Revisiting the Cost of Inflation: Identification Using Vintage Moments*

Michael Jenuwine†

November 2026

Abstract

Measuring the welfare cost of inflation requires determining whether observed movements in relative prices are driven by nominal rigidities or by changes in firms' fundamentals. This paper introduces a new empirical approach to measuring the cost of inflation in sticky price models: conditioning the cross-section of price changes on *vintage*, or duration of the price spell. Moments derived from these conditioned distributions (vintage moments) contain information absent from those derived from the unconditioned cross-section. Prices with low vintages have accumulated few shocks to fundamentals, while those at high vintages have accumulated many. Comparing price changes across vintages thus helps identify the magnitude and persistence of those shocks – an identification not possible with cross-sectional data alone. I document vintage moments using micro price data, and estimate a quantitative sticky price model to match them. Compared to an analogous model estimated to match cross-sectional moments, I recover idiosyncratic shocks with lower persistence. In the model that matches vintage moments, the welfare cost of inflation is marginally higher at a 2% steady-state inflation rate, and responds twice as strongly to changes in steady-state inflation.

*I am deeply grateful to Mark Aguiar (chair), Richard Rogerson, and Karthik Sastry for their guidance and support. I thank John Grigsby, Nicholas Hommel, Yuyang Jiang, Aleksei Oskolkov, Mikkel Plagborg-Moller, Joris Tielens, Eric Qian, Gianluca Violante, Yinuo Zhang, and seminar participants at Princeton for their insightful conversations and feedback. This research benefited from financial support from the International Economics Section at Princeton University. The views expressed in this paper are those of the author and do not necessarily reflect the views of the Federal Reserve Board of Governors, the Federal Reserve System, or any other institution with which the author is affiliated.

†Federal Reserve Board of Governors. E-mail: michael.a.jenuwine@frb.gov

1 Introduction

Nominal rigidities are the primary friction of workhorse models used to study monetary policy and inflation. The welfare cost of inflation in these models is determined by the gaps between firms' listed prices and ideal prices – the latter being those that would maximize profits in a flexible-price benchmark. Properly measuring this welfare cost requires decomposing observed price changes into an efficient component driven by movements in ideal prices, and a costly component driven by nominal rigidities. This is a challenge in practice, as neither ideal prices nor nominal rigidities are directly observed. Measurement typically requires estimating a model of firm price setting to match moments of observed data on prices, and calculating the cost of inflation implied by that estimated model.

This paper presents a new empirical approach to estimating models of firm price setting: matching vintage-conditional moments. These are moments derived from distributions of price changes conditioned on the *duration of the preceding price spell*, which I refer to as vintage.[1] I demonstrate with an illustrative example that vintage-conditional moments (henceforth "vintage moments") contain useful information for estimating sticky price models, and that this information is absent from the moments typically used, which are derived from a single pooled cross-section of price changes. In particular, vintage moments better identify the persistence of shocks to firms' ideal prices. Next, I document systematic empirical patterns in vintage moments that sticky price models fail to reproduce, given previously estimated parameters. Finally, I estimate a standard New Keynesian general equilibrium model to match vintage moments and calculate the associated welfare cost of inflation. I compare this welfare cost to that implied by the same model, but estimated to match moments of the cross-section of price changes.

I begin with a stylized model. Firms minimize the present-discounted quadratic distance between their price and their ideal price, subject to nominal rigidities in the form of random menu costs. Ideal prices follow a stochastic process with idiosyncratic shocks. This stripped-down setup spotlights the core difficulty of model estimation: jointly determining the distribution of menu costs and the stochastic process for idiosyncratic shocks when only prices are observed. It is thus an ideal setting to illustrate how and why vintage moments assist model estimation.

I demonstrate the usefulness of vintage moments over cross-sectional moments through a simple example with exponentially distributed menu costs and an AR(1) shock process with persistence ρ. Comparing the opposing cases of a permanent-

[1] For example, if a product changes its price two months in a row, say in January and in February, the second price change would have a vintage of one. Alternatively, if a product changes its price in January and then in April, the second price change would have a vintage of three.

shocks economy ($\rho = 1$) and a transitory-shocks economy ($\rho = 0$), I show that both can exactly match cross-sectional data on the overall frequency and variance of price changes. Furthermore, the cross-sectional distributions of price changes implied by each case are nearly identical, so examining additional cross-sectional moments offers little benefit for differentiating between the two.

On the other hand, examining vintage moments provides a simple heuristic for distinguishing between the two cases. In the permanent shocks economy, the vintage-conditional probability of price adjustment and vintage-conditional variance of price changes both monotonically increase in vintage. This is because additional shocks continuously push firms' ideal prices away from their current prices, resulting in larger and more frequent price changes over time. However, in the transitory shocks economy, adjustment probabilities and variances both monotonically decrease in vintage. This is because firms with a large incentive to change their price select into doing so early. Additional shocks do not disperse ideal prices, so this selection effect dominates and price changes become smaller and less frequent over time. I demonstrate numerically that vintage moments show similar patterns for intermediate values of $\rho \in (0, 1)$.

The cross-sectional data performs poorly because it cannot distinguish to what degree a price change is motivated by a firm's current shock as opposed to accumulated past shocks. Incorporating vintage is beneficial because it adds information about how long it takes firms to select into changing prices. Firms at low vintage have been hit by few idiosyncratic shocks, whereas those at high vintages have been hit by many such shocks. Comparing price change distributions across vintage is thus intuitively informative about the stochastic process governing those shocks.

Next, I document the behavior of vintage moments in the panel of retailers' prices underlying the United Kingdom's Consumer Price Index. Within narrow product categories, I highlight two systematic patterns that previously estimated sticky price models do not reproduce. First, the vintage-conditional probability of price adjustment, or hazard rate, declines as vintage increases, with occasional spikes at 4- and 12-month intervals. Second, depending on the product, the vintage-conditional variance of price changes either declines monotonically as vintage increases, or exhibits a single early peak at a vintage of 2-4 months before declining. Through the lens of the stylized model, both of these facts indicate that idiosyncratic shocks are transitory.

To quantify the importance of using vintage moments in estimation, I move to a New Keynesian general equilibrium setting that more closely mirrors the types of quantitative models used for policymaking. Firms compete monopolistically and face idiosyncratic shocks to productivity that follow an AR(1) process. Technology is linear in labor and nominal wages grow at a constant rate of inflation. In order to change prices, firms must pay a random menu cost in units of labor, which is

distributed exponentially.

I estimate this model twice, targeting different sets of moments. First, to stand in for a typical calibration procedure, I match a set of commonly used cross-sectional moments. Second, I match a set of vintage moments. Comparing the two calibrations, the one matching vintage moments recovers idiosyncratic productivity shocks with larger innovations and a lower persistence (a monthly autocorrelation of .15 as opposed to .79), as well as lower menu costs on average.

For each set of estimated parameters, I solve the quantitative model at various rates of steady-state inflation and compute welfare loss as consumption-equivalent units relative to the flexible price economy. At the common inflation target of 2%, the model matching vintage moments implies a welfare cost of inflation that is .1 percentage points higher compared to that of the model matching cross-sectional moments. This welfare cost is also more sensitive to changes in steady-state inflation in the model matching vintage moments: moving from an annualized inflation rate of 0% to 10% raises welfare loss by about .17 percentage points, compared to a rise of only .09 percentage points in the model matching cross-sectional moments.

These welfare results are driven by differences in the productivity shock process across the two calibrations. In the calibration matching vintage moments, firms respond less aggressively to the transitory productivity shocks, understanding that they will mean-revert. This drives the difference in level of welfare loss. As the inflation rate rises, it is easier for firms facing persistent shocks to change prices in a way that adjusts to those shocks while simultaneously keeping pace with inflation. This drives the difference in slopes.

Related Literature. This paper primarily contributes to a literature concerning how to estimate sticky price models using micro-price data. Models with menu costs are primarily estimated to match moments of the cross-section of price changes (Nakamura and Steinsson, 2008, 2010). Analytical mappings between model primitives and price change data can be derived in particular settings (Alvarez and Lippi, 2014; Alvarez et al., 2016). Notably, Alvarez et al. (2021) show how to identify the distribution of random menu costs that rationalizes any observed cross-section price changes, provided that idiosyncratic shocks follow a random walk.

Outside of the cross-section, Alvarez et al. (2011) and Baley and Blanco (2019) use vintage-conditional hazard rates to discipline learning behavior in price setting models with imperfect information. Baley and Blanco (2021) observe that the covariance between vintage and firm state can identify asymmetric adjustment costs. These studies, however, rely on non-stationary idiosyncratic shocks to firms' states.

My contribution is to introduce vintage moments as an important feature of micro-price data to match when estimating sticky price models with stationary idiosyncratic

shocks, and to show that these moments identify the persistence of those shocks.

Second, I contribute to a literature measuring the cost of inflation in sticky price models. Nakamura et al. (2018) demonstrate how the nature of price stickiness impacts the welfare cost of inflation, while Golosov and Lucas (2007) highlight the size of shocks to firms' ideal prices and Midrigan (2011) highlights the shape of that shock distribution as playing major roles. Others attempt to empirically determine the cost of inflation by using correlations between the inflation rate and dispersion in price changes (Lach and Tsiddon, 1992) or dispersion in prices (Sheremirov, 2020), or by exploiting varying trends in relative prices across products (Adam and Weber, 2023; Adam et al., 2023).

I contribute to this literature by demonstrating that the persistence of shocks to ideal prices is an important determinant of the cost of inflation. In particular, when shocks to ideal prices are more transitory, the welfare cost of inflation implied by sticky price models is larger and increases more quickly as inflation rises.

Additionally, I connect to a literature documenting dynamic features of micro-price data.[2] In both old surveys (Klenow and Malin, 2010) and recent papers describing micro-price data (Gautier et al., 2024), these features receive little attention. Vintage-conditional hazard rates are the exception – they have been documented and are found to decrease in vintage (Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2008; Vavra, 2010).[3] Campbell and Eden (2014) also measure differences in the distribution of relative prices at young and old vintages. In the same UK CPI micro-price data I use, Bunn and Ellis (2012a,b) document facts concerning vintage-conditional hazard rates and sizes of price changes.

My work adds to the evidence that hazard rates decline in vintage, and documents a new fact that the vintage-conditional variance of price changes declines in vintage as well.

The rest of the paper is organized as follows. Section 2 presents a stylized model of firm price setting, and provides intuition for the usefulness of vintage moments. Section 3 documents the behavior of vintage moments in the data underlying the UK CPI, and Section 4 presents the benefit of using vintage moments for estimation and its implications for the welfare cost of inflation. Section 5 concludes.

[2] I call vintage moments "dynamic" in the sense that they exploit variation *over time within a panel unit*. One can also exploit time variation by examining how the cross-section of price changes evolves over the business cycle, as in Gagnon (2009), Vavra (2013) and Berger and Vavra (2018).

[3] This finding has recently been called into question – Alvarez et al. (2025) and Karadi et al. (2023) measure hazard rates that increase in vintage.

2 Theoretical Framework

The theoretical framework is a standard discrete-time, random menu cost model with idiosyncratic shocks in the tradition of Caballero and Engel (1993). This framework contains the essential features necessary to highlight the core difficulty in estimating models of firm price setting with menu costs, yet remains general enough to nest most typical specifications of such models.

2.1 Setup.

Consider a measure 1 continuum of firms indexed by i, living in discrete time. These firms face a dynamic, forward-looking price setting problem subject to random menu costs.

Following Alvarez et al. (2016), I model a firm's per-period loss as $B(p_{it} - p_{it}^*)^2$: a quadratic function of the deviation between its current log price, p_{it}, and its "ideal" log price, p_{it}^*. This quadratic loss function can be microfounded as a second-order approximation of firm profits around p_{it}^*, normalized by the firm's non-stochastic steady-state profits. In this quadratic approximation, B represents the curvature of the firm's profit function and the ideal price represents the firm's *static, profit-maximizing price*.

I model ideal prices p_{it}^* as composed of a firm fixed effect, an aggregate component that follows a deterministic trend, and an idiosyncratic component that evolves by some first-order Markov process.

$$p_{it}^* = p_{i0}^* + \underbrace{\pi t}_{\substack{\text{aggregate} \\ \text{component}}} + \underbrace{a_{it}}_{\substack{\text{idiosyncratic} \\ \text{component}}} .$$

The trend π reflects the aggregate inflation rate, and the idiosyncratic component a_{it} reflects movements in firm-specific supply and demand. This idiosyncratic shock evolves according to $G(a_{it}|a_{it-1})$, the conditional CDF of a_{it} given a_{it-1}. I assume this process is either a martingale, or stationary with an unconditional mean of zero. This admits the two most commonly used processes for idiosyncratic shocks – a random walk and an AR(1) – as well as any finite-state first-order Markov process. I do not permit aggregate shocks to ideal prices for expositional simplicity, but they can be accommodated.[4]

I model nominal frictions as random menu costs, κ_{it}, drawn independently over time and across firms from distribution $H(\kappa)$. This admits common pricing frictions, such as the Calvo friction, a fixed menu cost as in Golosov and Lucas (2007), or the

[4]This specification as written also assumes that state a is a scalar, but one can easily extend it to a scalar-valued function of some state vector that follows a Markov process.

Calvo-plus friction of Nakamura and Steinsson (2010). It rules out specifications in which firms' menu costs are serially correlated, or in which the distribution of menu costs is time dependent, such as Taylor pricing.

Firm problem. At the start of each period, the firm observes its menu cost draw and its new ideal price. The firm then decides whether to pay the menu cost and update its price, or leave its price unchanged.

