular residual—which is large when a few observations are disproportionately influential—accounts for only 3 percent of aggregate inflation. However, post-1990, its contribution rises to 23 percent, and to 33 percent in 2020-2021. This rise in importance suggests a heightened importance of idiosyncratic components to the inflationary process.

Third, the inflation readings from mean and robust measures of aggregate inflation, such as trimmed mean or median inflation, are directly related to the importance of granularity and the systematic changes in the extremes of the distribution. In particular, the relative ranking between mean and robust measures of inflation has reversed several times in our data period: During the high-inflation regimes before 1990, mean inflation exceeds median and trimmed mean inflation by one percentage point, on average. During the ensuing low-inflation regime, median and trimmed mean inflation have typically been higher, especially since 2010, by approximately half a percentage point.[2]

Finally, a systematic change in the variance-covariance structure complements the characterization of changes in the distribution of inflation rates over time: The covariance of disaggregated inflation rates has decreased more than the variance over time. This results follows from a decomposition of the variance in headline inflation into two components. The first is the variance of disaggregated inflation rates and the

[1] Other measures of skewness such as kelley and bowley skewness have also changed over time. (Verbrugge and Zaman, 2022)
[2] Mean inflation has again been above median inflation since March of 2021 (Ocampo et al., 2022).

second is the covariance of inflation terms. We find that the covariance of disaggregated inflation rates has substantially declined over time and then exhibits a stable distribution for the last 30 years. While the overall variance has also declined, its decline has been less than the decline of the covariance. Together, these findings suggest that granular, idiosyncratic shocks have likely been playing an increasingly important role for the inflationary process over time, and for the readings of inflation from various aggregate measures.

A heterogeneous production model with idiosyncratic shocks can rationalize these facts, while also providing new insights into the stabilization properties of monetary policy. We show these two results based on a model following Pasten et al. (2020). In this setup, sectors differ in size, the degree of price rigidity and input-output linkages. Sectors may be subject to idiosyncratic as well as an aggregate productivity shock. Monetary policy follows a Taylor-type monetary policy rule that targets headline inflation. We set the same parameters for the economy as in Pasten et al. (2020), and run two model exercises. The first exercise establishes that a model with idiosyncratic shocks is needed to match the four facts presented. The second exercise then shows how both the choice of monetary policy rule and targeted measure of inflation can affect the distribution of inflation rates, and lead to different, but additive stabilization gains.

The first exercise establishes the need to use a heterogeneous, multi-sector model of the economy with idiosyncratic shocks to match the four facts presented. We analyze six different calibrations to arrive at this result. Each calibration represents a combination of an aggregate shock and/or idiosyncratic shocks that hit either one, all, or a subset of sectors. We find that only the inclusion of idiosyncratic shocks allows us to match all four facts. There is no need to include aggregate shocks. In fact, aggregate shocks alone cannot replicate the fat tails in the distribution and the larger drop in the covariance than the average variance. A direct implication of these exercises is that a conventional (one-sector) model with an aggregate shock cannot replicate the features of the inflationary process as in the data.

The second exercise shows the importance of these insights for modeling the inflationary process and for the strategies monetary policymakers may pursue when they face a cycle driven by select idiosyncratic shocks. Building on a setup with idiosyncratic shocks, we show how both individually but also in combination two choices in particular affect the distribution of inflation rates and the stabilization of inflation through monetary policy. These choices concern the policy regime—a Taylor rule versus AIT—and the targeted measure of inflation—headline versus core inflation. Our analysis in particular focuses on a shock to the oil producing sectors, but we show the results hold for shocks to other sets of industries. We find three results in our calibrated economy: First, AIT stabilizes the distribution more than a Taylor-type policy rule, especially by compressing the center of the distribution more. Moving from a Taylor-type rule to AIT reduces inflation volatility overall by about a factor of one-half while keeping inflation fluctuations stable. This finding is in line with the well-known stabilizing properties of average inflation targeting. Second, we find that targeting either core inflation or inflation in the set of industries not hit by the shock achieves most of the benefits for inflation stabilization as AIT, while leaving consumption volatility unchanged. Third, stabilization of core inflation and also shifting to an AIT regime creates additional inflation stabilization benefits. The fact that key welfare metrics improve relative to a Taylor rule and a size-weighted mean measure of aggregate inflation is not surprising given the limitations of a Taylor rule (see e.g.) and the further complications of optimal policy in multisector models as discussed for example in Rubbo (2020) or Jennifer and Tahbaz-Salehi (2020).

Our paper connects to several strands of the literature on pricing and monetary economics. First, a literature too large to summarize has both theoretically and empirically studied various moments of pricing mainly in the cross section, including Nakamura and Steinsson (2008), Bils and Klenow (2004), Bhattarai and Schoenle (2014), Gagnon (2009), Alvarez et al. (2016), Midrigan (2011), Bonomo et al. (2020), or Karadi et al. (2021). We add to this focus on moments an explicit focus on interaction of the entire distribution with aggregate inflation measures, and the evolution of this relationship over time. Vavra (2014), Luo and Villar (2021) and Nakamura et al. (2018) take a related perspective by considering variation over time in terms of volatility, skewness, frequency and size of CPI micro price changes. Second, our focus on granularity in the inflationary process is motivated by the seminal work of Gabaix (2011). A subsequent literature has shown the pervasive relevance of granularity for mostly real macro variables to which we add a focus on price inflation. Finally, our paper connects to the work on heterogeneous production networks, such as Acemoglu et al. (2012), Acemoglu et al. (2017), Pasten et al. (2020, 2021), Baqaee and Farhi (2020), or Rubbo (2020). We focus on the role monetary policy plays via the distribution of inflation rates for measures of aggregate inflation.

2 Data and Methodology

The primary data for this project come from the underlying data supplements of the National Income and Product Accounts Personal Consumption Expenditure (PCE) data release (Bureau of Economic Analysis, 2022). The PCE data provide highly disaggregated price indexes and expenditure weights that cover the entire U.S. consumption basket. The data we use span the period from 1959 to 2021. The Bureau of Economic Analysis (BEA) attempts to construct series under a consistent methodology from 1959 onwards through regular revision.

Figure 1: Official and Constructed Headline Inflation

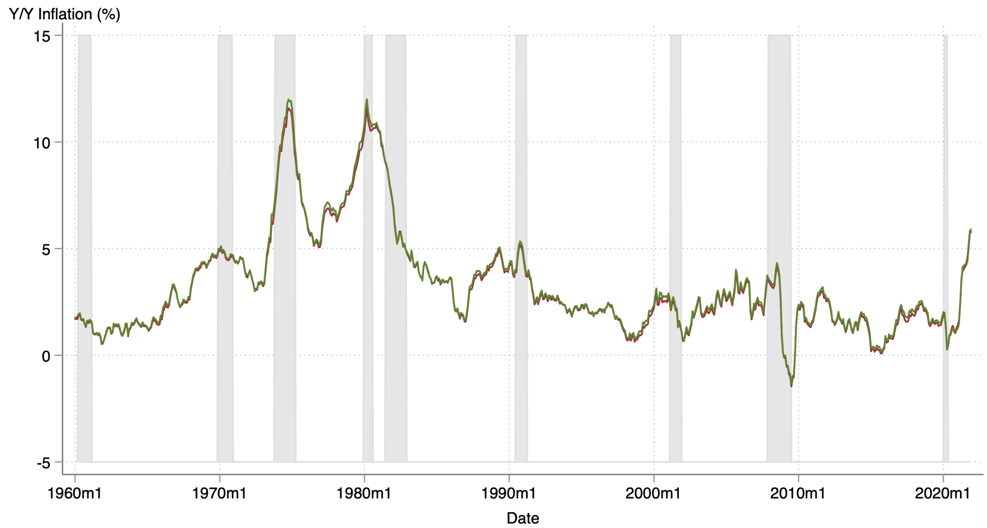

Notes: Authors' calculations from the PCEPI data. The figure plots headline inflation as reported in the PCEPI (series DPCERG) against the authors' calculations of headline inflation as a Laspeyres index (see equation 1) of the 183 component categories. The average difference between the two lines is 0.11 with a maximum difference of 0.54 percentage points. Shaded areas indicate recessions as defined by the NBER.

We select a set of 183 spending categories that are consistent over time and partition consumer spending. Some series are introduced in later years because they contain new goods such as mobile phones. When this occurs we assume the spending on this category was 0 prior to its inclusion in the PCE price index (PCEPI) and that its price index was undefined prior to inclusion. Our PCE categories cover 98 percent of the aggregate consumption basket, excluding only final consumption expenditures of nonprofit institutions (series DNPIRC and its subcategories).

The PCE data contain price indexes for each of the 183 consumption categories and estimates of consumer spending in those categories. To confirm that these underlying data are valid for our analysis, we replicate the official published PCE headline inflation series. To do so, we construct headline inflation as a Laspeyres Index of these underlying components. The headline inflation rate we construct based on this methodology and our data is extremely close to the headline PCE inflation rate published by the BEA, as Figure 1 shows. Our numbers differ only slightly from the official headline series because some components of headline inflation are aggregated before the BEA publishes their data and the BEA uses a Fischer Index. See Appendix A.1 for a comparison of the official headline inflation series and inflation series constructed using various index formulas.

We compute this measure of headline inflation over a twelve month period ending in month t as a function of the prices of individual components, p_{it}, and the associated weights of those components, w_{it-12},

$$\pi_t^{12} = \left(\sum_i w_{it-12} \left(\frac{p_{it}}{p_{it-12}} - 1 \right) \right) * 100 \tag{1}$$

where $t-12$ denotes a twelve-month lag and p_{it} price indexes in each consumption category i. Our subsequent analysis centers on the individual inflation components $\pi_{it} = \left(\frac{p_{it}}{p_{it-12}} - 1 \right) * 100$. We describe the construction of our moments of interest where appropriate in the text below and the appendix. We focus on twelve month inflation rates to mitigate concerns about seasonality in our analysis.

Table 1 presents summary statistics for the average, standard deviation, and percentiles for our sample. We separately present numbers for the entire period of our study and for two separate periods. Average inflation falls from 4.50 percent per year to 2.12 percent per year in the second period. All of the percentiles fall between the two periods. Even the 10th percentile of inflation is positive on average in the first half of the sample.

Table 1: Inflation Summary Statistics

Statistic	1960-2021	1960-1989	1990-2021
Mean	3.31	4.50	2.12
S.D.	2.47	2.92	0.99
Skewness	2.56	3.90	0.59
Kurtosis	67.03	90.65	31.49
1st Percentile	−9.56	−7.24	−11.74
10th Percentile	−0.79	0.10	−1.62
25th Percentile	1.74	2.62	0.92
Median	3.43	4.45	2.48
75th Percentile	5.04	6.41	3.75
90th Percentile	7.37	9.21	5.65
99th Percentile	19.92	23.33	16.72
Granular Residual Mean	0.33	0.14	0.49
Granular Residual S.D.	0.47	0.51	0.32
Equal Weighted Mean	2.98	4.36	1.63
Equal Weighted S.D.	2.42	2.68	0.97

Notes: Numbers are summary statistics for π_{it} over a given time period. The mean and standard deviation (s.d.) refer to aggregate inflation over the time period. Skewness and kurtosis from the distribution of individual series' inflation rates over time. The percentiles are calculated for each month on a weighted basis and then averaged over time. The granular residual and equal weighted components are calculated according to equation 4 and averaged over the time period. Authors' calculations from the PCE.

2.1 Contribution of a Consumption Category to Inflation

As part of the subsequent analysis, we calculate the contribution of a consumption category to the headline rate of inflation. We do so in two ways. We primarily use the weight of each item in the headline index multiplied by its change in price normalized by the headline rate of inflation,

$$c_{it} = \frac{w_{it-12}(\pi_{it} + 100)}{\pi_t^{12} + 100}. \tag{2}$$

This approach has the advantage that all categories have a positive contribution and areas with more inflation or higher weight will have larger contributions. However, any item with deflation will have a small, but positive, contribution because $\pi_{it} + 100$ is always greater than 0. Therefore, we also measure what we call the absolute contribution of each item to the overall index as

$$\bar{c}_{it} = \left| w_{it-12}\pi_{it} \right|. \tag{3}$$

This approach leads to categories having large contributions when their weight is large and prices are changing, whether they are increasing and decreasing.

2.2 Additional Data

We use additional data when we calibrate our model. The model features producers of goods setting prices who use the output of other sectors in their production as intermediate inputs. To calibrate the production structure of the economy, we focus on 341 industries. These sectors are defined based on the North American Industrial Classification System (NAICS) codes and the BEA input-output industry codes. Some of these industries correspond closely to components of the PCE, but in other cases, particularly in services, direct links to the PCE are not possible.

We discipline the size of each sector and the input-output network using data on the input-output structure of the U.S. economy constructed in Pasten et al. (2020): These data, based on the BEA input-output tables, include industry-by-industry trade flows and sectoral GDP shares. We also account for the fact that industries differ in the frequency with which they change their prices. This dimension of heterogeneity has been shown to be important in Pasten et al. (2020, 2021) from whom we use data for 341 sectoral frequencies of price changes.

3 Facts

The cross-sectional distribution of disaggregated inflation rates in the PCE has systematically changed between 1959 and 2021: First, extreme increases in inflation have become more rare, and extreme decreases have appeared. Second, inflation is granular, and importance of granularity has increased over time. Third, the ranking of headline inflation versus robust measures of inflation inverts several times in our data. Fourth, the covariance of inflation rates decreases more than the variance over time. These empirical findings provide important modeling guidance suggesting that the inflationary process post-1990 has been driven by granular, idiosyncratic shocks rather than aggregate shocks.

3.1 Extreme Increases in Inflation More Rare, Decreases Appearing

As the level of headline inflation has decreased there have been significant changes to the cross-sectional distribution of inflation as shown in Figure 2. In this figure, we rank all PCE categories from lowest to highest inflation rates in each month. We then create a percentile X by choosing the inflation rate of the category such that $X\%$ of the expenditure in the PCE has an inflation rate less than or equal to that category inflation rate. Before 1990 there were large positive shocks to inflation. The 99th and 90th percentile would experience occasional periods of very high inflation. The 1st and 10th percentile of inflation were fairly constant and rarely very negative. After about 1990 this pattern changed with the larger percentiles of the inflation distribution being smaller, particularly the 90th percentile. On the other hand, there were more large declines in the 1st percentile of inflation.

These two features of the distribution which we have highlighted have a very powerful implication: Which measure of aggregate inflation we choose—means or robust measures—in order to read off aggregate inflation will be heavily influenced by the entire distribution. For example, as Figure B.1 in the appendix illustrates, robust, central measures of CPI and PCE inflation were consistently above two percent between 2010 and 2020. At the same time, average-based measures exhibit protracted periods of low and near zero percent inflation. What do these differences mean for judging the success of central banks in hitting a two percent inflation target? What models should we use to incorporate insights into the inflation distribution? What effects will different monetary policy rules have on the distribution and hence inflationary readings? We next turn to more comprehensive study of the inflation distribution both in the cross section and the time series, before use a model to study these questions.

Figure 2: Percentiles of Inflation Distribution

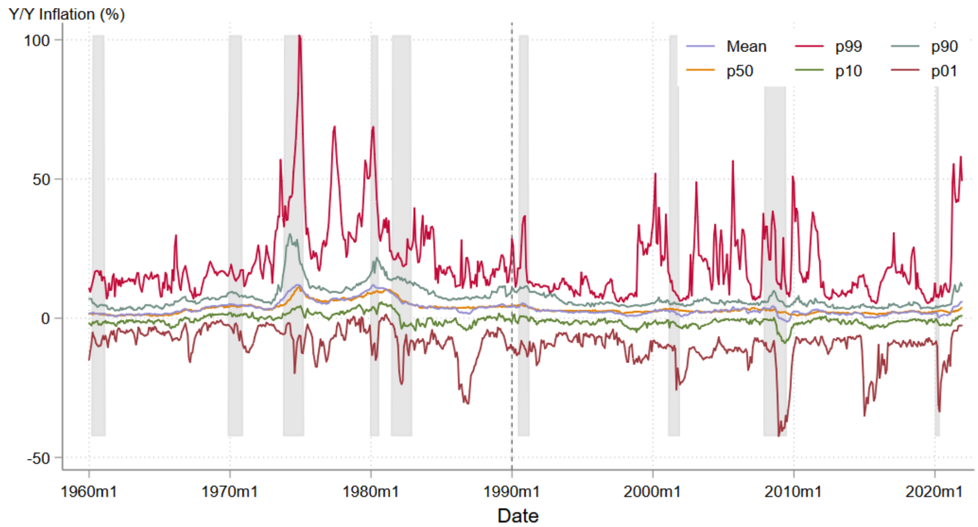

Notes: Lines plot the percentiles of category PCEPI yearly inflation changes. Percentiles are calculated by ranking categories from lowest to highest inflation. Each category is assigned a weight according to the fraction of spending associated with the category. Then, the inflation rate is assigned to each percentile is such that X percent of the weight in the PCE has inflation less than or equal to that percentile. Authors calculations from the PCE. The order of the lines follows the order of the legend with the exception that the mean and is always near p50 (which indicates the 50th percentile). Shaded areas indicate recessions as defined by the NBER. See Appendix B for more information.

3.2 Inflation is Granular

The distribution of inflation rates is "granular." We use different methods to illustrate this fact. A common definition of granularity would suggest that the contributions of a few categories to headline inflation are disproportionately larger than others. A stricter definition of granularity is in terms of fat tails (see, for example, Gabaix (2011)): When the distribution of (absolute) inflation rates exhibits fat tails, then some sectoral rates are disproportionately large and granular at any level of disaggregation.

Disaggregated inflation rates are granular according to both of these definitions. First, their contributions to headline inflation are concentrated among a few categories. Second, a log-normal distribution approximates the distribution well at the category level.

Fact 1 *Inflation rates are granular:*

- *The top 10 out of 183 personal consumption expenditure (PCE) categories account for 41 percent of PCE inflation in 2019. The top 15 categories account for 50 percent.*

- *The distribution contributions of categories to headline inflation has fat tails, in particular, it is approximately log-normal.*

The first sense in which aggregate inflation is granular is that the contributions to aggregate inflation are highly concentrated among a few PCE categories. We make this point using Lorenz curves in Figure 3. Lorenz curves summarize the inequality in the contribution to aggregate inflation across the PCE categories. The 45-degree line in the figure indicates perfect equality in these contributions. The area between the 45-degree line and the Lorenz curves is the Gini coefficient, measuring the extent of inequality in the contributions to the aggregate. We find that in all periods, as well as on average across periods, concentration of contributions to inflation is high. For example, we find that the 10 (out of 183) largest contributors to headline PCE inflation account for about 41 percent of PCE inflation in 2019, and the top 15 percent of PCE categories account for 50 percent. The top 10 contributors to headline PCE inflation are relatively constant over time and include owner-occupied housing, non-profit hospitals, and physician services.

Inflation is also granular in a statistically more rigorous sense: The distribution of category contributions to headline inflation is fat-tailed (Gabaix, 2011) and, in particular, well approximated by a log-normal distribution. A simple way to illustrate this point is to look at a Q-Q (quantile-quantile) plot, in which we plot the actual percentiles of the log absolute contributions against a set of percentiles from a simulated log-normal distribution with the same mean and variance. If both sets of percentiles come from the same distribution, the plotted points should line up along the 45-degree line. Figure 4 shows that this is the case—the scatter points roughly follow the 45-degree line across the entire distribution. Near zero the percentiles fall above the 45-degree line which indicates the actual distribution of inflation rates has somewhat less mass near zero than would be implied by a log-normal distribution. In the early period the percentiles eventually fall slightly below the 45-degree line which implies the tails are slightly less fat than a log-normal distribution. In the later period the points are on or above the 45-degree line which implies the tails of the distribution are at least as fat as a log-normal. We perform the same comparison in Appendix C.1 to a normal distribution and show the inflation distribution has fatter tails than a comparable normal distribution.