I define a firm's *price gap* as $x_{it} \equiv p_{it-1} - p^*_{it}$, and write the firm's value function in recursive form as

$$\begin{aligned} V(x,a,\kappa) = \min \Big\{ & Bx^2 + \beta \mathbb{E}\left[V\left(x + a - \pi - a', a', \kappa'\right) | x, a\right], \\ & \kappa + \min_{\bar{x}} B\bar{x}^2 + \beta \mathbb{E}\left[V\left(\bar{x} + a - \pi - a', a', \kappa'\right) | \bar{x}, a\right] \Big\}, \\ & a' \sim G(a'|a), \quad \kappa \sim H(\kappa). \end{aligned}$$

where the law of motion for price gaps is inherited from the law of motion for ideal prices, and setting a new price is equivalent to choosing a new price gap.[5]

For a firm with state (x, a) that selects into changing its price, optimal reset gap $\bar{x}$ is determined by the FOC

$$2B\bar{x} + \beta \mathbb{E}\left[\frac{\partial V}{\partial x}\left(\bar{x} + a - \pi - a', a', \kappa'\right) | a\right] = 0.$$

At this level of generality, a closed-form expression for $\bar{x}$ cannot be written. However, $\bar{x}$ is a function only of the current idiosyncratic shock, a. I therefore define the reset gap policy function $\bar{x}(a)$ as the solution to this FOC, and $V^c(a)$ as the expected value of changing price:

$$V^c(a) = B\bar{x}(a)^2 + \beta \mathbb{E}[V(x', a', \kappa') | \bar{x}(a), a]$$

I similarly define the expected value of not changing price for a firm with state (x, a):

$$V^n(x,a) \equiv Bx^2 + \beta \mathbb{E}[V(x', a', \kappa') | x, a].$$

The adjustment hazard function, or the probability of a firm with state (x, a)

[5]Note that in cases where idiosyncratic shocks a follow a martingale, the value function can be written in terms of just price gap x as a state, as can subsequent derivations of stationary distributions and hazard functions.

changing its price, is

$$\begin{aligned}\Lambda(x,a) &= Pr\left(\kappa + V^c(a) < V^n(x,a)\right) \\ &= H\left(V^n(x,a) - V^c(a)\right).\end{aligned}$$

Note that the probability of price adjustment is a function of both the firm's price gap and the realization of the idiosyncratic state.

The primary object of interest in this model is the distribution of firm states *after* price change decisions are made: $f(x,a)$. This object, which I call the *ex-post distribution* of firm states, describes the size of all price gaps in the economy, and in a richer model would map to welfare loss. It is closely related to the *ex-ante distribution* of firm states from *before* price change decisions are made, denoted $\hat{f}(x,a)$. The two of them are defined together as follows:

$$f(x,a) = \underbrace{(1-\Lambda(x,a))\hat{f}(x,a)}_{\text{non price-changers}} + \underbrace{\delta_0(x-\bar{x}(a))\int_{-\infty}^{\infty}\Lambda(y,a)\hat{f}(y,a)dy}_{\text{price changers}},$$

$$\hat{f}(x',a') = \int_{-\infty}^{\infty} g(a'|a)f(x'-a+\pi+a',a)da.$$

$g(a'|a)$ is the PDF of idiosyncratic state a' conditional on previous state a, and $\delta_0(\cdot)$ is the Dirac delta function.

These distributions depend primarily on two objects: the stochastic process for idiosyncratic states, and the distribution of menu costs. The goal is therefore to estimate the parameters of these two functions given panel data on the prices set by firms. When faced with this estimation problem, the literature typically only considers two features of these data to match: the frequency of price changes, and the cross-sectional distribution of price changes. I illustrate that using vintage-conditional features of the data is superior to using cross-sectional data alone.

2.2 A simple example.

Consider a simple example in which $\pi = 0$, and $\beta = 0$. Suppose that idiosyncratic shocks follow an AR(1) process with persistence ρ and normal innovations, and that random menu costs are exponentially distributed.[6]

$$G(a'|a) = \Phi\left(\frac{a'-\rho a}{\sigma}\right), \quad H(\kappa) = 1 - \exp\left(-\frac{\kappa}{2K}\right)$$

[6]These are reasonable assumptions, with the exception of $\beta = 0$. Exponentially distributed menu costs produce empirically reasonable distributions of price changes and $\pi = 0$ is a good approximation of a low inflation environment. It is strong to assume firms are myopic with $\beta = 0$, but necessary for some analytical expressions I derive. The assumption of myopia will later be relaxed.

To best highlight the key message, I focus on two polar cases: a permanent-shocks economy in which $\rho = 1$, and a transitory-shocks economy in which $\rho = 0$. Given a panel of prices, I ask which features of the data best differentiate between the permanent-shocks economy and the transitory-shocks economy.

Each economy has two free parameters: σ and K. It is standard practice for sticky price models to match the frequency of price changes and some notion of the size of price changes. I show that for each economy, I can analytically map the frequency of price changes and the variance of price changes to the two free parameters.

Proposition 1. *Given data on the frequency of price changes, $\hat{\lambda}$, and the variance of price changes $\widehat{Var}(\Delta p)$:*

- *There exists unique $\{\sigma_p, K_p\}$ such that the frequency and variance of price changes implied by the permanent shocks model exactly match these data.*
- *There exists unique $\{\sigma_t, K_t\}$ such that the frequency and variance of price changes implied by the transitory shocks model exactly match these data.*

Even though the two economies feature fundamentally different processes for ideal prices, both can match these commonly used moments. If one restricts oneself to looking at the cross-section of price changes, it is unclear which additional feature of that distribution should be used to distinguish the two economies. While one show analytically that the permanent and transitory shocks economies imply different cross-sectional distributions of price changes, no particular moment of those distributions stands out. This can be seen in the left panel of Figure 1, which plots an example of those two model-implied cross-sections.[7] The two do not perfectly coincide, but it is unclear how to exploit the gaps between them for identification.

A better method is to move beyond the cross-section of price changes. I propose including *vintage*, or the number of periods since a firm last changed its price, as an observable. This exploits information in the panel of prices that cross-sectional data ignores. I examine vintage-conditional analogues of the moments already used: vintage-conditional hazard rates, or probabilities of price change, and vintage-conditional variances of price changes. Given these vintage-conditional moments, I prove that there is a simple heuristic for telling the two economies apart.

Proposition 2.

- *In the permanent shocks model, the vintage-conditional hazard rate and the vintage-conditional variance of price changes are both monotonically increasing in vintage.*

[7] These are constructed to have frequency of price changes equal to .088 and variance of price changes equal to .025, which correspond to those statistics in micro-data underlying the UK CPI.

- *In the transitory shocks model, the vintage-conditional hazard rate and the vintage-conditional variance of price changes are both monotonically decreasing in vintage.*

The intuition for these results is as follows. In the permanent shocks model, all price-changing firms set their price gaps to $\bar{x} = 0$. As shocks are permanent, the distribution of future price gaps is centered at 0 for every firm, and spreads out as shocks accumulate over time. This means the distribution of realized price gaps across firms becomes more dispersed as time passes. Firms in the tails of this distribution are more likely to change their price, and make larger price changes when they do. So, as vintage increases, the tails of the price gap distribution become relatively fatter, which leads to both more frequent and larger price changes.

The machinery of the transitory shocks model is more subtle. I explain it in two steps. All price changing firms set their price gaps to $\bar{x} = 0$. As shocks are transitory, the distribution of future price gaps is constant for every firm. Suppose for the moment that this distribution is centered at 0 for every firm. If this were the case, the distribution of realized price gaps across firms would be the same over time. As a result, the hazard rate and variance of price changes would be constant across vintage.

In actuality, the distribution of future price gaps is not centered at 0 for every firm. It is instead centered around each firm's idiosyncratic shock when it originally changed price. This creates a selection effect. Firms with an extreme shock are likely to immediately face a large movement in their price gap and change their price. Over time, the remaining firms become those with less extreme initial shocks. Due to this effect, the distribution of realized price gaps across firms becomes less dispersed as time passes. So, as vintage increases, the tails of the price gap distribution get relatively smaller, which leads to both less frequent and smaller price changes.

The right panel of Figure 1 plots the model-implied vintage-conditional variances of price changes, for vintages of 12 months or fewer. The model-implied vintage-conditional hazard rates follow similar profiles (see Appendix D.1). The difference between the two economies is stark and immediate. Vintage-conditional moments are highly informative for differentiating between these two economies.[8]

[8]The informativeness of vintage moments as opposed to higher-order moments of the cross-sectional distribution is further corroborated by Monte Carlo simulations. See appendix B for details.

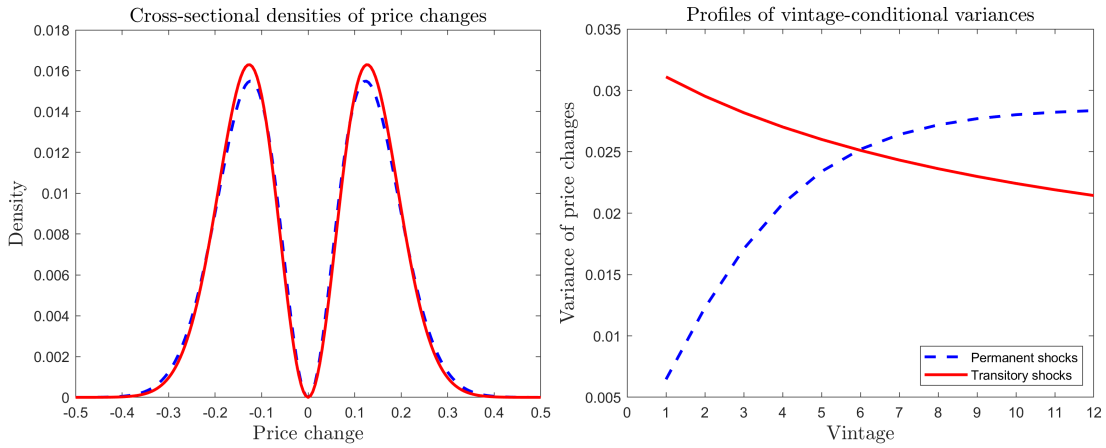


Figure 1: Cross-sections vs vintage-conditional variances

The left panel plots example cross-sections of price changes for the permanent and transitory shocks economies. The right panel plots the vintage-conditional variance of price changes against vintage for the same two economies.

I now show that vintage-conditional moments are also informative in a less stylized setting. In particular, I relax the assumption that firms are myopic and allow for non-extreme values of the AR parameter, ρ. I set realistic values of monthly discount rate $\beta = (.96)^{\frac{1}{12}}$ and inflation rate $\pi = .0017$ (roughly 2% annualized). Then, for various fixed values of $\rho \in (0, 1)$, I select K and σ to match frequency of price changes $\hat{\lambda} = .088$ and variance of price changes $\widehat{Var}(\Delta p) = .025$. These correspond to the aggregate frequency and variance of price changes in the micro-data underlying the United Kingdom's Consumer Price Index.

Given these parameters, I then solve and simulate the models to find the implied cross-sectional distribution of price changes, and profiles of vintage-conditional hazard rates and variances. For selected values of ρ, Figure 2 plots simulated cross-sectional distributions of price changes (left), and profiles of vintage-conditional variances (right). As with the extreme cases, the simulated cross-sectional distributions differ, but the differences are small, while the profiles of vintage-conditional variances vary widely.

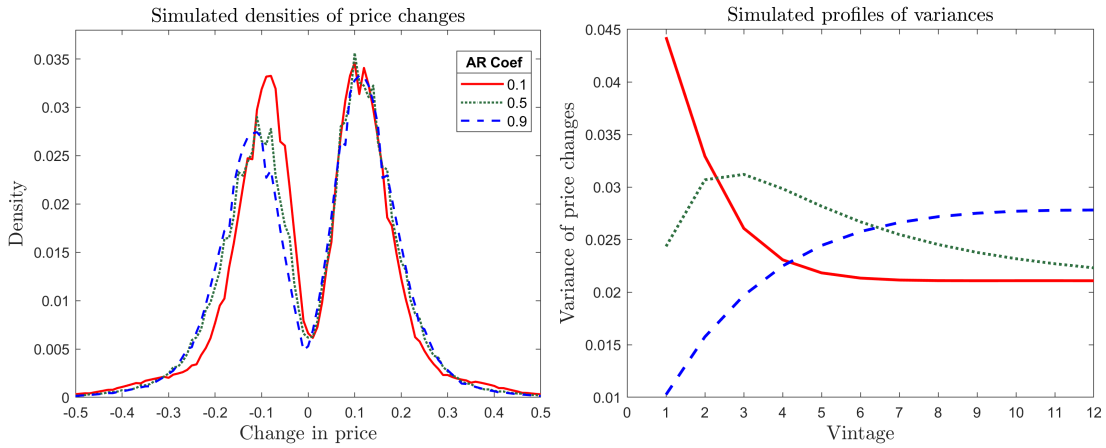


Figure 2: Simulated price change cross-sections

The left panel plots simulated cross-sections of price changes in the economy for selected AR coefficients. The right panel plots the vintage-conditional variance of price changes against vintage for the same economies.

Furthermore, the profiles of the simulated vintage moments mirror the theoretical results from the extreme cases. As ρ approaches 0 and shocks become more transitory, the variance of price changes approaches a monotonically decreasing profile. As ρ approaches 1 and shocks become more permanent, the variance of price changes approaches a monotonically increasing profile. These patterns are also produced in the vintage-conditional hazard rates (see Appendix D.1). This reinforces the informativeness of vintage-conditional moments, as well as the heuristic of slope for determining whether shocks are permanent or transitory.

2.3 Identifying model primitives.

Vintage-conditional data appear to identify the ideal price process more effectively than cross-sectional data. This is unsurprising in some sense: vintage-conditional data contain information about both the size of price changes and the duration of price spells, whereas cross-sectional data contain information only about the former. However, to understand why vintage data are useful *specifically* for identifying the ideal price process, it helps to look at the model-implied cross-sectional and vintage-conditional price change distributions.