Third, we show that some PCE categories explain a disproportionately large share in variation of headline PCE inflation. We accomplish this goal by decomposing inflation into two components, an equal weighted average and a "granular" component as in Foerster et al. (2011). The equal weighted component indicates what inflation would be in the absence of granularity in consumer spending. The granular component is the residual of actual inflation and the equal weighted component. It will be large if a few large categories are subject to large inflation rates. If the direction of inflationary shocks is uncorrelated with the size of

Figure 3: Gini Plot of Contribution to Overall Inflation

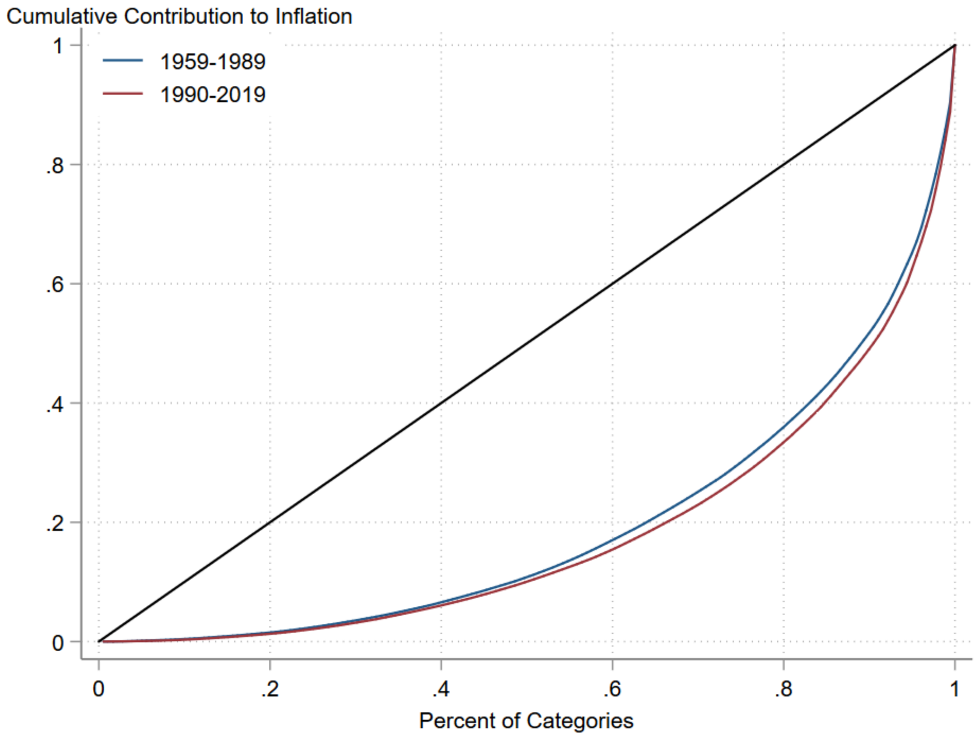

Notes: The figure plots the share of inflation (y-axis) accounted for by the smallest X percent of PCE categories. Categories are ranked from smallest to largest average contribution for each time period. Shifts to the right indicate a small number of categories are more important. The contribution of each series is calculated according to equation 2.

series then the granular component should be zero on average. Following the notation in equation 1 the decomposition is

$$\pi_t = \overbrace{\sum_{i=1}^{N} \frac{1}{N} \frac{p_{it}}{p_{it-12}}}^{\text{Equal Weights}} + \underbrace{\sum_{i=1}^{N} \left(w_{it} - \frac{1}{N} \right) \frac{p_{it}}{p_{it-12}}}_{\text{Granular Residual}} . \tag{4}$$

The granular residual component is positive in our sample, with an average of 0.33 and a standard deviation 0.47. The equal-weighted component is also positive in our sample, with an average of 2.98 and a standard deviation 2.42.

The importance of granularity for inflation has increased over time according to our various measures of granularity.

Figure 4: Comparison of PCEPI Inflation Percentiles to Log Normal Distribution

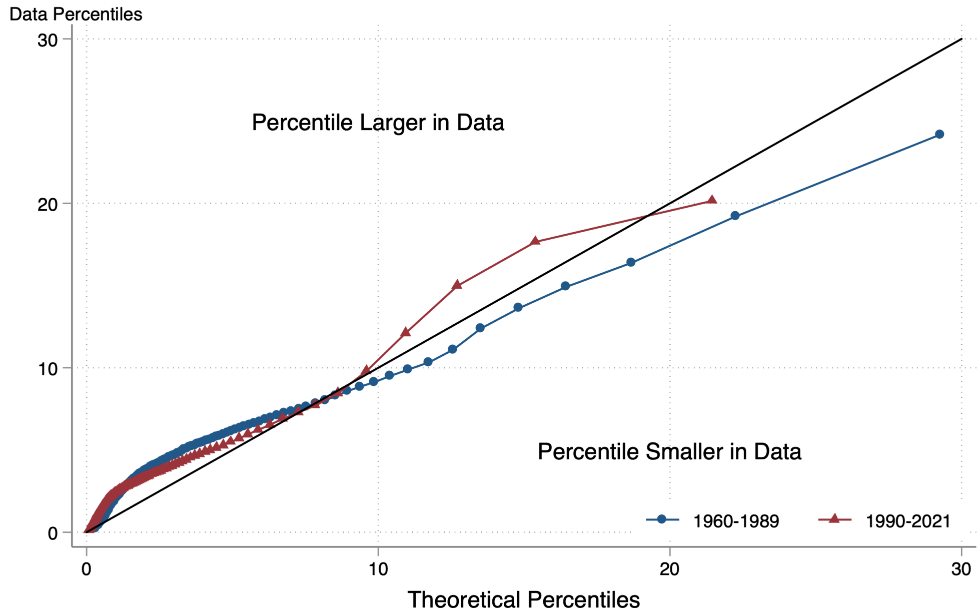

Notes: Authors' calculations from the PCE. The figure plots the percentiles of the observed cross-sectional distribution of inflation changes averaged over time, calculated as $|\pi_{it}|$. Each percentile is computed in each month and then averaged over the relevant time period, either 1960-1989 or 1990-2019. These values are plotted against the percentiles of a log normal distribution with the same mean and standard deviation as the inflation distribution in each time period (x-axis) which imply log-normal shape parameters of 1.09 and 1.09 for 1960-1989 and 0.31 and 0.31 for 1990-2019.

Fact 2 *Inflation granularity has increased:*

- *The top 10 out of 183 personal consumption expenditure (PCE) categories account for 35 percent of inflation in 1960, increasing to 41 percent by 2019. The top 15 categories account for half of inflation in 2019.*

- *The granular residual has become an equally important contributor to aggregate inflation post-1990 while equal-weighted contributions account for the bulk of inflation pre-1990.*

- *The fat tails of category inflation rates have flipped sides over time.*

First, the contribution of a few PCE categories to headline inflation has increased over time. The Gini Coefficient for inflation contribution increases from .663 in 1960 to .698 in 2019. The top 10 out of 183 personal consumption expenditure (PCE) categories account for 35 percent of PCE inflation in 1960, and for 41 percent by 2019. The top 15 categories account for 50 percent of inflation.

Second, we find that the granular residual as defined in equation (4) has become substantially more important over time. Figure 5 shows this result. The figure plots the equal-weighted and granular-residual components of inflation using a simple average during a three-month rolling window. Two observations stand out: On the one hand, the granular residual has become an equally important contributor to aggregate inflation post-1990. In particular, in the last 10 years, it has sometimes even exceeded the contribution of the equal-weighted component. This observation, on the other hand, contrasts with the pre-1990 experience. During that period, the equal-weighted contributions account by far for the majority of inflation.

Third, granularity as defined by fat tails of the category inflation distribution has also changed over time. As our discussion of the overall shape of the distribution pointed out, the skewness of the distribution exhibits a shift from right to left skew over time. This shift is also relevant for the nature of fat tails. Figure C.1 in the appendix illustrates this implication of the shift by considering a Q-Q normal plot in the pre-1990 and post-1990 periods. Clearly, the early period is characterized by fat tails on the positive side, and the more recent period by fat tails on the negative side. The fat tails of the distribution change location over time.

Figure 5: Inflation Decomposition - Y/Y Inflation

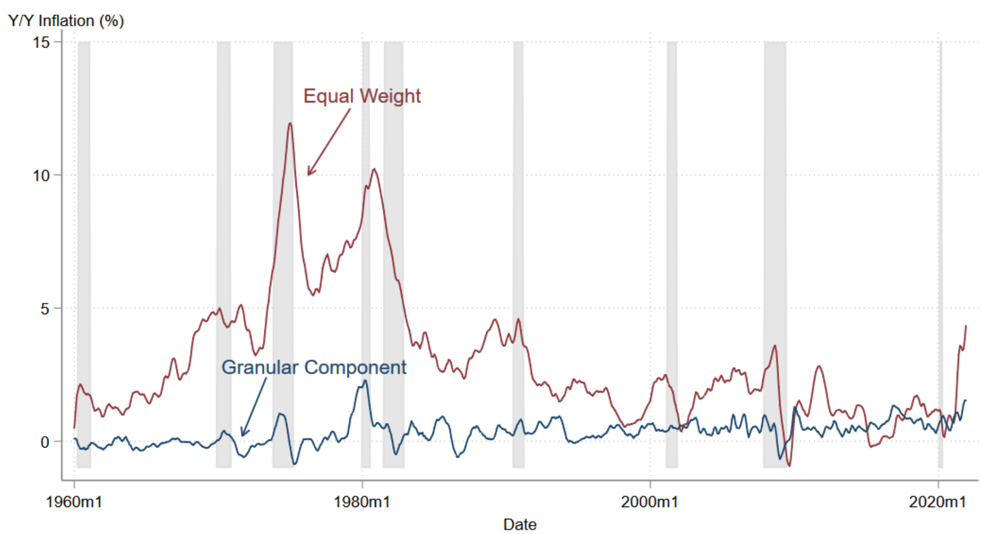

Notes: The figure plots the equal weight and granular components of inflation as defined in equation 4. The numbers are authors' calculations from the PCEPI using 12 month inflation rates. Numbers are calculated monthly and then smoothed via a simple average of a rolling 3 month window. Shaded areas indicate recessions as defined by the NBER.

3.3 Ranking of Robust Measures and Headline Inflation Changes

We argue that considering the entire distribution of highly disaggregated inflationary movements can improve our understanding of headline inflation. In particular, a few observations matter disproportionately in this context. We study more comprehensively the granularity of the inflation distribution both in the cross section and the time series in the next section, as well as its connection to volatility of headline inflation.

We are not the first to recognize that particular observations in the inflation distribution may matter for understanding the first moments of headline inflation. In fact, central banks and statistical agencies around the world have long published measures of headline inflation based on omitting certain categories. By omitting outliers, such measures of inflation can provide a better signal of the underlying inflation trend than measures that comprehensively include all categories. For example, the Federal Reserve Banks of Cleveland and Dallas publish measures of inflation such as median and trimmed mean inflation rates to provide more signals of what is happening with inflation movements. The performance of these measures has been explored in Ocampo, Schoenle and Smith (2022). Work by Rich et al. (2022) studies how to improve forecasting properties of robust measures.

Our analysis adds two insights to these existing efforts: First, one may believe that over short time periods, comprehensive, average-based inflation measures and robust measures typically tell similar, albeit not the same stories about the level of inflation. However, over longer time horizons there are important changes in the distribution that matter for the calculation of first moments: We find that the relationship between average and robust measures of headline inflation has changed. For example, the ranking between mean and median inflation has flipped indicating that there has been a change in the skewness of the inflation distribution. Figure 6 shows that before 1990 mean inflation exceeded median inflation by one percentage point on average with occasional periods where it was more than 2 percentage points higher. Since 1990 median inflation has typically been higher, especially since 2010. On average median inflation exceeds mean inflation by about half a percentage point. This change indicates that inflation has moved towards a relatively less positive skewness over the last 60 years. In fact, when we compute skewness in the data, skewness of category inflation rates is 3.90 pre-1990 and 0.59 post-1990. At the same time, kurtosis increases from 90.65 to 31.49 in these periods, also indicating relatively more mass in the tails. Non-parametric skewness, the difference between the mean and median normalized by the standard deviation, has on average turned negative post 1990, as Figure 6 illustrates. These findings are in line with other research showing the distribution of inflation rates in the U.S. is skewed (Bryan, Cecchetti and Wiggins, 1997) and that the skewness changes over time (Carroll and Verbrugge, 2019; Rich, Verbrugge and Zaman, 2022; Verbrugge and

Zaman, 2022). They are also in line with Luo and Villar (2021) who find that skewness and inflation have a positive correlation in the U.S. Consumer Price Index.

Figure 6: Mean vs Median Inflation

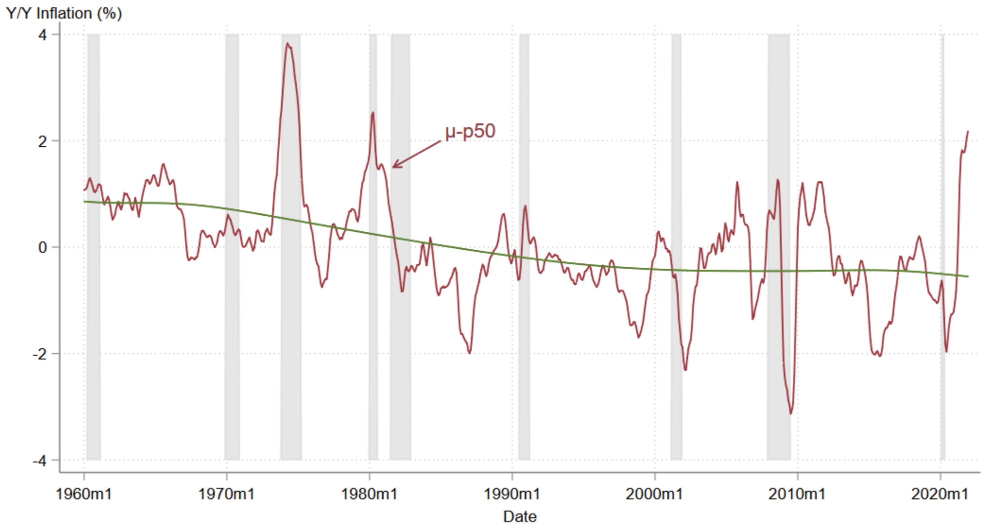

Notes: Authors' calculations from the PCEPI data. The figure plots the difference between inflation calculated as a weighted mean of individual categories and the median inflation category as calculated using the methodology of the Federal Reserve Bank of Cleveland (Federal Reserve Bank of Cleveland, 2021). In this methodology the median inflation category is selected for each month and then the rates for each category in the previous 12 months are combined to obtain an annual rate. The difference is smoothed using a 3-month rolling average. The trendline is calculated as a smoothed average using a Lowess regression (bandwidth = 0.8). Shaded areas indicate recessions as defined by the NBER. A plot of mean and median inflation levels is available in Figure B.1.

3.4 Covariance Decreases More than Variance

We next examine the evolution of second moments in the distribution of inflation rates. This focus leads us to calculate not only the variance of headline inflation, but also decompose it into the part that comes from the variance of individual components and the part that comes from the covariance between the series that make up the distribution.

Our analysis documents several pronounced trends for the second moments of inflation:

Fact 3 *Variance and covariance decrease over time:*

- *A decrease in the variance of individual series leads to a decrease of the variance of mean inflation over time.*

- *The average covariance and variance of inflation series decline substantially.*

The first trend we thus uncover relates to the variance of mean inflation: Over time, the variance of headline inflation has decreased, mostly due to a decrease in the variance of individual series. Second, we find that the average of covariances between the components of the PCEPI has decreased over time. We calculate the covariance between each pair of series in the PCEPI. Figure 7 plots this covariance, the average of variances of the series, and the average of the absolute values of the covariances of all pairs of series.

The average of variances and covariances decreases over time indicating that on average inflation is less volatile over time. The average of the variances falls about 78 percent from a peak at 55.71 in the 1970s to 12.37 at the end of 2019, and the average of covariances falls over 97 percent from a peak of 8.32 around 1975 to 0.28 by the end of 2019. This result is in line with the experience of high, volatile inflation in the 1970 and 1980s, and low and stable rates. The decrease in average covariance could be due to declines in positive covariances or an increase in the number of negative covariance terms. If the decline in the average covariance was caused by more negative covariance terms, then the average absolute covariance should not fall and could even increase. We find that this is not the case. The average absolute covariance series falls 89 percent from its peak of 14.38 to 1.51 in 2019. The 88 percent fall is smaller than the 98 percent fall in the average covariance term which implies that some covariance terms have shifted to be slightly more negative, but these results also indicate that covariance terms are closer to zero in general.

Figure 7: Variance and Covariance of Inflation Series

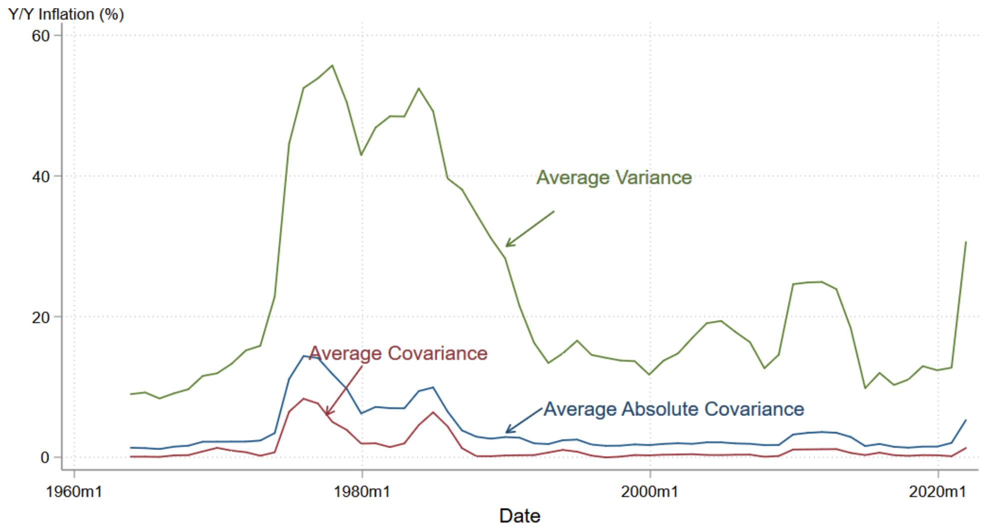

Notes: Authors' calculations from the PCEPI. Average variance is an unweighted average of the variance of each inflation series over the previous 60 months. Average covariance is the average pairwise covariance of inflation series over the previous 60 months. Average absolute covariance is the average of the absolute value of the covariance terms over the previous 60 months.

4 Model

This section presents a heterogeneous production model with idiosyncratic shocks at the sector level, as well as aggregate shocks. The model features heterogeneity in sectoral input-output linkages, degrees of nominal price rigidity and importance in the consumption basket. We show that a model version with idiosyncratic shocks can rationalize the preceding empirical findings. It also provides new insights into the stabilization properties of monetary policy absent in conventional one-sector models.

4.1 Households

The representative household in this setup solves

$$\max_{\{C_t, L_{kt}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\sigma} - 1}{1-\sigma} - \sum_{k=1}^{K} g_k \frac{L_{kt}^{1+\varphi}}{1+\varphi} \right),$$

subject to

$$\begin{aligned}\sum_{k=1}^{K} W_{kt} L_{kt} + \sum_{k=1}^{K} \Pi_{kt} + I_{t-1} B_{t-1} - B_t &= P_t^c C_t \\ \sum_{k=1}^{K} L_{kt} &\leq 1,\end{aligned}$$

In this closed economy model with no investment and no government spending, C_t can be interpreted either as consumption or GDP and P_t^c as the consumer price index or the GDP deflator. L_{kt} and W_{kt} are labor employed and wages paid in sector $k = 1, ..., K$. Households own firms and receive net income, Π_{kt}, as dividends. Bonds, B_{t-1}, pay a nominal gross interest rate of I_{t-1}. Total labor supply is normalized to 1.

C_t aggregates from sectoral GDP, C_{kt}, and in turn from households' final demand for each good, C_{jkt}, according to

$$C_t \equiv \left[\sum_{k=1}^{K} \omega_{ck}^{\frac{1}{\eta}} C_{kt}^{1-\frac{1}{\eta}}\right]^{\frac{\eta}{\eta-1}}, \tag{5}$$

$$C_{kt} \equiv \left[n_k^{-1/\theta} \int_{\Im_k} C_{jkt}^{1-\frac{1}{\theta}} \, dj\right]^{\frac{\theta}{\theta-1}}. \tag{6}$$

A continuum of goods indexed by $j \in [0, 1]$ exists with total measure 1. Each good belongs to one of the K sectors in the economy. Mathematically, the set of goods is partitioned into K subsets $\{\Im_k\}_{k=1}^{K}$ with associated measures $\{n_k\}_{k=1}^{K}$ such that $\sum_{k=1}^{K} n_k = 1$.[3] We allow the elasticity of substitution across sectors η to differ from the elasticity of substitution within sectors θ.

The first key ingredient of our model is the vector of weights $\Omega_c \equiv [\omega_{c1}, ..., \omega_{cK}]$ in equation (5). Households' sectoral demand

$$C_{kt} = \omega_{ck} \left(\frac{P_{kt}}{P_t^c}\right)^{-\eta} C_t \tag{7}$$

determines the interpretation as sectoral GDP shares as in steady state, when all prices are identical, $\omega_{ck} \equiv \frac{C_k}{C}$ (variables without a time subscript denote steady-state levels.) Thus, the vector Ω_c satisfies $\Omega_c' \iota = 1$, where ι denotes a column-vector of 1s. Away from steady state, sectoral GDP shares depend on sectoral prices relative to the aggregate price index,

$$P_t^c = \left[\sum_{k=1}^{K} \omega_{ck} P_{kt}^{1-\eta}\right]^{\frac{1}{1-\eta}}. \tag{8}$$

[3]The sectoral subindex is redundant, but it clarifies exposition. We can interpret n_k as sector size measured or gross output share.