In the model, a firm with ex-ante state (x, a) changes price with probability $\Lambda(x, a)$. The model-implied frequency of price changes is therefore the average of $\Lambda(x, a)$ over the stationary distribution of ex-ante firm states, $\hat{f}(x, a)$,

$$\lambda = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Lambda(x, a) \hat{f}(x, a) dx \, da$$

When a firm with state (x, a) changes its price, the size of that price change is $\Delta p(x, a) = \bar{x}(a) - x$. The model-implied cross-sectional distribution of price changes

is therefore

$$q(\Delta p) = \frac{1}{\lambda}\int_{-\infty}^{\infty} \Lambda(\bar{x}(a) - \Delta p, a)\hat{f}(\bar{x}(a) - \Delta p, a)da$$

When $\bar{x}(a)$ is non-constant, there is no one-to-one mapping between the size of a price change and the underlying state of a firm.[9] This is why the cross-section of price changes contains an integration over idiosyncratic states. In this case, the size of price changes alone cannot distinguish whether a price change is driven by the current shock realization, a, or the accumulated price gaps from shock realizations, x. As a result, differing stochastic processes for a can generate similar cross-sectional data.

Using vintage-conditional moments ameliorates this issue. Defining τ as vintage, with $\tau = 0$ representing firms that just changed their price in the current period, we have the hazard rate and distribution of price changes, conditional on vintage,

$$\lambda_\tau = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \Lambda(x,a)\hat{f}_\tau(x,a)\, dx\, da,$$
$$q_\tau(\Delta p) = \frac{1}{\lambda_\tau}\int_{-\infty}^{\infty} \Lambda(\bar{x}(a) - \Delta p, a)\hat{f}_\tau(\bar{x}(a) - \Delta p, a)da,$$

where $\hat{f}_\tau(x,a)$ is the ex-ante distribution of firm states for firms that last changed their price τ periods ago. Vintage data are useful because this distribution of firm states changes in a predictable, recursive manner across vintage. At $\tau = 0$, price changers reset their price gaps according to policy function $x = \bar{x}(a)$. This provides an initial distribution of firm states $f_0(x,a)$:

$$f_0(x,a) = \frac{\delta_0(x - \bar{x}(a))\int_{-\infty}^{\infty}\Lambda(y,a)\hat{f}(y,a)dy}{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\Lambda(y,a)\hat{f}(y,a)dy\, da}.$$

The numerator of this expression is the measure of firms setting new price gap $\bar{x}(a)$, scaled by the overall measure of price changers.

Given the ex-post distribution of firm states for firms who last changed price $\tau - 1$ periods ago, $f_{\tau-1}(x,a)$, we can apply new idiosyncratic shocks to get the ex-ante distribution of firm states for those who last changed price τ periods ago, but have not yet decided whether to change price today,

$$\hat{f}_\tau(x',a') = \int_{-\infty}^{\infty} g(a'|a)f_{\tau-1}(x' - a + \pi + a', a)da,$$

and then apply the adjustment hazard function to find the ex-post distribution of firm states for firms who ultimately decide not to change their price, and now last

[9] In the special case where idiosyncratic shocks follow a random walk, $\bar{x}(a)$ is constant and such a one-to-one mapping exists. Alvarez et al. (2021) show how this mapping can be used to determine all model primitives from the cross-sectional data alone.

changed price τ periods ago:

$$f_\tau(x,a) = \frac{(1-\Lambda(x,a))\hat{f}_\tau(x,a)}{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}(1-\Lambda(y,a))\hat{f}_\tau(y,a)dy\,da}.$$

We know that firms with vintage τ have been hit by τ idiosyncratic shocks before choosing to change price. Differences in the distribution of price changes from one vintage to the next are therefore informative about transition probabilities in the idiosyncratic state, which are exactly what is needed to identify the ideal price process.

3 Data Description and Facts

I now document properties of vintage-conditional distributions in the price quote micro-data underlying the United Kingdom's Consumer Price Index (CPI), collected by the Office of National Statistics (ONS). I am particularly interested in vintage-conditional hazard rates and variances, as these are the moments that are heuristically informative about the persistence of idiosyncratic shocks. I document two facts. First, consistent with previous work, hazard rates decrease in vintage. Second, and less appreciated, the variance of price changes also broadly decreases in vintage, and either does so monotonically or has a single peak in the 2 to 4 month range.

3.1 Data description and cleaning.

The UK CPI microdata are a rotating panel of price quotes collected monthly from retailers across the country. Each observation contains characteristics about the retail outlet being sampled from, the product being sampled, and the product's price.[10] I consider data collected between January 1996 and December 2019. From these data, I want to extract price changes and their associated vintages.

In order to measure vintage, I need to follow the same product over time. ONS data collectors are instructed to record prices for the same exact product over time within a retailer, so I construct product identifiers from the data using product categories and retailer characteristics. I drop products that cannot be uniquely identified. The ONS records when a product changes from one month to the next, including changes in package size and quality substitutions. When substitutions occur, future prices continue to be collected for the new (substitute) product. In such cases, I assign all post-substitution prices to a new product identifier. I drop prices that fail to satisfy ONS validation procedures, as well as those associated with missing and

[10] For additional description of the data, including sampling procedure, methods of price collection, and construction of weights, please refer to the ONS Consumer Price Indices Technical Manual (Office for National Statistics, 2024).

out of stock products. The remaining prices constitute my sample, which appears to be representative as the year-over-year change in CPI constructed from the sample closely tracks that of the official ONS release (See Figure C.1 in Appendix C).

I must also account for sales in the data. Sales make up a large portion of changes in retailers' prices, especially short term and small price changes (Nakamura and Steinsson, 2008; Kehoe and Midrigan, 2015). It is unlikely that sales reflect changes in retailers' ideal prices, so the theory is better suited to describe non-sale or "regular" prices. I identify sales using two approaches. First, I use an ONS variable that flags when a product is on sale or "recovering" from a sale. Second, I apply a 3-month V-shaped filter as in Kehoe and Midrigan (2015). This filter finds instances in which the price for a product falls, stays at some level for at most 3 months, and then returns to the same pre-decrease level. I flag prices satisfying either condition as "sale prices," and replace these with the most recent non-sale price to generate a panel of regular prices. After this, the final sample contains about 23 million price quotes and 2 million regular price changes that span 1,285 product categories. Additional details about sample construction and summary statistics can be found in Appendix C.

In this panel of regular prices, I determine vintage as follows. For each panel unit, I assign a vintage of 1 to periods following a price change, and increment the vintage by one every month until a price change occurs. As data are left-censored, vintage is unknown prior to the first price change of each panel unit, and I do not impute over gaps in the panel. Vintage is unknown for 869,773 price changes, of which 736,208 are the first price change in a panel unit. Restricting the sample to prices and price changes for which vintage is observed may result in selection toward products that change prices more frequently. In Appendix C I report summary statistics for price changes in the sample of all price changes and the restricted sample of price changes where vintage is observed. In the restricted sample, price changes are more frequent and are slightly smaller in absolute size. Overall, I do not view selection on observing vintage as a major concern.

3.2 Vintage-conditional hazard rates.

I first examine vintage-conditional hazard rates. Figure 3 plots aggregate hazard rates, which are constructed for each vintage by dividing the number of price change observations by the number of price observations. The aggregate hazard rates decline in vintage, with spikes at regular 4- and 12-month intervals.

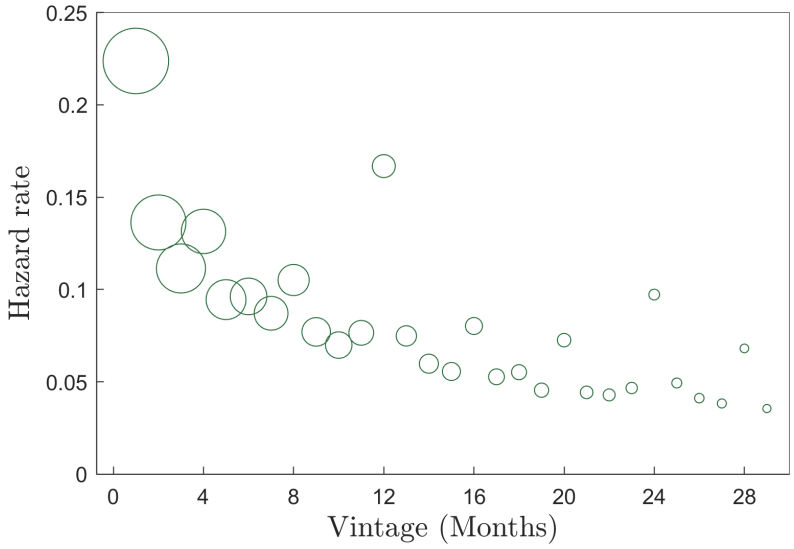


Figure 3: Vintage-conditional aggregate hazard rates

This figure plots the vintage-conditional hazard rate against vintage. The size of each marker is proportional to the number of observations at that vintage.

These aggregate statistics do not control for changes in the composition of products across vintage. If products exhibit heterogeneity in their frequency of price change, composition bias across vintage could generate a declining profile in aggregate hazard rates. To account for this, I calculate the vintage-conditional hazard rate within each vintage-product category pair[11] and estimate the following relationship between vintage and hazard rate,

$$\lambda_{j,\tau} = \alpha_j + \beta\tau + \delta_4\tau_4 + \delta_{12}\tau_{12} + u_{j,\tau},$$

where $\lambda_{j,\tau}$ is the hazard rate for products in category j with vintage τ, the α_j are product category fixed effects, and τ_4 and τ_{12} are indicator variables taking value 1 when vintage is a multiple of 4 or 12, respectively. These indicators capture price changes that occur at regular intervals, as seen in the aggregate data. Table 1 reports the results.

[11] Product categories are defined according to an internal ONS classification system and correspond roughly to COICOP micro-classes (6 digit categories).

	(1) $\lambda_{j,\tau}$	(2) $\lambda_{j,\tau}$	(3) $\lambda_{j,\tau}$
δ_4	.0105		.0143
	(.00138)		(.00136)
δ_{12}	.0671		.084
	(.00326)		(.00325)
β		-.00192	-.00274
		(.0000999)	(.0001)
Fixed Effects	Product Category	Product Category	Product Category
Observations	17172	17172	17172
R-Squared	.73	.725	.742

Standard errors in parentheses

Table 1: Vintage and hazard rate

The estimated coefficients on τ are significant and negative, indicating that the hazard rate declines in vintage on average across product categories. Not having changed price for an additional month predicts, on average, a 0.3 percentage point decrease in the probability of changing price in the current month.

In addition to the overall decline in the hazard rate with respect to vintage, prices also change significantly more frequently at regular 4- and 12-month intervals. In these data, Bunn and Ellis (2012a) show this behavior is almost entirely confined to services producers.[12] As labor is the primary input in service industries, and most wage changes occur quarterly or annually, rather than monthly (Grigsby et al., 2021), I interpret this behavior as reflecting longer intervals between firm-level cost shocks.[13] When estimating the same model as above with the time scale of vintage changed from every month to every 4 or 12 months, I obtain a similar declining relationship between hazard rate and vintage. These results are reported in Appendix D.2.

The regression above is restrictive in that it assumes product-level heterogeneity is captured by a level shift, and assumes the relationship between vintage and hazard rate is linear. To account for the former, I estimate a specification in which product-level heterogeneity is captured by a shift in the log of hazard rate, rather than level. Results are qualitatively similar to those above, and can be found in Appendix D.2. To account for the latter, I estimate a specification with vintage-specific fixed effects

[12] I reproduce this result in a sample of only market goods in Appendix D.2. This Appendix also contains subsequent empirical analyses on the sample of market goods, all of which are qualitatively similar to the full sample.

[13] An alternative view is that firms price services using time-dependent rules, as in Taylor pricing.

rather than a linear trend.

$$\lambda_{j,\tau} = \alpha_j + \sum_{t=1}^{24} \beta_t \mathbb{1}\{\tau = t\} + u_{j,\tau}$$

Figure 4 plots the estimated β_t coefficients. In addition to the overarching downward trend and spikes every 4 and 12 months, this specification also reveals that much of the decline occurs at early vintages.

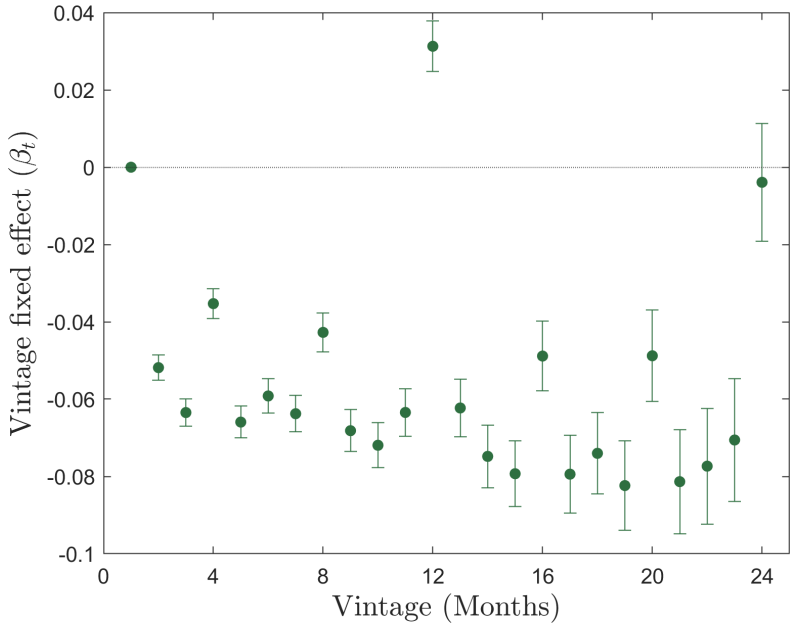


Figure 4: Coefficients from fixed effects regression

This figure plots the estimated vintage fixed effects from the regression $\lambda_{j,\tau} = \alpha_j + \sum_{t=1}^{24} \beta_t \mathbb{1}\{\tau = t\} + u_{j,\tau}$. Estimates are plotted relative to $\beta_1 = 0$. Bars represent a 95% confidence interval.