Household demand for goods within a sector is given by

$$C_{jkt} = \frac{1}{n_k}\left(\frac{P_{jkt}}{P_{kt}}\right)^{-\theta} C_{kt} \text{ for } k = 1, ..., K. \tag{9}$$

Firms within a sector equally share the production of goods in steady state. Away from steady state, the gap between a firm's price, P_{jkt}, and the sectoral price, P_{kt}, distorts the demand for goods within a sector.

Sector k's price is defined as

$$P_{kt} = \left[\frac{1}{n_k}\int_{\Im_k} P_{jkt}^{1-\theta} dj\right]^{\frac{1}{1-\theta}} \text{ for } k = 1, ..., K. \tag{10}$$

The household first-order conditions determine labor supply and the Euler equation

$$\frac{W_{kt}}{P_t^c} = g_k L_{kt}^{\varphi} C_t^{\sigma} \text{ for all } k, j, \tag{11}$$

$$1 = \mathbb{E}_t\left[\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\sigma} I_t \frac{P_t^c}{P_{t+1}^c}\right]. \tag{12}$$

We implicitly assume sectoral segmentation of labor markets, so labor supply in equation (11) holds for a sector-specific wage $\{W_{kt}\}_{k=1}^{K}$. We choose the parameters $\{g_k\}_{k=1}^{K}$ to ensure a symmetric steady state across all firms.

4.2 Firms

A continuum of monopolistically competitive firms exists, each producing a single good. To facilitate exposition, firms are indexed by the good $j \in [0, 1]$ they produce and the sector, $k = 1, ...K$ they belong to. The production function is

$$Y_{jkt} = e^{a_{kt}} L_{jkt}^{1-\delta} Z_{jkt}^{\delta}, \tag{13}$$

where a_{kt} is an i.i.d. productivity shock to sector k with $\mathbb{E}[a_{kt}] = 0$ and $\mathbb{V}[a_{kt}] = v^2$ for all k, L_{jkt} is labor, and Z_{jkt} is an aggregator of intermediate inputs

$$Z_{jkt} \equiv \left[\sum_{k'=1}^{K} \omega_{kk'}^{\frac{1}{\eta}} Z_{jk}(k')^{1-\frac{1}{\eta}}\right]^{\frac{\eta}{\eta-1}}. \tag{14}$$

$Z_{kjt}(k')$ is the amount of goods firm jk demands as inputs at time t from sector k'.

The second key ingredient of our model is heterogeneity in aggregator weights $\{\omega_{kk'}\}_{k,k'}$. We denote these weights in matrix notation as Ω, satisfying $\Omega\iota = \iota$. The demand of firm jk for goods produced in sector k' is given by

$$Z_{jkt}(k') = \omega_{kk'} \left(\frac{P_{k't}}{P_t^k}\right)^{-\eta} Z_{jkt}. \tag{15}$$

We interpret $\omega_{kk'}$ as the steady-state share of goods from sector k' in the intermediate input use of sector k, which determines the input-output linkages across sectors in steady state. Away from the steady state, the gap between the price of goods in sector k' and the aggregate price relevant for a firm in sector k, P_t^k, distorts input-output linkages

$$P_t^k = \left[\sum_{k'=1}^{K} \omega_{kk'} P_{k't}^{1-\eta}\right]^{\frac{1}{1-\eta}} \quad \text{for } k = 1, ..., K. \tag{16}$$

P_t^k uses the sector-specific steady-state input-output linkages to aggregate sectoral prices.

The aggregator $Z_{jk}(k')$ gives the demand of firm jk for goods produced in sector k'

$$Z_{jk}(k') \equiv \left[n_{k'}^{-1/\theta} \int_{\Im_{k'}} Z_{jkt}(j', k')^{1-\frac{1}{\theta}} \, dj'\right]^{\frac{\theta}{\theta-1}}. \tag{17}$$

Firm jk's demand for an arbitrary good j' from sector k' is

$$Z_{jkt}(j', k') = \frac{1}{n_{k'}} \left(\frac{P_{j''k't}}{P_{k't}}\right)^{-\theta} Z_{jk}(k'). \tag{18}$$

In steady state, all firms within a sector share the intermediate input demand of other sectors equally. Away from steady state, the gap between a firm's price and the sectoral price index distorts the firm's share in the production of intermediate inputs. Our economy has $K + 1$ different aggregate prices, one for the household sector and one for each of the K sectors. By contrast, the household sector and all sectors face unique sectoral prices.

The third key ingredient of our model is sectoral heterogeneity in price rigidity. For quantitative purposes, we model price rigidity a la Calvo with parameters $\{\alpha_k\}_{k=1}^K$ such that the pricing problem of firm jk is

$$\max_{P_{jkt}} \mathbb{E}_t \sum_{s=0}^{\infty} Q_{t,t+s} \alpha_k^{\tau} \left[P_{jkt} Y_{jkt+s} - MC_{kt+s} Y_{jkt+s}\right]. \tag{19}$$

Marginal costs are $MC_{kt} = \frac{1}{1-\delta} \left(\frac{\delta}{1-\delta}\right)^{-\delta} e^{-a_{kt}} W_{kt}^{1-\delta} \left(P_t^k\right)^{\delta}$ after imposing the optimal mix of labor and intermediate inputs

$$\delta W_{kt} L_{jkt} = (1-\delta) P_t^k Z_{jkt}, \tag{20}$$

and $Q_{t,t+s}$ is the stochastic discount factor between periods t and $t+s$.We assume the elasticities of substitution across and within sectors are the same for households and all firms. This assumption shuts down price discrimination across different customers, and firms choose a single price P_{kt}^*

$$\sum_{\tau=0}^{\infty} Q_{t,t+\tau} \alpha_k^s Y_{jkt+\tau} \left[P_{kt}^* - \frac{\theta}{\theta-1} MC_{kt+\tau} \right] = 0, \tag{21}$$

where $Y_{jkt+\tau}$ is the total production of firm jk in period $t+\tau$.

We define idiosyncratic shocks $\{a_{kt}\}_{k=1}^K$ at the sector level, and it follows the optimal price, P_{kt}^*, is the same for all firms in a given sector. Thus, aggregating among all prices within sector yields

$$P_{kt} = \left[(1-\alpha_k) P_{kt}^{*1-\theta} + \alpha_k P_{kt-1}^{1-\theta} \right]^{\frac{1}{1-\theta}} \text{ for } k = 1, ..., K. \tag{22}$$

4.3 Monetary Policy, and Equilibrium Conditions

The choice of monetary policy is crucial for the effects of shocks on the distribution and dynamics of inflation rates in the economy. Our baseline monetary policy follows a Taylor-type rule which can be represented as

$$I_t = \frac{1}{\beta} \left[\left(\frac{P_t^c}{P_{t-1}^c} \right)^{\phi_\pi} \left(\frac{C_t}{\phi^* C_t^* + (1-\phi^*) C} \right)^{\phi_y} \left(\frac{C_t}{C_{t-1}} \right)^{\phi_{gc}} \right], \tag{23}$$

with a degree of monetary policy smoothing of ρ_i, a systematic response ϕ_π to inflation, P_t^c / P_{t-1}^c, a systematic response ϕ_y to deviations of GDP from a weighted average between the frictionless and steady-state GDP levels, $C_t / [(\phi^* C_t^* + (1-\phi^*) C]$, and a systematic response ϕ_{gc} to GDP growth.

To further gauge the stabilizing effects of different choices of monetary policy rules for inflation and in particular, the distribution of inflation rates, we consider the following alternative specifications of monetary policy: First, we consider a specification which focuses on average inflation targeting, AIT, that targets an equal-weighted average of six lags of headline inflation:

$$I_t = \frac{1}{\beta} \left[\left(\prod_{j=0}^{j=6} \frac{P_{t-j}^c}{P_{t-j-1}^c} \right)^{\frac{\phi_\pi}{6}} \right], \tag{24}$$

Second, we also allow for a response of monetary policy to the shock that is lagging one period behind: That is, policy does not react to a contemporaneous inflation vector, but a one-period lag of it. This specification is meant to capture a scenario where a policymaker reacts with delay to idiosyncratic shocks.

To focus on the interaction of monetary policy rules and inflation measures. When shocks are idiosyncratic, we do not only consider CPI inflation, but also "complementary inflation" denoted by $P_t^{comp}/P_{t-1}^{comp}$. This measure of inflation is defined by the inflation rates in the complement of sectors hit with idiosyncratic shocks. For example, if there are idiosyncratic shocks to food and gas prices only, complementary inflation is essentially given by core inflation. We also consider various sets of idiosyncratic shocks discussed below in the model analysis.

Finally, bonds are in zero net supply, $B_t = 0$, labor markets clear, and goods markets clear such that

$$Y_{jkt} = C_{jkt} + \sum_{k'=1}^{K} \int_{\Im_{k'}} Z_{j'k't}\left(j,k\right) dj', \tag{25}$$

implying a wedge between gross output Y_t and GDP C_t.

5 Model Analysis

This section presents two theoretical analyses. The first shows that a setup with idiosyncratic shocks is needed to match the empirical facts. The second highlights the implications of the choice of the monetary policy rule and target for the distribution of inflation rates. We set model parameters for the U.S. economy following the parameter choices in Pasten et al. (2020). These parameters include the persistence of shocks, the discount factor, the strength of input-output linkages, pricing frictions, sectoral consumption shares and elasticities of substitution. Table D.1 in the appendix summarizes the parameter choices.

5.1 The Importance of Idiosyncratic Shocks

In the first analysis, we use six different model calibrations to establish that idiosyncratic shocks are essential to match our empirical findings. Each of these calibrations represents a combination of an aggregate productivity shock and/or idiosyncratic productivity shocks that hit either one, all, or a subset of sectors. The first such calibration only features an aggregate productivity shock; the second calibration 341 independent unit-productivity shock simultaneously hitting each sector; the third calibration 341 unit-productivity shocks hitting only one sector at a time; the fourth calibration 341 combinations of a unit aggregate shock and a unit sector shock at a time; the fifth calibration a unit aggregate shock and 341 simultaneous, independent unit sector shocks; the sixth calibration a unit shock that hits the oil-producing sectors in the US (NAICS codes: 211000 (Oil and Gas Extraction, 213111 (Drilling Oil and Gas Wells), 324110 (Petroleum Refineries), 325110 (Petrochemical Manufacturing), and 325120 (Industrial Gas Manufacturing)) and the seventh calibration a unit aggregate shock and a unit shock that hits the oil-producing sectors in the US.