Overall, the findings here comport with the consensus established in previous work. After controlling for product-level heterogeneity, the hazard rate is downward sloping in vintage and most of the decrease comes in the first couple months (Nakamura and Steinsson, 2008; Cavallo and Rigobon, 2011; Campbell and Eden, 2014). Bunn and Ellis (2012a) also measure hazard rates in these same data, focusing on how vintage-conditional hazard rates differ across broad product groups. I instead focus on the downward trend that is common across product categories after removing heterogeneity.

3.3 Vintage-conditional variances of price changes.

I now document the vintage-conditional variances of price changes. To my knowledge, this paper is the first to examine this statistic rigorously. To construct these variances, I first standardize price changes for products i within each product category j by subtracting the average price change within the category and dividing by

the standard deviation of price changes within the category: $\Delta\tilde{p}_{it} = \frac{\Delta p_{it} - \mathbb{E}[\Delta p_{it}|j]}{\sqrt{Var(\Delta p_{it}|j)}}$. This standardization eliminates changes in the scale of aggregate variance as the composition of products changes across vintage. In subsequent within product-category analyses, this transformation has no effect. Figure 5 plots the aggregate variance of standardized price changes against vintage. Similar to the hazard rates, the variance of price changes declines in vintage before leveling off.

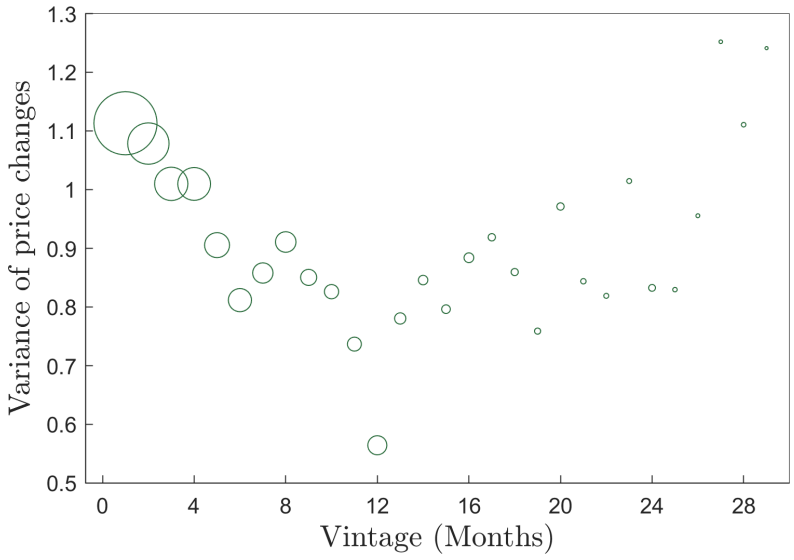


Figure 5: Vintage-conditional variances of price changes: aggregate

This figure plots the vintage-conditional variance of price changes against vintage. The size of each marker is proportional to the number of observations at that vintage.

As before, I want to control for heterogeneity in the relationship between the variance of price changes and vintage across product categories. To do so, I estimate

$$Var(\Delta\tilde{p}_{it}|j,\tau) = \alpha_j + \beta_j\tau + u_{j,\tau}$$

within each product category, where $Var(\Delta\tilde{p}_{it}|j,\tau)$ is the variance of price changes within product category j and vintage τ. Figure 6 plots a histogram of the estimated β_j coefficients, along with those that are significantly different from 0 at the 10% level. While the estimated coefficients exhibit considerable heterogeneity across product categories, the variance of price changes declines in vintage for the vast majority of cases. β_j is negative for 83% of product categories, and among the few product categories with a positive β_j, only 8% are statistically different from zero at the 10% level. Broadly speaking, the variance of price changes is greater at low vintages than at high vintages.

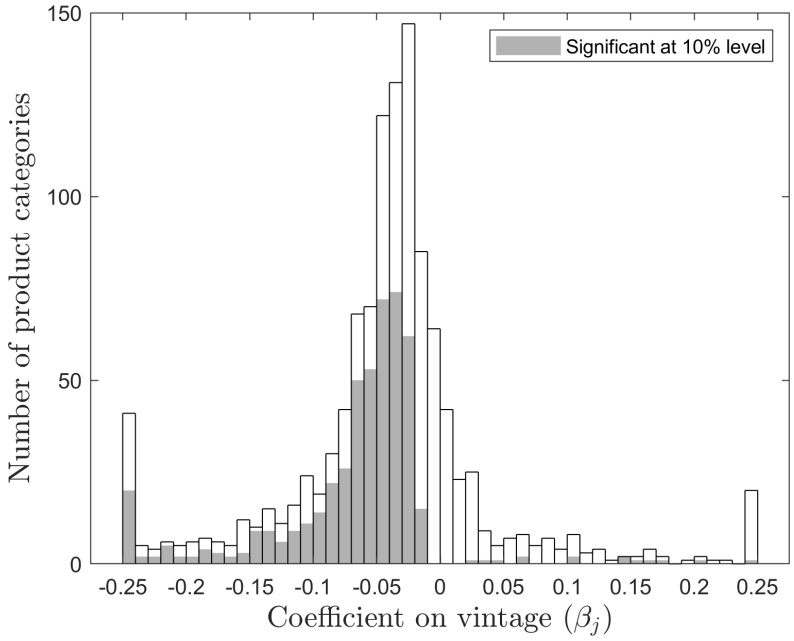


Figure 6: Estimated coefficients on vintage

This is a histogram of the estimated β_j coefficients from the regression $Var(\Delta \tilde{p}_{it}|j,\tau) = \alpha_j + \beta_j \tau + u_{j,\tau}$. Coefficients that are significant at the 10% level are overlaid in grey.

This regression specification assumes a linear relationship between variance of price changes and vintage, but the true relationship need not be linear, and indeed need not be monotonic. It is difficult to discern how the variance of price changes moves vintage-by-vintage within each product category, as some product-vintage pairs contain few observations. So, for each product category j and vintage $\bar{\tau}$ I pool price changes into two groups: those with a younger vintage ($\tau \leq \bar{\tau}$) and those with an older vintage ($\tau > \bar{\tau}$). For a given threshold vintage and product category, I say that the variance of price changes is "locally decreasing" if the variance of price changes is lower among the vintages above the threshold than among the vintages below the threshold. I define "locally increasing" analogously.

Figure 7 plots how often the variance of price changes is locally decreasing, either over thresholds within a product category (left panel) or over product categories within a threshold (right panel).

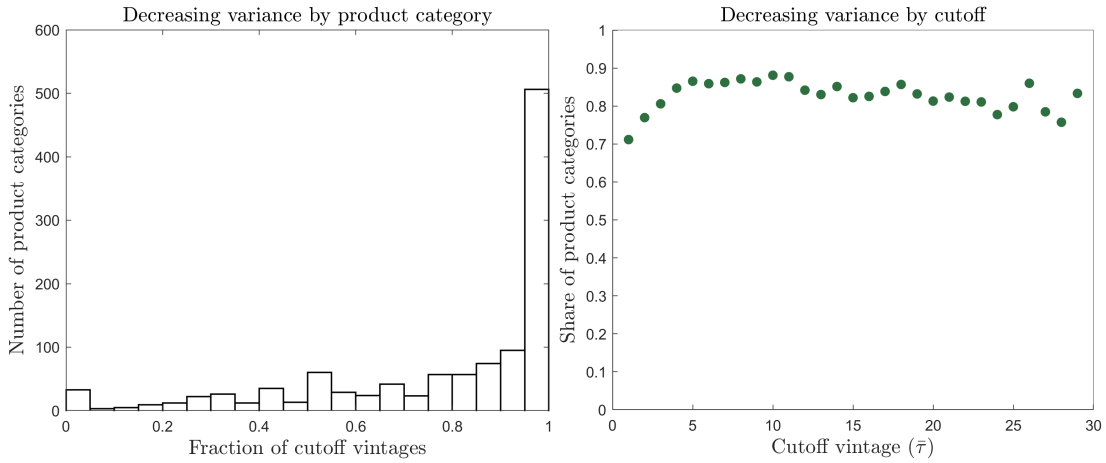


Figure 7: Prevalence of locally decreasing variance of price changes

The left panel is a histogram of the number of product categories by the share of cutoff vintages $\bar{\tau}$ for which the variance of price changes is locally decreasing in vintage. The right panel shows the fraction of product categories for which the variance of price changes is higher among vintages younger than $\bar{\tau}$ than among vintages older than $\bar{\tau}$, where $\bar{\tau}$ is vintage on the x-axis.

The left panel shows that the variance of price changes is locally decreasing for most vintages in most product categories. For 45% of product categories, the variance is locally decreasing at every vintage. Conversely, only 15% of product categories exhibit a locally increasing variance of price changes for 50% or more vintages. The right panel shows that in the few instances when the variance of price changes is locally increasing, it is typically at early threshold vintages of 2 to 4 months. Taken together, these results indicate that the variance of price changes either decreases monotonically in vintage, or rises briefly at low vintages before declining.

3.4 Unobserved heterogeneity.

The preceding analysis accounts for heterogeneity across product categories, which roughly correspond to COICOP-6 categories. There is likely unobserved heterogeneity within product categories that is not controlled for. As stated earlier, I are primarily concerned about unobserved firm types that generate composition bias across vintage.

For each panel unit i in product category j, I observe N_i price changes and vintages. This allows me to construct an empirical distribution function (EDF) of vintage for each unit. If these EDFs are identically distributed across every panel unit within a product category, then unobserved heterogeneity is not an issue. Statistics estimated within a product category can be interpreted as averages of those statistics across panel units.

As I lack the observations to test whether vintage EDFs are identically distributed for any pair of panel units, I instead test whether the EDF of each panel unit is distributed identically to the pooled EDF of vintage for all other panel units within the

same product category. In particular, I test whether the means of the two distributions and the variances of the two distributions are the same.[14] I then drop panel units for which either moment is significantly different at the 10% level. This removes about 75% of panel units and 55% of price change observations.

I repeat the exercises above on this limited sample and report results in Appendix D.3. The results are qualitatively similar. The one major difference is that, with the linear regression specification, I estimate that hazard rates are very slightly increasing in vintage. However, in the specification regressing hazard rates on vintage fixed effects, I still find that the estimated fixed effects decline markedly in the early vintages. I continue to estimate that the variances of price changes decline in vintage for the vast majority of product categories.

Summary Overall, in the data both hazard rates and variances of price changes are declining in vintage. Hazard rates decline especially in the first months following a price change, and the variance of price changes either declines monotonically or exhibits a peak around a vintage of 2 to 5 months before subsequently declining. This is the case even after controlling for heterogeneity. In the simulations of Section 2, both of these facts are produced when idiosyncratic shocks have a low persistence. I therefore interpret these empirical patterns as suggestive evidence that the idiosyncratic shocks faced by firms are transitory in nature.

4 Quantitative Model

Having documented the empirical behavior of vintage-conditional moments, I now assess the quantitative importance of using them for model estimation. To do so, I move to a simple general equilibrium setting. This is for two reasons. First, I want to assess the benefit of using these moments in a setting that more closely mirrors a typical quantitative model used in the literature. Second, I want to examine what the estimation procedure implies about the welfare cost of inflation, and as such need a setting where welfare is well defined.

4.1 Model setup.

The general equilibrium setup I use is essentially the canonical New Keynesian model (Galí, 2015), but with nominal frictions in the form of random menu costs.[15] An in-

[14] I test mean and variance because the unit level EDFs are observationally exponential or single peaked, and thus well-approximated by a negative binomial distribution. The negative binomial distribution is defined by two parameters which can be identified from the first and second moments.

[15] Similar random menu cost settings have been studied by Dotsey et al. (1999); Costain and Nakov (2011) and Dotsey and Wolman (2020).

finitely lived, representative household maximizes expected utility over consumption and labor

$$\mathbb{E}_t \left[\sum_{k=t}^{\infty} \beta^k \left(\log(C_{t+k}) - \alpha L_{t+k} \right) \right],$$

where C_t is a composite consumption good and L_t is labor supply. The composite good is a CES aggregate over individual differentiated goods C_{it}, $C_t = \left(\int C_{it}^{\frac{\nu-1}{\nu}} di \right)^{\frac{\nu}{\nu-1}}$, where ν is the elasticity of substitution. The household can purchase a one-period bond B_t at price Q_t, and is subject to per-period budget constraints,

$$\int P_{it} C_{it} di + Q_t B_t = W_t L_t + \Pi_t + B_{t-1},$$

where P_{it} are the prices of the individual goods, W_t is the nominal wage and Π_t is remitted firm profits.

Taking the prices of individual goods and the wage as given, household optimality conditions result in demand curves for individual products,

$$C_{it} = \left(\frac{P_{it}}{P_t} \right)^{-\nu} C_t, \quad P_t = \left(\int P_{it}^{1-\nu} di \right)^{\frac{1}{1-\nu}},$$

and intratemporal optimality condition $\alpha C_t = \frac{W_t}{P_t}$.

There is a continuum of firms, each of which operates a technology linear in labor to produce one individual product, $Y_{it} = A_{it} L_{it}$. Changing prices requires paying a random menu cost κ, measured in units of labor. Using the firm's production function and demand schedule, flow profits for firm i are

$$\Pi_{it} = P_{it} Y_{it} - W_t L_{it} - W_t \kappa_{it} I_{it} = \left(P_{it}^{1-\nu} - \frac{W_t}{A_{it}} P_{it}^{-\nu} \right) P_t^{\nu} C_t - W_t \kappa_{it} I_{it},$$

where I_{it} is an indicator variable taking value 1 when a firm changes its price.

Firms set prices to maximize the present value of real profits Π_{it}/P_t discounted at rate β. In the absence of menu costs, the firm's static, profit maximizing ideal price would be $P_{it}^* = \frac{\nu}{\nu-1} \frac{W_t}{A_{it}}$. The firm profit maximization problem can then be written in terms of ratios between set prices and ideal prices,

$$\max_{\{P_{it+k}\}} B \mathbb{E}_t \left[\sum_{k=t}^{\infty} \beta^k \left(A_{it+k}^{\nu-1} \left(\left(\frac{P_{it+k}}{P_{it+k}^*} \right)^{1-\nu} - \frac{\nu-1}{\nu} \left(\frac{P_{it+k}}{P_{it+k}^*} \right)^{-\nu} \right) - \hat{\kappa}_{it+k} I_{it+k} \right) \right],$$

where $B = (\frac{\nu}{\nu-1} \alpha \bar{C})^{1-\nu} \bar{C}$ and $\hat{\kappa} = (\frac{\nu}{\nu-1} \alpha \bar{C})^{\nu-1} \alpha \kappa$, and $\bar{C}$ is steady-state consumption.

I assume that firms' idiosyncratic productivity follows an AR(1) process in logs

$$\log(A_{it}) = \rho \log(A_{it-1}) + \sigma\varepsilon_{it}, \quad \varepsilon \sim N(0,1)$$

and that menu costs are distributed exponentially with a mass point at 0.[16]

$$\hat{\kappa} \sim H(k), \quad H(k) = \begin{cases} \kappa_0 & k = 0 \\ 1 - (1-\kappa_0)\exp\left(-\frac{k}{K}\right) & k > 0 \end{cases}$$

Finally, a monetary authority targets a constant rate of inflation, which manifests as the nominal wage W_t growing at a constant rate:

$$\log(W_t) = \pi + \log(W_{t-1}).$$

When writing the firm value function in its recursive form, the relation to the stylized model of Section 2 is plain to see. Defining $X_{it} = \frac{P_{it-1}}{P_t^*}$,

$$\begin{aligned} V(X, A, \kappa) = \max \Bigg\{ & A^{\nu-1}\left(X^{1-\nu} - \frac{\nu-1}{\nu}X^{-\nu}\right) + \beta\mathbb{E}\left[V\left(\frac{e^{-\pi}XA'}{A}, A', \kappa'\right)|X, A\right], \\ & -\hat{\kappa} + \max_{\bar{X}} A^{\nu-1}\left(\bar{X}^{1-\nu} - \frac{\nu-1}{\nu}\bar{X}^{-\nu}\right) + \beta\mathbb{E}\left[V\left(\frac{e^{-\pi}\bar{X}A'}{A}, A', \kappa'\right)|\bar{X}, A\right] \Bigg\}, \\ & \log(A') = \rho\log(A) + \sigma\varepsilon, \quad \varepsilon \sim N(0,1), \quad \hat{\kappa} \sim H(k). \end{aligned}$$

4.2 Calibration.

I calibrate this model to match features of the aggregate price data in the United Kingdom. There are eight parameters in the model. The first two, $\{\nu, \beta\}$, are governed by preferences and set to standard values of $\nu = 4$ and $\beta = 0.96^{\frac{1}{12}}$. I set $\alpha = 2.25$ to recover labor supply $L = .33$ in the frictionless steady state, and set $\pi = .0017$ to equal the observed average monthly inflation rate in the UK from 1996 to 2019. The final four parameters, $\{\rho, \sigma, K, \kappa_0\}$ are estimated within the model by matching moments.

I compare the results of two calibrations – one that targets moments from the cross-sectional distribution of price changes, and one that targets vintage-conditional moments. For the cross-sectional calibration, I target the frequency of price changes, the share of price changes that are decreases, the average absolute price change, and the variance of price changes. These are moments that have been used in the past to calibrate similar sticky price models (Nakamura and Steinsson, 2008; Costain

[16] This assumption produces a hazard function that is approximately quadratic in the log price gap, which is a good fit to the data (Alvarez et al., 2021; Gagliardone et al., 2025).

and Nakov, 2011).[17] For the vintage-conditional calibration, I target the vintage-conditional hazard rate and variance of price changes for vintages of 1 month and 10 months.[18] These are the moments that heuristically identify the persistence of idiosyncratic productivity shocks, and contain information about the overall frequency and variance of price changes. The parameters recovered by each calibration are reported in Table 2. Each model's fit to the targeted cross-sectional and vintage-conditional moments is presented in Table 3 and Figure 8, respectively. The full cross-sections of price changes relative to the data are reported in Appendix D.4.

Parameter	Description	Cross-Section Calibration	Vintage Calibration
ν	Elasticity of substitution	4	
β	Discount factor (monthly)	$0.96^{\frac{1}{12}}$	
α	Disutility of labor	2.25	
π	Monthly trend inflation	0.0017	
ρ	Autocorrelation of productivity shocks	.79	.15
σ	S.d. of innovation to productivity	.10	.14
K	Average non-zero menu cost	.61	.28
κ_0	Probability of free price adjustment	.012	.000

Table 2: Calibrated Parameters

The cross-sectional calibration recovers parameter values typical of menu cost models, although the standard deviation of shock innovations is somewhat high. In comparison, the vintage calibration recovers a significantly lower autocorrelation of idiosyncratic shocks, larger shock innovations and smaller menu costs.

The low autocorrelation of shocks is needed to match the downward sloping profiles of the vintage moments, and the other parameter values follow from there. When shocks are more transitory, it becomes less costly for a firm to maintain its original price in response to a shock. In order to match the level of the vintage-conditional hazard rates, it must also be less costly for a firm to change its price in response to a shock, so menu costs must be lower on average. Similarly, larger shock innovations are needed to match the vintage-conditional variances.

[17] Other candidate moments are percentiles of the price change distribution (Midrigan, 2011; Dotsey and Wolman, 2020), and the kurtosis of price changes (Alvarez and Lippi, 2014).

[18] For both sets of moments, I calculate aggregate statistics as follows. I first find the statistics within each product category. Then I take medians across product categories. This aggregation method follows Nakamura and Steinsson (2008).

Moment	Data	Cross-Section Calibration	Vintage Calibration
Freq of price changes	.088	.088	.079
Share of price decreases	.437	.435	.415
Avg absolute price change	.116	.136	.120
Var of price changes	.025	.023	.020
Untargeted Moments			
Kurt of price changes	5.286	2.105	2.757
Share of small (<.05) price changes	.317	.130	.131

Table 3: Cross-Sectional Moments

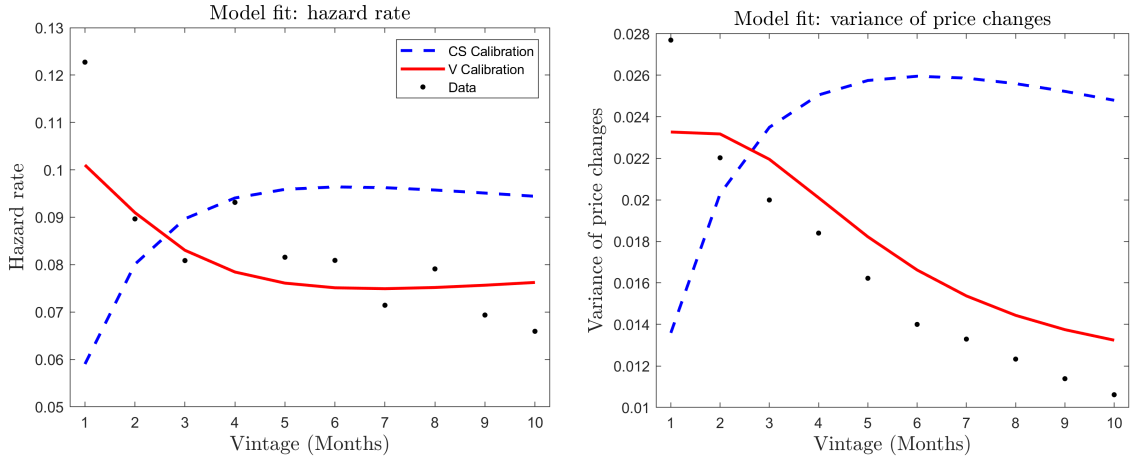


Figure 8: Vintage-Conditional Moments

These figures plot the model-implied profiles of vintage-conditional hazard rates (left) and vintage conditional variances (right) for the model calibrated to cross-sectional moments and the model calibrated to vintage-conditional moments.

Although neither calibrated model matches the targeted moments exactly, both fit the cross-sectional moments fairly well. However, the model calibrated to the cross-sectional data fares poorly at matching the vintage-conditional moments. This is unsurprising. The simulations in Section 2.2 indicate that vintage moments pin down the persistence of idiosyncratic shocks, and other parameters can be chosen to match the cross-section of price changes given that persistence. Furthermore, while vintage-conditional moments implicitly contain information about their cross-sectional analogues, the reverse is not true.

Conversely, it is surprising how well the vintage calibration can match the entire profiles of vintage moments with only four free parameters, especially as menu cost models typically have difficulty producing hazard rates that decline in vintage. Nakamura and Steinsson (2008) examine the behavior of the hazard rates in a fixed menu cost model where firms face both stationary idiosyncratic shocks and a non-stationary aggregate shock. They find that hazard rates decline in vintage only when the variance of idiosyncratic shocks becomes unreasonably large, but perform this

analysis assuming those idiosyncratic shocks are fairly persistent.[19] This calibration indicates that if the mean-reversion of idiosyncratic shocks is sufficiently low, one can match hazard rates declining in vintage with reasonably sized shock innovations.

4.3 Welfare results.

Given that these calibration procedures recover different parameters, it is natural to ask whether the model-implied costs of inflation differ as well. I measure welfare loss as the percentage change in consumption needed to equate the welfare of the household in the sticky price economy with that of the household in a corresponding flexible price economy. This is given by Λ such that

$$\log((1+\Lambda)\bar{C}) - \alpha\bar{L} = \log(\bar{C}_{flex}) - \alpha\bar{L}_{flex}$$

where $\bar{C}$ and $\bar{L}$ are equilibrium consumption and labor supply.[20] Note the lack of an expectation operator, as there is no aggregate uncertainty.

For each set of parameters estimated above, I solve the model varying the trend inflation rate, π, from 0% annually to 15% annually. I then compute the implied $\Lambda(\pi)$ at each inflation rate. These are plotted in Figure 9.

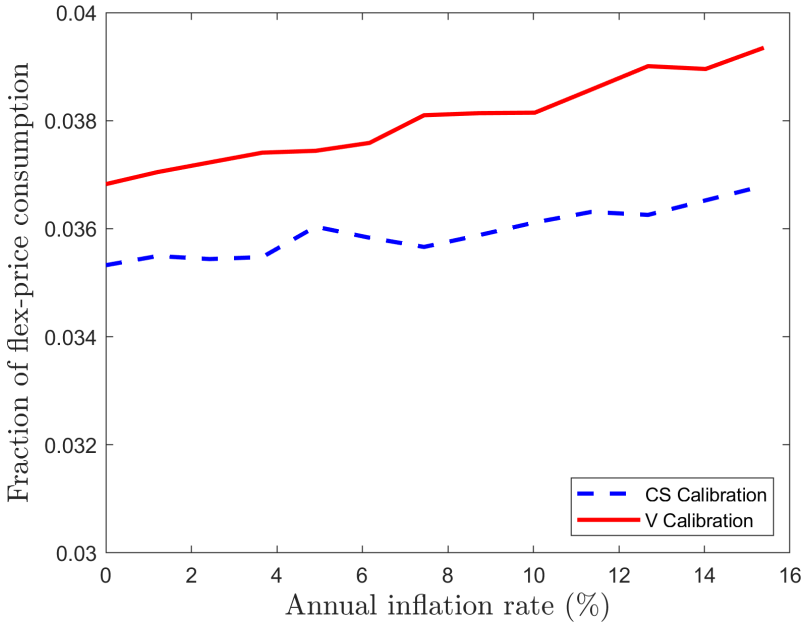


Figure 9: Welfare loss

This figure plots the model-implied welfare loss in consumption equivalent units at varying levels of steady-state inflation for the model calibrated to match cross-section (CS) moments and the model calibrated to match vintage (V) moments.

It is clear that when using parameters from the vintage calibration, welfare loss

[19] The AR parameter of their productivity process is fixed at $\rho = .66$.

[20] An exact expression for Λ in terms of model primitives and the steady-state distribution of price gaps can be found in Appendix A.3.

is higher and increases more quickly as steady-state inflation rises. At an annual inflation target of 2%, the model-implied welfare cost of inflation is 0.1 percentage points higher in the vintage calibration than in the cross-section calibration. Moving from 0% to 10% annual inflation increases welfare loss by about .17 percentage points in the vintage calibration, compared to a rise of only .09 percentage points in the cross-section calibration. In this region of low to moderate inflation, the slope of welfare loss with respect to steady-state inflation is about twice as steep in the vintage calibration.

Welfare loss in this economy comes from two sources. First, the gaps between actual prices and ideal prices lead to misallocation of productive labor across firms. This reduces the allocative efficiency of the economy compared to the flexible price benchmark. Second, labor is spent on paying menu costs to change prices, rather than producing goods for consumption. In Appendix D.4 I decompose welfare loss into the contribution made by each of these channels. The majority of the difference between the two calibrations stems from the allocative efficiency channel, rather than from the labor paid to menu costs.

In sticky price models with menu costs, the welfare cost of inflation is typically small and near-constant for low to moderate inflation rates (Nakamura et al., 2018; Alvarez et al., 2018). While both calibrations recover welfare losses that are small in magnitude, these losses are only near-constant in the cross-sectional calibration with persistent idiosyncratic shocks. In the vintage calibration with transitory idiosyncratic shocks, there is an appreciable rise in welfare loss as inflation increases.

In order to understand which parameters drive the difference in welfare results, I examine two intermediate calibrations - one in which productivity process parameters (ρ and σ) come from the cross-sectional calibration while the menu cost parameters (K and κ) come from the vintage calibration, and one with the opposite. As above, I compute welfare loss for each of these cases at a variety of steady-state inflation rates. The levels and slopes of welfare loss are reported in Table 4.