To gauge if a calibration is able to match our facts, we compute the following statistics using a simulation of the model. We first compute whether there are fat tails on the negative side of the distribution of inflation rates. Specifically, we conclude that the distribution has fat left tails if the first and fifth percentile are smaller than that of a normal distribution with the same standard deviation of inflation rates observed in the data. Second, we analogously compute whether there are fat tails on the positive side of the distribution of inflation rates using the 95th and 99th percentile. Third, as a derivative of these calculations, we conclude that a calibration delivers fat tails in general if both negative and positive fat tails are present. Fourth, we conclude that the ranking of headline inflation and median inflation reverses if both a month exists whether the headline rate exceeds the median inflation rate and a month exists where the headline rate is less than the median rate. Fifth, we compute the granular residual as defined in equation 4. If the granular residual is ever non-zero, we conclude that the model features granularity in inflation rates. Sixth, we conclude the granular residual is important if its variance is at least 20 percent of the variance of headline inflation. Finally, we simulate the model twice with different size shocks and determine whether the change in average sectoral variance exceeds the changes in average sectoral covariance.

A clear message emerges when we compare these model-simulated statistics and our empirical findings: Idiosyncratic shocks are necessary to match the empirical cross-section and time-series properties of the distribution of inflation rates we have uncovered. Table 2 summarizes how we arrive at this conclusion: Whenever we match an empirical fact (given by a separate row). We indicate such a conclusion in two ways: Either an “X” marks a model success, or we make an addition to the count of successes when there are 341 possible sub-calibrations. Comparing across columns which embody the various calibrations described above, it is evident that only calibrations with idiosyncratic shocks match all facts with seven “X”s, or a large number of successes (Columns 2, 3 and 6).

More specifically, we see the following patterns: A calibration with only aggregate shocks (Column 1) overall fails both along cross sectional and time series dimensions. On the positive side, even an aggregate shock generates reversals of mean and median aggregate inflation rates, as well as a role for granularity. These findings originate due to the heterogeneity in the production structure, such as differences in pricing frictions that generate a distribution of inflation rates with the appropriate dynamics (while not the question

Table 2: Matching Facts

	(1) Agg	(2) Idio all	(3) Idio 1-ind	(4) Agg + 1-ind	(5) Agg + all	(6) Oil	(7) Agg + oil
Negative Fat Tails		X	309			X	X
Positive Fat Tails		X	309			X	X
Fat Tails		X	309			X	X
Mean and Median Flip	X	X	341	341	X	X	X
Granular Residual	X	X	341	341	X	X	X
Important Granular Residual	X	X	191	341	X	X	X
Larger Cov than Var Drop		X	158			X	

Notes: Results are authors' calculations. Standard deviations are taken over the 12 months following the shock. The inflation impact is the change in inflation in the month of the shock and is expressed as a yearly rate.

at hand, we have checked that a homogeneous production model fails at generating these facts). However, in the cross section, the pure aggregate shock calibration does not generate fat tails on either side. This failure arises despite the presence of heterogeneity of input-output linkages in the model which in principle might amplify the effects of the aggregate shock asymmetrically for some sectors, leading to asymmetric price responses. In the time series dimension, to match the relatively larger drop in covariance relative to average variance, aggregate and hence perfectly correlated shocks are, not surprisingly, also not sufficient in generating a differential movement in second moments.

This pattern of failure of aggregate shocks stands in stark contrast with the success of idiosyncratic shocks in Columns 2, 3, and 6: When idiosyncratic shocks independently affect sectoral fluctuations (Column 2), the model can match all empirical facts. Idiosyncratic shocks by nature have asymmetric effects and in this heterogeneous model economy, calibrated to the U.S., they generate both fat tails and variation in the variance-covariance matrix over time. Even if only one idiosyncratic shock at a time hits the economy (Column 3), many of the 341 possible calibrations generate statistics in line with the facts. Naturally, not all calibrations succeed at generating patterns observed in the data. While idiosyncratic shocks can have an asymmetric effect in our setup, this exercise demonstrates that they need not to, depending on how large a sector is, how flexible its prices are and how connected it is to other sectors. However, as Column 6 illustrates, a specific sector in our empirical calibration to the US economy succeeds at matching all facts: when the oil-producing sectors are subject to a common unit-productivity shock, the properties that characterize them in the U.S. economy lead to a response of the distribution of inflation rates in line with the facts we observe.

A final insight derives from the calibrations that mix an aggregate shock with idiosyncratic shocks. As

Columns 4, 5 and 7 show, the predictions of an aggregate shock rather than the predictions of idiosyncratic shocks generally dominate. The intuitive reason is that an aggregate, unit-productivity shock is much more impactful relative to 341 independent unit-productivity shocks that average out to some extent, or relative to a single idiosyncratic, unit-productivity shock. This insight re-enforces the findings of the literature such as Pasten et al. (2021) that we need relatively large idiosyncratic shocks to generate an effect on inflation.

5.2 Monetary Policy and the Distribution of Inflation Rates

Based on the above results, a second part of the analysis builds on a multi-sector model with idiosyncratic shocks. We illustrate in our calibrated model for the U.S. economy what happens to inflation and inflation stabilization when a policymaker faces fluctuations driven by idiosyncratic shocks. We consider the implications of the policymaker following either a standard Taylor Rule or average inflation targeting (AIT). We also consider three different measures of inflation that the policymaker can target. Headline inflation is the standard measure of inflation, core inflation which strips out food and energy sectors, and complementary inflation which targets inflation in the sectors unaffected by the shock.

Both average inflation targeting and targeting either core or complementary inflation generally stabilize the movements in inflation and compress the distribution of rates. Moreover, the stabilization gains from these policy rules interact, leading to further gains when both are used. We show that these results apply in particular when a shock to the oil-producing industries hits the economy and when shocks occur to broad sectors of the economy.

Our first results come from shocking the five oil-producing sectors with a negative productivity shock that is calibrated to produce a one percent increase in headline inflation under a Taylor rule targeting headline inflation. Then we consider how the path and variability of inflation and consumption change if instead the central bank uses average inflation targeting or targets a different measure of inflation. As a result, the analysis considers six different candidate combinations of two monetary policy rules and three inflation targets.

Table 3 shows the results of these exercises. Each row represents one of the calibrations. The first and second columns show the inflation response on impact and after 12 months under the six scenarios. Next, we show the standard deviation of inflation. Finally, we show the standard deviation of consumption. Both standard deviations are computed over the first year and make up the heart of conventional loss functions used in welfare analysis.

New insights into the stabilization properties of monetary policy arise regarding the strategies central

Table 3: Impacts of Different Monetary Policy Rules - Oil Shock

Inflation Measure	Policy Rule	π Impact	π 12-month	$\sigma(\pi)$	$\sigma(C)$
Headline	Taylor	1.00	0.39	0.0133	0.000931
Headline	AIT	0.44	0.23	0.0134	0.000719
Core	Taylor	0.58	0.37	0.0134	0.000981
Core	AIT	0.35	0.22	0.0135	0.000645
Complementary	Taylor	0.99	0.38	0.0123	0.000902
Complementary	AIT	0.44	0.22	0.0121	0.000700

Notes: Results are authors' calculations. Standard deviations are taken over the 12 months following the shock. π impact is the change in inflation in the month of the shock and is expressed as a yearly rate. π 12-month is the inflation rate 12 months after the shock expressed at a yearly rate.

bankers may choose when they face a cycle driven by select idiosyncratic shocks. We find three results in our calibrated economy. These results hold when considering a shock to the oil sector, but also more generally by considering one idiosyncratic at a time and a policy that reacts to complementary inflation for each shock, as outlined above.

First, the analysis shows that AIT stabilizes the distribution more than a Taylor-type policy rule. As the top panel of Table 3 indicates for an oil price shock, moving from a Taylor-type rule to AIT reduces inflation volatility overall by about a factor of one-half while keeping inflation fluctuations stable. This finding is in line with the well-known stabilizing properties of average inflation targeting. By considering a multi-sector model, however, we can also locate the incidence of the policy. We find that AIT works not only by generating lower levels of inflation, but by compressing the center of the distribution more than a Taylor-type rule would. Figure 8 illustrates this insight. Additionally, Table 4 indicates, these conclusions generally hold for sectoral shocks.

Second, we find that stabilization of core inflation—ignoring the shocks to the oil-producing sectors and only reacting to their complement—achieves most of the benefits for inflation stabilization as AIT, while leaving consumption volatility unchanged. Again, as Table 4 indicates, these conclusions may generalize to more broad sets of shocks shock.

Third, our analysis suggests that stabilization of core inflation—while policy simultaneously shifts to an AIT regime—creates additional inflation stabilization benefits in the case of an oil-price shock. The fourth row of Table 3 shows this result. But in fact, policy can do even better by precisely targeting complementary inflation and switching to an AIT regime. As the last row in Table 3 indicates, this combination of complementary inflation and an AIT regime generates the lowest impact, the lowest 12-month response and

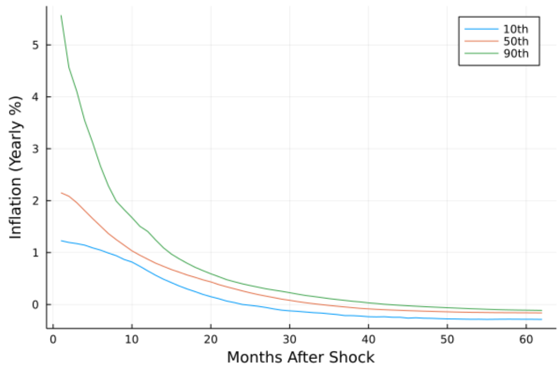

(a) Taylor Rule

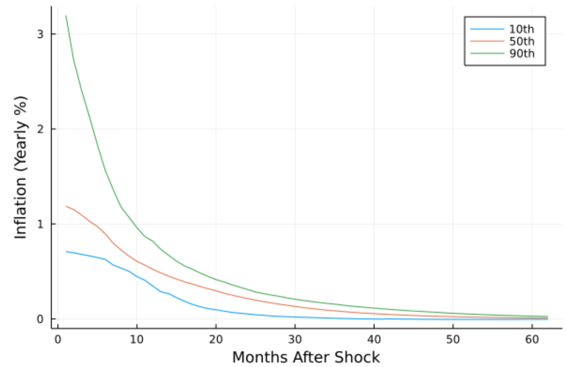

(b) AIT

Figure 8: Distribution of Inflation Rates after Shock

the lowest inflation and output volatility. Table 4 confirms that this result generally holds for a policy that reacts to complementary inflation for any idiosyncratic shock, and that also switches to AIT. The fact that key welfare metrics improve relative to a Taylor rule and a size-weighted mean measure of headline inflation is not surprising given the limitations of a Taylor rule (see e.g. Woodford (2001) or Giannoni (2014)) and the further complications of optimal policy in multisector models (see e.g. Rubbo (2020) or Jennifer and Tahbaz-Salehi (2020)).