Taking the cross-sectional calibration as a baseline and comparing it to the calibration with menu cost parameters from the vintage calibration (Column 2), I find that welfare loss becomes both smaller and less sensitive to changes in the inflation rate. This is because menu costs are smaller on average in the vintage calibration, so this intermediate step just moves the model closer to the frictionless case. Comparing the baseline to the calibration with productivity shock parameters from the vintage calibration (Column 3), welfare loss is higher and much more sensitive to changes in the inflation rate. With larger, more transitory shocks, firms understand that their productivity is mean reverting and a given deviation from optimal price is therefore less costly long term. Firms are willing to tolerate a greater degree of mispricing in any period, knowing their price gap will likely mean-revert and avoiding the risk of

making two costly price changes. As a result, the level of welfare loss is higher.

	(1) Cross-Section Calibration	(2) Vintage MC Params	(3) Vintage Prod Params	(4) Vintage Calibration
Welfare loss at $\pi = 2\%$	3.549	2.945	3.966	3.705
Δ Welfare, $\pi = 0\%$ to $\pi = 10\%$	.099	.035	.266	.175

Table 4: Welfare loss comparison

The connection between the persistence of shocks and the slope of welfare loss is more subtle. It is helpful to first recall the basic intuition behind the interaction of welfare loss, idiosyncratic shocks, and the inflation rate in menu cost models. When the inflation rate is near zero, pricing behavior is dominated by idiosyncratic shocks. An increase in inflation thus has little effect on firm behavior or welfare loss. Conversely, when the inflation rate is extremely high, pricing behavior is dominated by the inflation rate and it is as if there are no idiosyncratic shocks. Welfare loss increases at a constant rate with respect to inflation.[21] Features of the shock process thus govern the effect of a rise in inflation on welfare for intermediate levels of inflation. In particular, the persistence of shocks determines the degree to which those shocks dominate firm pricing behavior as trend inflation rises.

When shocks are persistent, future shocks are near the current shock in expectation. A price changing firm thus has a strong incentive to adjust to its current shock, while limiting its response to trend inflation. Furthermore, because trend inflation also behaves as a permanent shock, firms are able to slightly adjust to inflation without inducing a larger price gap when future shocks arrive. As a result, shocks continue to dominate firm pricing behavior for a greater range of inflation rates.

However, when shocks are transitory, firms expect future shocks to mean revert. A price changing firm thus has relatively less incentive to fully adjust to the current shock it faces, and relatively more incentive to adjust to trend inflation, which accumulates permanently. Trend inflation thus begins to alter firm pricing behavior, and with it cause welfare loss, at lower levels. This causes welfare loss to rise more quickly with inflation in a calibration with transitory shocks.

5 Conclusion

This paper argues that vintage-conditional moments should be matched when estimating sticky price models. I demonstrate in a simple example that these moments identify the persistence of idiosyncratic shocks to firms' ideal prices, and

[21] These intuitions are formalized by Alvarez et al. (2018) using a continuous time, fixed menu cost model with permanent idiosyncratic shocks.

derive a heuristic to determine whether those shocks are permanent or transitory. Using the price micro-data underlying the UK CPI, I document that both hazard rates and variances of price changes decline in vintage. This indicates that idiosyncratic shocks are relatively transitory. Compared to a random menu-cost model estimated using standard, cross-sectional moments, the same model estimated to match vintage-conditional moments recovers a lower persistence of idiosyncratic shocks and a higher welfare cost of inflation. Additionally, the cost of inflation rises substantially more rapidly as inflation rises. The majority of the divergence between the two model calibrations stems from the difference in estimated ideal price processes, rather than the difference in estimated menu costs.

6 References

Adam, Klaus and Henning Weber, "Estimating the Optimal Inflation Target from Trends in Relative Prices," *American Economic Journal: Macroeconomics*, July 2023, *15* (3), 1–42.

__ , Andrey Alexandrov, and Henning Weber, "Inflation Distorts Relative Prices: Theory and Evidence," Technical Report 18088, CEPR Press April 2023.

Alvarez, F. E., F. Lippi, and L. Paciello, "Optimal Price Setting With Observation and Menu Costs," *The Quarterly Journal of Economics*, November 2011, *126* (4), 1909–1960.

Alvarez, Fernando and Francesco Lippi, "Price Setting With Menu Cost for Multiproduct Firms," *Econometrica*, 2014, *82* (1), 89–135.

__ , __ , and Aleksei Oskolkov, "The Macroeconomics of Sticky Prices with Generalized Hazard Functions," *The Quarterly Journal of Economics*, November 2021, *137* (2), 989–1038.

__ , Hervé Le Bihan, and Francesco Lippi, "The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach," *American Economic Review*, October 2016, *106* (10), 2817–2851.

__ , Katarína Borovičková, and Robert Shimer, "Consistent Evidence on Duration Dependence of Price Changes," *American Economic Review*, October 2025, *115* (10), 3322–3366.

__ , Martin Beraja, Martín Gonzalez-Rozada, and Pablo Andrés Neumeyer, "From Hyperinflation to Stable Prices: Argentina's Evidence on Menu Cost Models," *The Quarterly Journal of Economics*, September 2018, *134* (1), 451–505.

Baley, Isaac and Andrés Blanco, "Firm Uncertainty Cycles and the Propagation of Nominal Shocks," *American Economic Journal: Macroeconomics*, January 2019, *11* (1), 276–337.

__ and __ , "Aggregate Dynamics in Lumpy Economies," *Econometrica*, 2021, *89* (3), 1235–1264.

Berger, David and Joseph Vavra, "Dynamics of the U.S. price distribution," *European Economic Review*, April 2018, *103*, 60–82.

Bunn, Philip and Colin Ellis, "Examining the Behaviour of Individual UK Consumer Prices," *The Economic Journal*, January 2012, *122* (558), F35–F55.

__ and __ , "How Do Individual UK Producer Prices Behave?," *The Economic Journal*, January 2012, *122* (558), F16–F34.

Caballero, R. J. and E. M. R. A. Engel, "Microeconomic Adjustment Hazards and Aggregate Dynamics," *The Quarterly Journal of Economics*, May 1993, *108* (2), 359–383.

Campbell, Jeffrey R. and Benjamin Eden, "Rigid Prices: Evidence from U.S. Scanner Data," *International Economic Review*, April 2014, *55* (2), 423–442.

Cavallo, Alberto and Roberto Rigobon, *The Distribution of the Size of Price Changes* February 2011.

Costain, James and Anton Nakov, "Price Adjustments in a General Model of State-Dependent Pricing," *Journal of Money, Credit and Banking*, March 2011, *43* (2–3), 385–406.

Dotsey, M., R. G. King, and A. L. Wolman, "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *The Quarterly Journal of Economics*, May 1999, *114* (2), 655–690.

Dotsey, Michael and Alexander L. Wolman, "Investigating Nonneutrality in a State-Dependent Pricing Model with Firm-Level Productivity Shocks," *International Economic Review*, January 2020, *61* (1), 159–188.

Gagliardone, Luca, Mark Gertler, Simone Lenzu, and Joris Tielens, *Micro and Macro Cost-Price Dynamics in Normal Times and During Inflation Surges* February 2025.

Gagnon, Etienne, "Price Setting During Low and High Inflation: Evidence from Mexico," *Quarterly Journal of Economics*, 2009.

Galí, Jordi, *Monetary Policy, Inflation, and the Business Cycle*, second ed., Princeton University Press, 2015.

Gautier, Erwan, Cristina Conflitti, Riemer P. Faber, Brian Fabo, Ludmila Fadejeva, Valentin Jouvanceau, Jan-Oliver Menz, Teresa Messner, Pavlos Petroulas, Pau Roldan-Blanco, Fabio Rumler, Sergio Santoro, Elisabeth Wieland, and Hélène Zimmer, "New Facts on Consumer Price Rigidity in the Euro Area," *American Economic Journal: Macroeconomics*, October 2024, *16* (4), 386–431.

Golosov, Mikhail and Robert E. Lucas, "Menu Costs and Phillips Curves," *Journal of Political Economy*, April 2007, *115* (2), 171–199.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz, "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data," *American Economic Review*, February 2021, *111* (2), 428–471.

Karadi, Peter, Juergen Amann, Javier Sánchez Bachiller, Pascal Seiler, and Jesse Wursten, "Price setting on the two sides of the Atlantic - Evidence from supermarket scanner data," *Journal of Monetary Economics*, November 2023, *140*, S1–S17.

Kehoe, Patrick and Virgiliu Midrigan, "Prices are sticky after all," *Journal of Monetary Economics*, October 2015, *75*, 35–53.

Klenow, Peter J. and Benjamin A. Malin, "Microeconomic Evidence on Price-Setting," in Benjamin M. Friedman and Michael Woodford, eds., *Handbook of Monetary Economics*, Vol. 3, Elsevier, 2010, pp. 231–284.

__ **and Oleksiy Kryvtsov**, "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?," *Quarterly Journal of Economics*, August 2008, *123* (3), 863–904.

Lach, Saul and Daniel Tsiddon, "The Behavior of Prices and Inflation: An Empirical Analysis of Disaggregated Price Data," *Journal of Political Economy*, April 1992, *100* (2), 349–389.

Midrigan, Virgiliu, "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations," *Econometrica*, July 2011, *79* (4), 1139–1180.

Nakamura, Emi and Jón Steinsson, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, November 2008, *123* (4), 1415–1464.

__ **and** __ , "Monetary Non-neutrality in a Multisector Menu Cost Model," *The Quarterly Journal of Economics*, August 2010, *125* (3), 961–1013.

__ **,** __ **, Patrick Sun, and Daniel Villar**, "The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation," *The Quarterly Journal of Economics*, August 2018, *133* (4), 1933–1980.

Office for National Statistics, "Consumer Prices Indices Technical Manual, 2019," Technical Report, United Kingdom Office for National Statistics 2024.

Sheremirov, Viacheslav, "Price dispersion and inflation: New facts and theoretical implications," *Journal of Monetary Economics*, October 2020, *114*, 59–70.

Vavra, Joseph, "The Empirical Price Duration Distribution and Monetary Non-Neutrality," 2010.

__ , "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *The Quarterly Journal of Economics*, September 2013, *129* (1), 215–258.

A Theoretical Results

A.1 Simple Example: Permanent vs Transitory Shocks

This section presents proofs underlying the propositions given in Section 2.2.

Setup. Starting with the model presented in Section 2, I assume the following:

- $B = 1$, $\pi = 0$, and $\beta = 0$.
- $G(a'|a) = \Phi\left(\frac{a' - \rho a}{\sigma}\right)$.
- $H(\kappa) = 1 - \exp\left(-\frac{\kappa}{2K}\right)$.

Finally, I assume that the economy is either a permanent shocks economy, where $\rho = 1$, or a transitory shocks economy, where $\rho = 0$.

With these assumptions, the firm value function is

$$V(x_{it}, a_{it}, \kappa_{it}) = \min\left\{x_{it}^2,\ \kappa_{it} + \min_{\bar{x}} \bar{x}^2\right\}.$$

The optimal reset price gap is clearly $\bar{x} = 0$, and I can then derive hazard function

$$\Lambda(x_{it}, a_{it}) = H(x_{it}^2) = 1 - \exp\left(-\frac{x_{it}^2}{2K}\right).$$

Proof of Proposition 1

In the permanent shocks economy, the stationary distribution is defined by:

$$f_p(x) = e^{-\frac{x^2}{2K}} \hat{f}_p(x) + \delta_0(x) \int_{-\infty}^{\infty} (1 - e^{-\frac{y^2}{2K}}) \hat{f}_p(y) dy$$
$$\hat{f}_p(x) = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-z)^2}{2\sigma^2}\right) f_p(z) dz$$

By guessing and verifying, the following equations in terms of $\{\sigma, K\}$ satisfy these conditions:

$$f_p(x) = \sqrt{2\pi} c_0 \delta_0(x) + \sum_{i=1}^{\infty} \frac{c_i}{\sqrt{k_i}} \exp\left(-\frac{x^2}{2k_i}\right)$$
$$\hat{f}_p(x) = \sum_{i=0}^{\infty} \frac{c_i}{\sqrt{\sigma^2 + k_i}} \exp\left(-\frac{x^2}{2(\sigma^2 + k_i)}\right)$$
$$k_0 = 0, \quad k_i = \frac{K(\sigma^2 + k_{i-1})}{K + \sigma^2 + k_{i-1}},$$
$$c_0 = \frac{1}{\sqrt{2\pi}} - \sum_{i=1}^{\infty} c_i, \quad c_i = c_{i-1} \sqrt{\frac{K}{K + \sigma^2 + k_{i-1}}}$$

Given initial point $k_0 = 0$, the sequence $\{k_i\}$ is monotonically increasing and bounded above by $\bar{k} = \frac{\sqrt{\sigma^4+4K\sigma^2}-\sigma^2}{2}$, and sequence $\{c_i\}$ is monotonically decreasing and bounded below by 0. This assures that all infinite sums are converging, and that c_0 is determined by σ and K.

The model-implied distribution of price changes is

$$q_p(\Delta p) = \frac{1}{\sqrt{2\pi}c_0}(1 - e^{-\frac{\Delta p^2}{2K}})\hat{f}_p(\Delta p).$$

From this, the model-implied hazard rate and variance of price changes are

$$\begin{aligned}
\lambda_p &= \int_{-\infty}^{\infty} \Lambda_p(x)\hat{f}_p(x)dx = \int_{-\infty}^{\infty}(1 - e^{\frac{-x^2}{2K}})\hat{f}_p(x)dx \\
&= \sum_{i=0}^{\infty} \frac{c_i}{\sqrt{\sigma^2+k_i}} \int_{-\infty}^{\infty}(1 - e^{\frac{-x^2}{2K}})e^{-\frac{x^2}{2(\sigma^2+k_i)}}dx \\
&= \sum_{i=0}^{\infty} \frac{c_i\sqrt{2\pi}}{\sqrt{\sigma^2+k_i}} \left(\sqrt{\sigma^2+k_i} - \sqrt{\frac{K(\sigma^2+k_i)}{K+\sigma^2+k_i}}\right) \\
&= \sqrt{2\pi}\sum_{i=0}^{\infty}(c_i - c_{i+1}) = \sqrt{2\pi}c_0
\end{aligned}$$

$$\begin{aligned}
Var_p(\Delta p) &= \int_{-\infty}^{\infty} \frac{\Delta p^2}{\sqrt{2\pi}c_0}(1 - e^{\frac{-\Delta p^2}{2K}})\hat{f}_p(\Delta p)d\Delta p \\
&= \sum_{i=0}^{\infty} \frac{c_i}{\sqrt{2\pi}c_0\sqrt{\sigma^2+k_i}} \int_{-\infty}^{\infty} \Delta p^2(1 - e^{\frac{-\Delta p^2}{2K}})e^{-\frac{\Delta p^2}{2(\sigma^2+k_i)}}d\Delta p \\
&= \sum_{i=0}^{\infty} \frac{c_i}{c_0\sqrt{\sigma^2+k_i}} \left((\sigma^2+k_i)^{\frac{3}{2}} - (\frac{K(\sigma^2+k_i)}{K+\sigma^2+k_i})^{\frac{3}{2}}\right) \\
&= \frac{1}{c_0}\sum_{i=0}^{\infty} \left(c_i\sigma^2 + c_ik_i - c_i\sqrt{\frac{K}{K+\sigma^2+k_i}}(\frac{K(\sigma^2+k_i)}{K+\sigma^2+k_i})\right) \\
&= \frac{1}{c_0}\sum_{i=0}^{\infty} \left(c_i\sigma^2 + c_ik_i - c_{i+1}k_{i+1}\right) = \frac{1}{c_0}\sum_{i=0}^{\infty} c_i\sigma^2 = \frac{\sigma^2}{\sqrt{2\pi}c_0}
\end{aligned}$$

Taking these together, in order to match data $\hat{\lambda}$, $\widehat{Var}(\Delta p)$ I can set $\sigma_p^2 = \hat{\lambda}\widehat{Var}(\Delta p)$ and K_p such that $c_0 = \frac{\hat{\lambda}}{\sqrt{2\pi}}$, given σ_p^2. To see the existence of such a $K_p > 0$, observe that the condition on c_0 can be written, using the definitions above, as

$$\begin{aligned}
\frac{\hat{\lambda}}{\sqrt{2\pi}} = c_0 = \frac{1}{\sqrt{2\pi}} - \sum_{i=1}^{\infty} \frac{\hat{\lambda}}{\sqrt{2\pi}} \prod_{j=0}^{i-1} \sqrt{\frac{K}{K+\sigma^2+k_j}} \\
\implies \sum_{i=1}^{\infty} \prod_{j=0}^{i-1} \sqrt{\frac{K}{K+\sigma^2+k_j}} = \frac{1-\hat{\lambda}}{\hat{\lambda}}
\end{aligned}$$

The LHS of this expression is a strictly increasing continuous function, which is 0 for $K = 0$ and approaching infinity as $K \to \infty$. So, by the intermediate value theorem, for a given $\hat{\lambda} \in (0, 1)$ there exists $K_p > 0$ such that this equation is satisfied, and this K_p is unique. This gives us a unique pair $\{\sigma_p, K_p\}$ that rationalizes the data in the permanent shocks economy.

In the transitory shocks economy, the stationary distribution is defined by:

$$f_t(x,a) = e^{-\frac{x^2}{2K}}\hat{f}(x,a) + \delta_0(x)\int(1 - e^{-\frac{y^2}{2K}})\hat{f}_t(y,a)dy$$
$$\hat{f}_t(x,a) = \frac{1}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{a^2}{2\sigma^2}\right)\int f_t(x-z+a,z)dz$$

By guessing and verifying, the following equations in terms of $\{\sigma, K\}$ satisfy these conditions:

$$\begin{aligned} f_t(x,a) &= \frac{1}{2\pi\sigma^2}\exp\left(-\frac{(x+a)^2+a^2}{2\sigma^2} - \frac{x^2}{2K}\right) \\ &\quad + \delta_0(x)\frac{\exp\left(-\frac{a^2}{2\sigma^2}\right)}{\sqrt{2\pi\sigma^2}}\left(1 - \sqrt{\frac{K}{K+\sigma^2}}\exp\left(-\frac{a^2}{2(K+\sigma^2)}\right)\right) \\ \hat{f}_t(x,a) &= \frac{1}{2\pi\sigma^2}\exp\left(-\frac{(x+a)^2+a^2}{2\sigma^2}\right) \end{aligned}$$

The model-implied distribution of price changes is

$$q_t(\Delta p) = \frac{(1 - e^{-\frac{\Delta p^2}{2K}})}{1 - \sqrt{\frac{K}{K+2\sigma^2}}}\int \hat{f}_t(-\Delta p, a)da$$

From this, the model-implied hazard rate and variance of price changes are

$$\begin{aligned} \lambda_t &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\Lambda_t(x)\hat{f}_t(x,a)da\ dx \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}(1 - e^{\frac{-x^2}{2K}})\frac{1}{2\pi\sigma^2}e^{-\frac{(x+a)^2+a^2}{2\sigma^2}}da\ dx \\ &= \int_{-\infty}^{\infty}\frac{1}{\sqrt{2\pi\sigma^2}}(1 - e^{\frac{-x^2}{2K}})e^{-\frac{x^2}{4\sigma^2}}dx \\ &= 1 - \sqrt{\frac{K}{K+2\sigma^2}} \end{aligned}$$

$$
\begin{aligned}
Var_t(\Delta p) &= \int_{-\infty}^{\infty} \Delta p^2 \frac{(1-e^{-\frac{\Delta p^2}{2K}})}{1-\sqrt{\frac{K}{K+2\sigma^2}}} \int_{-\infty}^{\infty} \hat{f}_t(-\Delta p, a) da \, d\Delta p \\
&= \frac{1}{1-\sqrt{\frac{K}{K+2\sigma^2}}} \int_{-\infty}^{\infty} \Delta p^2 (1-e^{-\frac{\Delta p^2}{2K}}) e^{-\frac{\Delta p^2}{4\sigma^2}} \, d\Delta p \\
&= 2\sigma^2 \frac{1-\left(\frac{K}{K+2\sigma^2}\right)^{\frac{3}{2}}}{1-\sqrt{\frac{K}{K+2\sigma^2}}}
\end{aligned}
$$

Taking these together, in order to match data $\hat{\lambda}, \widehat{Var}(\Delta p)$ I can set $\sigma_t^2 = \frac{Var(\Delta p)}{2(\lambda^2-3\lambda+3)}$ and $K_t = \frac{2(1-\lambda)^2\sigma_t^2}{\lambda(2-\lambda)}$.

Proof of Proposition 2

In the permanent shocks economy, as all firms reset to $\bar{x} = 0$, we have initial distribution $f_{p,0}(x) = \delta_0(x)$.

Applying the definitions of $f_{p,\tau}$ and $\hat{f}_{p,\tau}$ recursively for $\tau \geq 1$, we have

$$
f_{p,\tau}(x) = \frac{e^{-\frac{x^2}{2\sigma_\tau^2}}}{\sqrt{2\pi\sigma_\tau^2}}, \quad \hat{f}_{p,\tau}(x) = \frac{e^{-\frac{x^2}{2(\sigma_{\tau-1}^2+\sigma^2)}}}{\sqrt{2\pi(\sigma_{\tau-1}^2+\sigma^2)}}
$$

with

$$
\sigma_0^2 = 0, \quad \sigma_\tau^2 = \frac{K(\sigma_{\tau-1}^2+\sigma^2)}{\sigma_{\tau-1}^2+\sigma^2+K}.
$$

The conditional distribution of price changes, conditional hazard rate and conditional variance are:

$$
q_{p,\tau}(\Delta p) = \frac{\sqrt{K+\sigma_{\tau-1}^2+\sigma^2}}{\sqrt{K+\sigma_{\tau-1}^2+\sigma^2}-\sqrt{K}}(1-e^{-\frac{\Delta p^2}{2K}}) \frac{e^{-\frac{\Delta p^2}{2(\sigma_{\tau-1}^2+\sigma^2)}}}{\sqrt{2\pi(\sigma_{\tau-1}^2+\sigma^2)}},
$$

$$
\lambda_{p,\tau} = 1-\sqrt{\frac{K}{K+\sigma_{\tau-1}^2+\sigma^2}}, \quad Var_{p,\tau}(\Delta p) = \frac{\sigma_{\tau-1}^2+\sigma^2}{\lambda_{p,\tau}} \left(1-(\frac{K}{K+\sigma_{\tau-1}^2+\sigma^2})^{\frac{3}{2}}\right).
$$

Observe first that $\sigma_\tau \in [0, \frac{\sqrt{\sigma^4+4K\sigma^2}-\sigma^2}{2})$ and is monotonically increasing in τ.

Given this, we have

$$\frac{\partial \lambda_{p,\tau}}{\partial \sigma^2_{\tau-1}} = \frac{1}{2K}\left(\frac{K}{K+\sigma^2_{\tau-1}+\sigma^2}\right)^{\frac{3}{2}} > 0$$

$$\frac{\partial Var_{p,\tau}(\Delta p)}{\partial \sigma^2_{\tau-1}} = \left(1+\frac{1}{2}\sqrt{\frac{K}{K+\sigma^2_{\tau-1}+\sigma^2}}\right)\left(1-(\frac{K}{K+\sigma^2_{\tau-1}+\sigma^2})^{\frac{3}{2}}\right) + \frac{\sigma^2_{\tau-1}+\sigma^2}{\lambda_{p,\tau}}\frac{3}{2K}(\frac{K}{K+\sigma^2_{\tau-1}+\sigma^2})^{\frac{5}{2}} > 0$$

Thus, $\lambda_{p,\tau}$ and $Var_{p,\tau}(\Delta p)$ are both monotonically increasing in τ.

In the transitory shocks case, as all firms reset to $\bar{x} = 0$, we have initial distribution

$$f_{t,0}(x,a) = \frac{\delta_0(x)}{1-\sqrt{\frac{K}{K+2\sigma^2}}}\frac{\exp\left(-\frac{a^2}{2\sigma^2}\right)}{\sqrt{2\pi\sigma^2}}\left(1-\sqrt{\frac{K}{K+\sigma^2}}\exp\left(-\frac{a^2}{2(K+\sigma^2)}\right)\right).$$

Applying the definitions of $f_{t,\tau}$ and $\hat{f}_{t,\tau}$ recursively for $\tau \geq 1$, we have

$$f_{t,\tau}(x,a) = w_\tau \frac{\exp\left(-\frac{1}{2}\begin{bmatrix} x & a\end{bmatrix}\Sigma^{-1}_{\tau-1}\begin{bmatrix} x \\ a\end{bmatrix}\right)}{2\pi\sqrt{|\Sigma_{\tau-1}|}} + (1-w_\tau)\frac{\exp\left(-\frac{1}{2}\begin{bmatrix} x & a\end{bmatrix}\Sigma^{-1}_{\tau}\begin{bmatrix} x \\ a\end{bmatrix}\right)}{2\pi\sqrt{|\Sigma_{\tau}|}}$$

$$\hat{f}_{t,\tau}(x,a) = w_{\tau-1} \frac{\exp\left(-\frac{1}{2}\begin{bmatrix} x & a\end{bmatrix}\hat{\Sigma}^{-1}_{\tau-1}\begin{bmatrix} x \\ a\end{bmatrix}\right)}{2\pi\sqrt{|\hat{\Sigma}_{\tau-1}|}} + (1-w_{\tau-1})\frac{\exp\left(-\frac{1}{2}\begin{bmatrix} x & a\end{bmatrix}\hat{\Sigma}^{-1}_{\tau}\begin{bmatrix} x \\ a\end{bmatrix}\right)}{2\pi\sqrt{|\hat{\Sigma}_{\tau}|}}$$

with

$$\Sigma_\tau = \frac{K\sigma^2}{K+\sigma^2}\begin{bmatrix} \frac{2K+(\tau+2)\sigma^2}{K+(\tau+2)\sigma^2} & -\frac{K+(\tau+1)\sigma^2}{K+(\tau+2)\sigma^2} \\ -\frac{K+(\tau+1)\sigma^2}{K+(\tau+2)\sigma^2} & \frac{K^2+(\tau+2)K\sigma^2+\sigma^4}{K^2+(\tau+2)K\sigma^2}\end{bmatrix}, \quad \hat{\Sigma}_\tau = \sigma^2\begin{bmatrix} \frac{2K+(\tau+2)\sigma^2}{K+(\tau+1)\sigma^2} & -1 \\ -1 & 1\end{bmatrix},$$

$$w_\tau = \frac{\sqrt{\frac{K+\sigma^2}{K+(\tau+1)\sigma^2}}}{\sqrt{\frac{K+\sigma^2}{K+(\tau+1)\sigma^2}}-\sqrt{\frac{K}{K+(\tau+2)\sigma^2}}}, \quad \sigma^2_\tau = \frac{\sigma^2(2K+(\tau+2)\sigma^2)}{K+(\tau+1)\sigma^2}.$$

The conditional distribution of price changes, conditional hazard rate and condi-

tional variance are:

$$q_{t,\tau}(\Delta p) = \frac{(1 - e^{-\frac{\Delta p^2}{2K}})}{\lambda_{t,\tau}} \left(\frac{w_{\tau-1} e^{-\frac{\Delta p^2}{2\sigma^2_{\tau-1}}}}{\sqrt{2\pi\sigma^2_{\tau-1}}} + \frac{(1 - w_{\tau-1}) e^{-\frac{\Delta p^2}{2\sigma^2_{\tau}}}}{\sqrt{2\pi\sigma^2_{\tau}}} \right),$$

$$\lambda_{t,\tau} = 1 - w_{\tau-1}\sqrt{\frac{K}{K + \sigma^2_{\tau-1}}} - (1 - w_{\tau-1})\sqrt{\frac{K}{K + \sigma^2_{\tau}}},$$

$$Var_{t,\tau}(\Delta p) = \frac{1}{\lambda_\tau} \left(w_{\tau-1}\sigma^2_{\tau-1} \left(1 - (\frac{K}{K + \sigma^2_{\tau-1}})^{\frac{3}{2}} \right) + (1 - w_{\tau-1})\sigma^2_{\tau} \left(1 - (\frac{K}{K + \sigma^2_{\tau}})^{\frac{3}{2}} \right) \right)$$