Table 4: Impacts of Different Monetary Policy Rules - Sectoral Shocks

Inflation Measure	Policy Rule	π Impact	π 12-month	$\sigma(\pi)$	$\sigma(C)$
Headline	Taylor	1.00	0.24	0.00135	0.000776
Headline	AIT	0.57	0.15	0.00063	0.000595
Complementary	Taylor	0.84	0.23	0.00134	0.000752
Complementary	AIT	0.52	0.18	0.00072	0.000776

Notes: Results are authors' calculations. Numbers are the average of three scenarios; shocks to agriculture and mining, shocks to manufacturing, and shocks to other industries. In each case the size of the shock is calibrated to generate a 1 percentage point increase in inflation with a standard Taylor Rule. Standard deviations are taken over the 12 months following the shock. π impact is the change in inflation in the month of the shock and is expressed as a yearly rate. π 12-month is the inflation rate 12 months after the shock expressed at a yearly rate.

6 Conclusion

Analyzing the distribution of disaggregated PCE category inflation rates, we detect systematic changes in the distribution between 1960 and 2021. Pre-1990, extreme positive tails characterize the distribution, but they moderate post-1990 while more negative tails appear. The distribution is granular, with an increasing importance of granularity over time. The ranking of mean inflation versus robust measures of inflation—medians and trimmed means—inverts several times for our times series. The covariance of disaggregated inflation rates decreases more than the variance over time.

These findings point to the use of multi-sector models when appropriately analyzing the stabilization properties of monetary policy. We show that a model with idiosyncratic shocks is necessary to match key features of the cross-sectional distribution of inflation rates. Then, we use oil price shocks and shocks to sectors of the economy to show how the choice of policy regime interacts with the distribution of inflation rates and the targeted measure of inflation. We find that average inflation targeting stabilizes more than a Taylor Rule. However, targeting core inflation with a Taylor Rule achieves many of the benefits of changing monetary policy regimes.

References

Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (2012) "The network origins of aggregate fluctuations," *Econometrica*, Vol. 80, No. 5, pp. 1977–2016.

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (2017) "Microeconomic Origins of Macroeconomic Tail Risks," *American Economic Review*, Vol. 107, No. 1, pp. 54–108.

Alvarez, Fernando, Hervé Le Bihan, and Francesco Lippi (2016) "The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach," *American Economic Review*, Vol. 106, No. 10, pp. 2817–51.

Baqaee, David Rezza and Emmanuel Farhi (2020) "Productivity and Misallocation in General Equilibrium," *The Quarterly Journal of Economics*, Vol. 135, No. 1, pp. 105–163.

Bhattarai, Saroj and Raphael Schoenle (2014) "Multiproduct firms and price-setting: Theory and evidence from U.S. producer prices," *Journal of Monetary Economics*, Vol. 66, pp. 178–192.

Bils, Mark and Peter Klenow (2004) "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, Vol. 112, No. 5, pp. 947–985.

Bonomo, Marco, Carlos Carvalho, Oleksiy Kryvtsov, Sigal Ribon, and Rodolfo Rigato (2020) "Multi-Product Pricing: Theory and Evidence from Large Retailers in Israel," staff working papers, Bank of Canada.

Bryan, Michael F, Stephen G Cecchetti, and Rodney L Wiggins (1997) "Efficient inflation estimation."

Bureau of Economic Analysis (2022) "Personal Consumpture Expenditure Data," https://apps.bea.gov/national/Release/XLS/Underlying/Section2All_xls.xlsx.

Carroll, Daniel R. and Randal J. Verbrugge (2019) "Behavior of a New Median PCE Measure: A Tale of Tails," *Economic Commentary (Federal Reserve Bank of Cleveland)*, No. 2019-10.

Federal Reserve Bank of Cleveland (2021) "Median CPI," technical report, Federal Reserve Bank of Cleveland.

Federal Reserve Bank of Dallas (2021) "Trimmed Mean PCE Inflation Rate [PCETRIM12M159SFFRBDAL]," https://fred.stlouisfed.org/series/PCETRIM12M159SFRBDAL.

Foerster, Andrew, Pierre Daniel Sarte, and Mark Watson (2011) "Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production," *Journal of Political Economy*, Vol. 119, No. 1, pp. 1 – 38.

Gabaix, Xavier (2011) "The Granular Origins of Aggregate Fluctuations," *Econometrica*, Vol. 79, No. 3, pp. 733–772.

Gagnon, Etienne (2009) "Price Setting during Low and High Inflation: Evidence from Mexico," *The Quarterly Journal of Economics*, Vol. 124, No. 3, pp. 1221–1263.

Giannoni, Marc P. (2014) "Optimal interest-rate rules and inflation stabilization versus price-level stabilization," *Journal of Economic Dynamics and Control*, Vol. 41, pp. 110–129.

Jennifer, La'O and Alireza Tahbaz-Salehi (2020) "Optimal monetary policy in production networks," NBER Working Papers w27464, National Bureau of Economic Research.

Karadi, Peter, Raphael Schoenle, and Jesse Wursten (2021) "Measuring price selection in microdata: it's not there," Working Paper Series 2566, European Central Bank.

Luo, Shaowen and Daniel Villar (2021) "The Skewness of the Price Change Distribution: A New Touchstone for Sticky Price Models," *Journal of Money, Credit and Banking*, Vol. 53, No. 1, pp. 41–72.

Midrigan, Virgiliu (2011) "Menu costs, multiproduct firms, and aggregate fluctuations," *Econometrica*, Vol. 79, No. 4, pp. 1139–1180.

Nakamura, E. and J. Steinsson (2008) "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, Vol. 123, No. 4, pp. 1415–1464.

Nakamura, Emi, Jón Steinsson, Patrick Sun, and Daniel Villar (2018) "The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation*," *The Quarterly Journal of Economics*, Vol. 133, No. 4, pp. 1933–1980.

Ocampo, Sergio, Raphael Schoenle, and Dominic A. Smith (2022) "How Robust are Robust Measures of PCE Inflation?", https://arxiv.org/abs/2207.12494.

Pasten, Ernesto, Raphael Schoenle, and Michael Weber (2020) "The propagation of monetary policy shocks in a heterogeneous production economy," *Journal of Monetary Economics*, Vol. 116, pp. 1–22.

Pasten, Ernesto, Raphael Schoenle, and Michael Weber (2021) "Heterogeneity in Nominal Price Rigidity and the Origin of Aggregate Fluctuations," Working Paper 2018-54, Friedman Institute for Economics.

Rich, Robert, Randal Verbrugge, and Saeed Zaman (2022) "Adjusting MEdian and Trimmed-Mean Inflation Rates for Bias Based on Skewness," https://www.clevelandfed.org/~/media/content/newsroom%20and%20events/publications/economic%20commentary/2022/ec%20202205/ec%20202205.pdf.

Rubbo, Elisa (2020) "Networks, Phillips Curves, and Monetary Policy," working paper, Harvard University, Department of Economics.

Vavra, Joseph (2014) "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *The Quarterly Journal of Economics*, Vol. 129, No. 1, pp. 215–258.

Verbrugge, Randal and Saeed Zaman (2022) "Improving Inflation Forecasts Using Robust Measures," https://fedinprint.org/item/fedcwq/94549/original.

Woodford, Michael (2001) "The Taylor Rule and Optimal Monetary Policy," *The American Economic Review*, Vol. 91, No. 2, pp. 232–237.

Appendices for Online Publication

Table of Contents

A Data Appendix

A.1 Matching Top Line Inflation

The BEA constructs their inflation estimates as Fischer indexes of multiple underlying components. The elementary components are not made available to data users. Instead, lower level Fischer indexes of these components are published. These Fischer indexes do not perfectly aggregate to the the top level inflation index. Figure A.1 shows plots of published year-over-year inflation against top line inflation calculated by combing the less detailed series using different index methods. Figure A.2 shows a box plot of the mean, interquartile range, and outliers of the difference between an aggregate index constructed according to each index number formula and the published aggregate index in the PCEPI.

Figure A.1: Comparison of Aggregation Methods

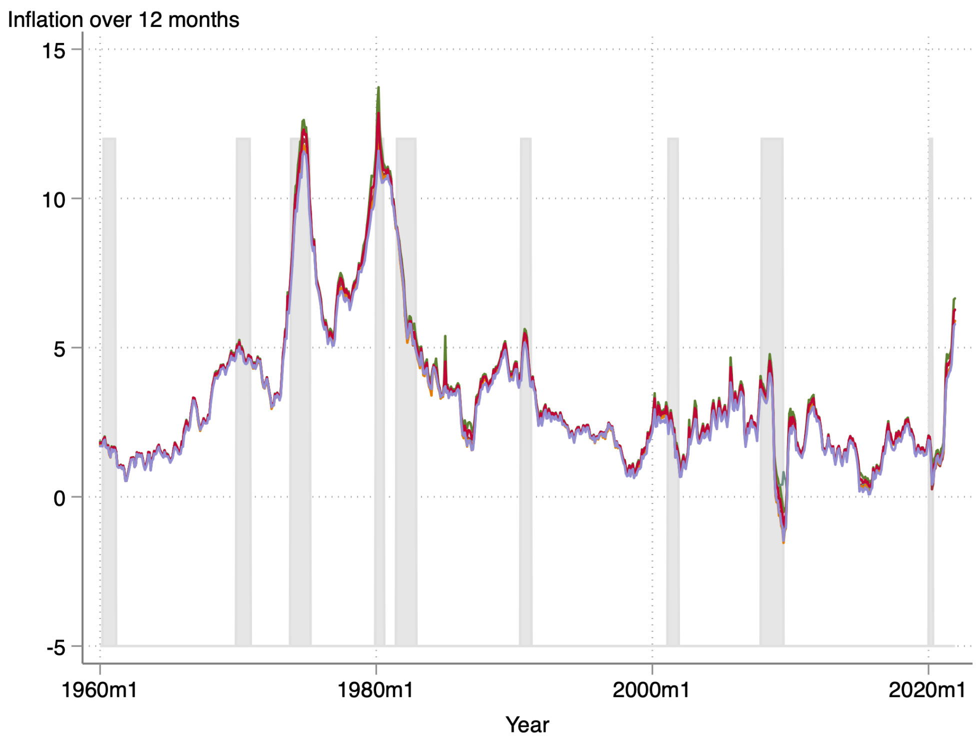

Notes: Plots year-over-year inflation from line one of the PCEPI against the highest level index implied by 183 series that partition consumer spending aggregated using a different methodology.

Figure A.2: Comparison of Aggregation Methods - Box Plot of Monthly Differences

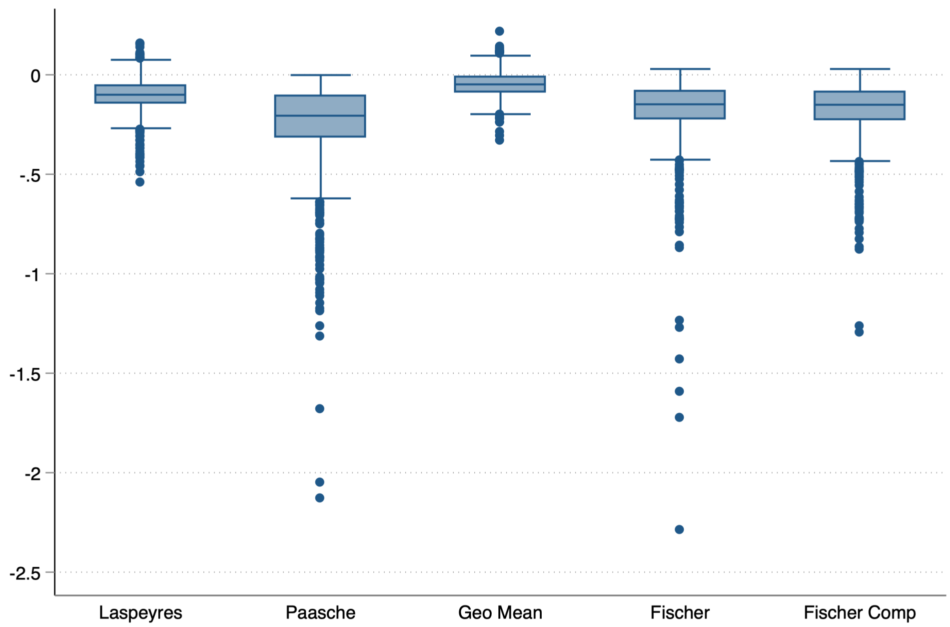

Notes: Plots y/y inflation from line one of the PCEPI against the highest level index implied by 183 series that partition consumer spending aggregated using a different methodology. The graph shows the mean, interquartile range and outliers of the difference between an aggregate index constructed according to each index number formula and the published aggregate index in the PCEPI.

Figure B.1: Trimmed Mean and Median Inflation

Notes: Notes: Median and mean inflation are authors' calculations from the PCEPI data. Trimmed mean inflation is calculated by the Federal Reserve Bank of Dallas (Federal Reserve Bank of Dallas, 2021). The figure plots year over year inflation calculated as a weighted mean of individual categories and the median inflation category as calculated using the methodology of the Federal Reserve Bank of Cleveland (Federal Reserve Bank of Cleveland, 2021). In this methodology the median inflation category is selected for each month and then the rates for each category in the previous 12 months are combined to obtain an annual rate. Shaded areas indicate recessions as defined by the NBER.

B Inflation Distribution Appendix

Figure B.1 presents mean, trimmed mean, and median inflation.

B.1 Calculating Inflation Percentiles

We calculate inflation percentiles in Figure 2 in a weighted manner by ranking inflation series from lowest to highest and defining a percentile, x, such that x percent of consumer spending is on inflation series with a lower rate of inflation. In Table 1 we present averages of unweighted percentiles which are calculated by assigning an equal weight to each inflation series.

C Facts Appendix

C.1 QQ Plots

Figure C.1: Normal Distribution QQ Plot

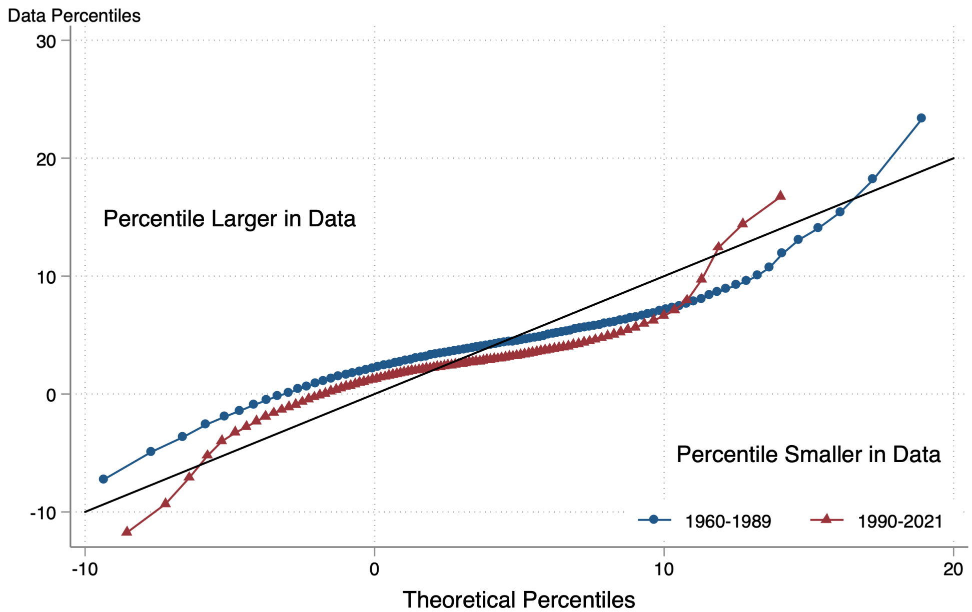

Notes: Notes: Authors' calculations from the PCEPI data. The graph plots the percentiles of detailed inflation rates against the percentiles of a normal distribution with the same mean and standard deviation.

Figure C.2: Top 10 Inflation vs Aggregate Inflation

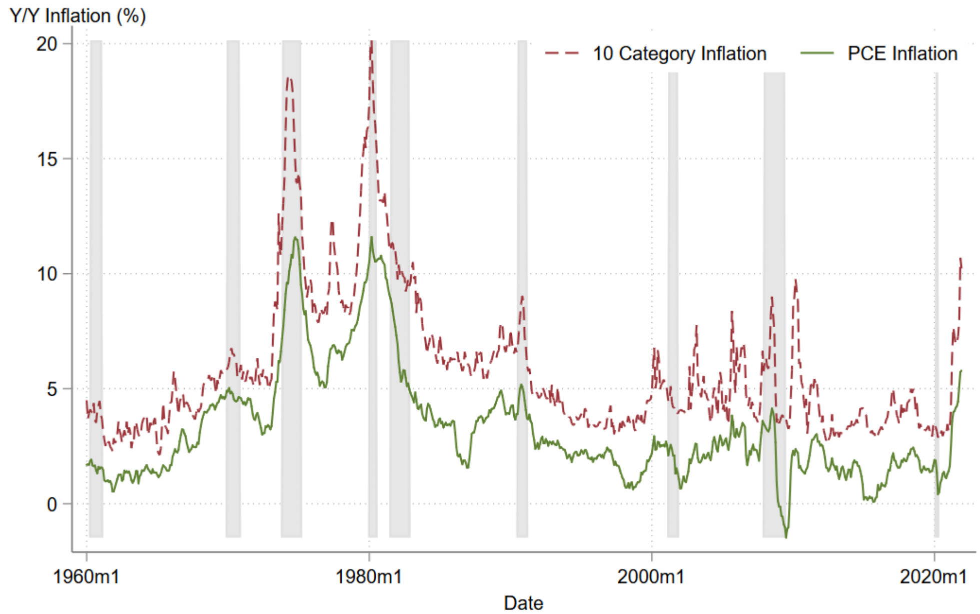

Notes: Notes: Authors' calculations from the PCEPI data. 10-category inflation is the average rate of inflation for the 10 largest inflation categories in terms of expenditure. PCE inflation is the headline inflation rate.

Figure C.3: Covariance and Variance of Contributions to Inflation

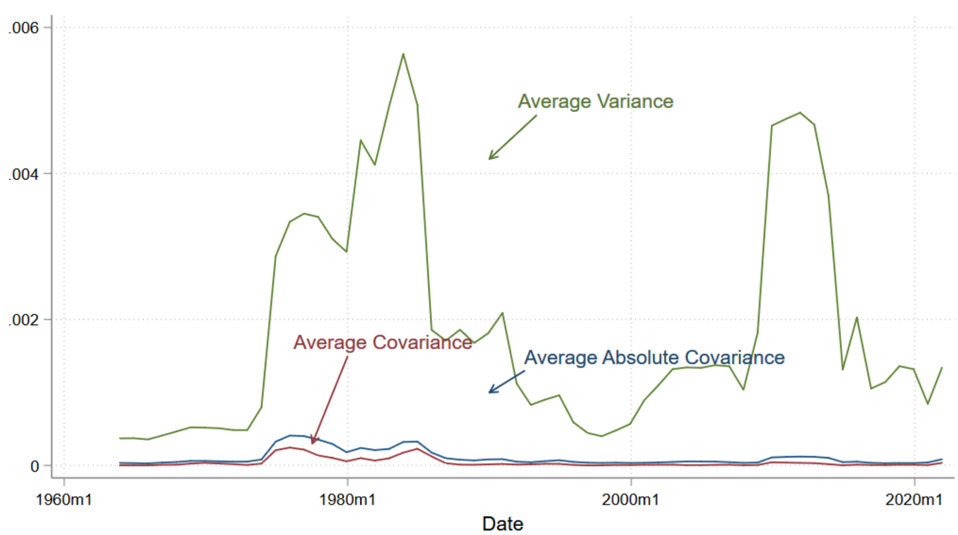

Notes: Notes: Authors' calculations from the PCEPI data. Numbers are the average covariance and variance of the contributions of detailed items to headline PCE.

D Calibration Appendix

Table D.1: Parameters Used in Estimation

Parameter	Value
δ	0.5
ρ_a	0.9
ρ_μ	0.9
ρ_k	0.9
β	0.9967
θ	6
ϕ_π	1.34
ϕ_y	0.33/12
ϕ_{g_c}	0
ρ_i	0
φ	2
σ	1
η	2

Notes: This table describes the parameters used in the model.