Note that, treating σ_τ and w_τ as continuous functions of τ, we can write

$$\sigma^2_\tau = \sigma^2 \left(1 + \frac{K + \sigma^2}{K + (\tau + 1)\sigma^2} \right), \quad w_\tau = \frac{1}{1 - \sqrt{\frac{K}{K + \sigma^2_\tau}}}$$

$$\implies \frac{\partial \sigma^2_\tau}{\partial \tau} = -\frac{\sigma^4(K + \sigma^2)}{(K + (\tau + 1)\sigma^2)^2}, \quad \frac{\partial^2 \sigma^2_\tau}{\partial \tau^2} = \frac{2\sigma^6(K + \sigma^2)}{(K + (\tau + 1)\sigma^2)^3}$$

$$\frac{\partial w_\tau}{\partial \tau} = -\left(\frac{K}{K + \sigma^2_\tau} \right)^{\frac{3}{2}} \frac{w^2_\tau \frac{\partial \sigma^2_\tau}{\partial \tau}}{2K}, \quad \frac{\partial^2 w_\tau}{\partial \tau^2} = \sqrt{\frac{K}{K + \sigma^2_\tau}} \frac{w^2_\tau \left((2w_\tau + 1)(\frac{\partial \sigma^2_\tau}{\partial \tau})^2 - 2(K + \sigma^2_\tau)\frac{\partial^2 \sigma^2_\tau}{\partial \tau^2} \right)}{4(K + \sigma^2_\tau)^2}$$

From this, some important properties are clear. First, σ^2_τ is strictly decreasing and convex in τ, with upper bound $2\sigma^2$ and lower bound σ^2. Second, w_τ is strictly increasing in τ with upper bound $\frac{1}{1-\sqrt{\frac{K}{K+\sigma^2}}}$ and lower bound $\frac{1}{1-\sqrt{\frac{K}{K+2\sigma^2}}}$. Less clear, but also important, is that w_τ is concave in τ. This can be shown algebraically and by noting that w_τ is bounded above by $\frac{1}{1-\sqrt{\frac{K}{K+\sigma^2}}}$.

From the concavity of w_τ, we have that

$$\frac{w_{\tau+1} + w_{\tau-1}}{2} < w_\tau \implies w_{\tau+1} - w_\tau < w_\tau - w_{\tau-1}$$

so $(w_\tau - w_{\tau-1})$ is decreasing in τ.

Thus,

$$\lambda_{t,\tau} = 1 - w_{\tau-1}\frac{w_{\tau-1} - 1}{w_{\tau-1}} - (1 - w_{\tau-1})\frac{w_\tau - 1}{w_\tau} = \frac{1 + w_\tau - w_{\tau-1}}{w_\tau}$$

As w_τ is increasing in τ and $(w_\tau - w_{\tau-1})$ is decreasing in τ, we have that $\lambda_{t,\tau}$ is decreasing in τ.

Furthermore,

$$Var_{t,\tau}(\Delta p) = \frac{1}{\lambda_{t,\tau}}\left(w_{\tau-1}\sigma^2_{\tau-1}(1-(\frac{w_{\tau-1}-1}{w_{\tau-1}})^3) + (1-w_{\tau-1})\sigma^2_{\tau}(1-(\frac{w_{\tau}-1}{w_{\tau}})^3)\right)$$
$$\frac{w_{\tau}\sigma^2_{\tau-1}(\frac{1-3w_{\tau-1}+3w^2_{\tau-1}}{w^2_{\tau-1}}) + (1-w_{\tau-1})\sigma^2_{\tau}(\frac{1-3w_{\tau}+3w^2_{\tau}}{w^2_{\tau}})}{1+w_{\tau}-w_{\tau-1}}$$

Defining $S_\tau = \sigma^2_\tau(\frac{1-3w_\tau+3w^2_\tau}{w^2_\tau})$, $A_\tau = w_\tau$ and $B_\tau = w_{\tau-1}-1$, I can write

$$Var_{t,\tau}(\Delta p) = \frac{A_\tau S_{\tau-1} - B_\tau S_\tau}{A_\tau - B_\tau}$$

The condition $Var_{t,\tau}(\Delta p) > Var_{t,\tau+1}(\Delta p)$ is then satisfied if and only if

$$\frac{A_\tau(A_{\tau+1}-B_{\tau+1})}{A_\tau A_{\tau+1} - B_\tau B_{\tau+1}}S_{\tau-1} + \frac{(A_\tau - B_\tau)B_{\tau+1}}{A_\tau A_{\tau+1} - B_\tau B_{\tau+1}}S_{\tau+1} > S_\tau$$

As A_τ and B_τ are both positive, and $A_\tau > B_\tau$, this condition is satisfied if S_τ is convex in τ. I can write S_τ as

$$S_\tau = \sigma^2_\tau\frac{1}{w^2_\tau}(1-3w_\tau) + 3\sigma^2_\tau$$

which, as σ^2_τ is convex and w^2_τ is positive and concave, is also convex. Therefore, the condition above is satisfied and $Var_{t,\tau}(\Delta p)$ is decreasing in τ.

A.2 GE Model

Setup. From Section 4, in the quantitative model the consumer problem is

$$\max_{\{C_{it+k},L_{t+k},B_{t+k}\}} \mathbb{E}_t\left[\sum_{k=t}^{\infty}\beta^k\left(\log(C_{t+k}) - \alpha L_{t+k}\right)\right]$$
$$\text{s.t. } \int P_{it}C_{it}di + Q_tB_t = W_tL_t + \Pi_t + B_{t-1}$$
$$C_t = \left(\int C_{it}^{\frac{\nu-1}{\nu}}di\right)^{\frac{\nu}{\nu-1}}$$

This results in first order conditions

$$\frac{\beta^t}{C_t} = \mu_t$$
$$\lambda_t P_{it} = \mu_t C_t^{\frac{1}{\nu}} C_{it}^{-\frac{1}{\nu}}$$
$$\beta^t\alpha = \lambda_t W_t$$
$$\lambda_t Q_t = \lambda_{t+1}$$

where λ_t is the Lagrange multiplier on the period t budget constraint, and μ_t is the Lagrange multiplier on the consumption aggregator. Defining $P_t = \left(\int P_{it}^{1-\nu} di\right)^{\frac{1}{1-\nu}}$ as the price of one unit of aggregate consumption, I have the following optimality conditions:

$$\alpha C_t = \frac{W_t}{P_t}, \quad \frac{C_{it}}{C_t} = \left(\frac{P_{it}}{P_t}\right)^{-\nu}, \quad \frac{\beta}{Q_t}\frac{P_t C_t}{P_{t+1} C_{t+1}} = 1$$

There is a continuum of firms, each of which operates a technology linear in labor to produce one individual product, $Y_{it} = A_{it} L_{it}$. Firms face downward sloping demand schedules from the consumer problem, $Y_{it} = C_{it} = \left(\frac{P_{it}}{P_t}\right)^{-\nu} C_t$.

Finally, a monetary authority targets a constant rate of inflation, which manifests as the nominal wage W_t growing at a constant rate:

$$\log(W_t) = \pi + \log(W_{t-1}).$$

Flexible prices. If firms; prices are flexible and costless to adjust, they will change every period to reflect movements in productivity and wages. The per-period profit maximization problem is

$$\max_{P_{it}, L_{it}} P_{it} Y_{it} - W_t L_{it}$$
$$\text{s.t. } Y_{it} = A_{it} L_{it}, \quad Y_{it} = \left(\frac{P_{it}}{P_t}\right)^{-\nu} C_t$$

Substituting in constraints and differentiating with respect to P_{it} yields FOC

$$(1-\nu) P_{it}^{-\nu} P_t^{\nu} C_t + \nu \frac{W_t}{A_{it}} P_{it}^{-\nu-1} P_t^{\nu} C_t = 0$$

and ideal price

$$P_{it}^* = \frac{\nu}{\nu - 1}\frac{W_t}{A_{it}}.$$

The flexible price equilibrium is given by the household and firms maximizing utility and profits, respectively, the path of wages set by the monetary authority, and labor market clearing.

Defining aggregate productivity $A_t = \left(\int A_{it}^{\nu-1} di\right)^{\frac{1}{\nu-1}}$, in the flexible price equilibrium I have

$$P_t = \left(\int \left(\frac{\nu}{\nu-1}\frac{W_t}{A_{it}}\right)^{1-\nu} di\right)^{\frac{1}{1-\nu}} = \frac{\nu}{\nu-1}\frac{W_t}{A_t}$$

and

$$L_t = \int \frac{C_{it}}{A_{it}} di = \int \left(\frac{P^*_{it}}{P_t}\right)^{-\nu} \frac{C_t}{A_{it}} di$$
$$= A_t^{-\nu} C_t \int A_{it}^{\nu-1} = \frac{C_t}{A_t} \implies C_t = A_t L_t$$

These, coupled with the optimal intratemporal household condition, determine steady state labor supply and consumption in the flexible price model.

$$\bar{C}_{flex} = \frac{\nu - 1}{\alpha\nu} \bar{A}, \quad \bar{L}_{flex} = \frac{\nu - 1}{\alpha\nu}$$

where $\bar{A}$ is steady state productivity, which is a primitive of the model pinned down by the stationary distribution of A_{it}.

Sticky prices If changing prices requires paying random menu cost κ in labor, firms' flow profits are

$$\Pi_{it} = P_{it} Y_{it} - W_t L_{it} - W_t \kappa_{it} I_{it} = \left(P_{it}^{1-\nu} - \frac{W_t}{A_{it}} P_{it}^{-\nu}\right) P_t^{\nu} C_t - W_t \kappa_{it} I_{it},$$

where I_{it} is an indicator variable taking value 1 when a firm changes its price.

Firms then set prices to maximize real profits Π_{it}/P_t discounted at rate β. Defining price deviations $X_{it} = \frac{P_{it-1}}{P^*_{it}}$, the firm problem becomes

$$\max_{\{P_{it+k}\}} B \mathbb{E}_t \left[\sum_{k=t}^{\infty} \beta^k \left(A_{it+k}^{\nu-1} \left(X_{it+k}^{1-\nu} - \frac{\nu - 1}{\nu} X_{it+k}^{-\nu}\right) - \hat{\kappa}_{it+k} I_{it+k}\right)\right],$$

where $B = (\frac{\nu}{\nu-1}\alpha\bar{C})^{1-\nu}\bar{C}$ and $\hat{\kappa} = (\frac{\nu}{\nu-1}\alpha\bar{C})^{\nu-1}\alpha\kappa$, and $\bar{C}$ is steady state consumption. Note that here I define X as the ex-post price gap following firms' decisions to change prices, whereas in the main text I define X as the ex-ante price gap prior to firms' decisions. Applying the firm policy function can easily convert from the latter to the former.

The sticky-price equilibrium is given by the household and firms maximizing utility and profits, respectively, the path of wages set by the monetary authority, and labor market clearing, accounting for labor spent on paying menu costs.

In equilibrium, the outcome of the firm maximization problem generates a stationary distribution over firms of $F = (X_i, A_i)$. I will use the notation $dF(i)$ to denote integrating over firms with respect to this distribution.

To contrast with the flexible price economy, I add hats to variables to denote that they are outcomes of the sticky-price equilibrium. Note, however, that the path of

wages W_t is the same between the two equilibria, as are ideal prices P_{it}^*.

In the sticky-price equilibrium, we have

$$\hat{P}_t = \left[\int (X_{it}P_{it}^*)^{1-\nu} dF(i)\right]^{\frac{1}{1-\nu}} = \frac{\nu}{\nu-1} W_t \left[\int X_{it}^{1-\nu} A_{it}^{\nu-1} dF(i)\right]^{\frac{1}{1-\nu}}$$

and

$$\hat{L}_t = \int \frac{\hat{C}_{it}}{A_{it}} dF(i) + \hat{L}_t^{pc} = \int \left(\frac{X_{it}P_{it}^*}{\hat{P}_t}\right)^{-\nu} \frac{\hat{C}_t}{A_{it}} dF(i) + \hat{L}_t^{pc}$$

$$= \frac{\int X_{it}^{-\nu} A_{it}^{\nu-1} dF(i)}{\left[\int X_{it}^{1-\nu} A_{it}^{\nu-1} dF(i)\right]^{\frac{\nu}{\nu-1}}} \hat{C}_t + \hat{L}_t^{pc} \implies \hat{C}_t = \frac{\left[\int X_{it}^{1-\nu} A_{it}^{\nu-1} dF(i)\right]^{\frac{\nu}{\nu-1}}}{\int X_{it}^{-\nu} A_{it}^{\nu-1} dF(i)} (\hat{L}_t - \hat{L}_t^{pc})$$

where $\hat{L}_t^{pc}$ is labor spent on price changes. Coupling these with the intratemporal household optimality condition, $\alpha \hat{C}_t = \frac{W_t}{\hat{P}_t}$, pins down steady state labor supply and consumption in the sticky-price model.

$$\bar{C} = \frac{\nu-1}{\alpha\nu} \left[\int X_i^{1-\nu} A_i^{\nu-1} dF(i)\right]^{\frac{1}{\nu-1}}, \quad \bar{L} = \frac{\nu-1}{\alpha\nu} \frac{\int X_i^{-\nu} A_i^{\nu-1} dF(i)}{\int X_i^{1-\nu} A_i^{\nu-1} dF(i)} + \bar{L}^{pc}$$

A.3 Welfare loss measure

The measure of welfare loss is the consumption equivalent units needed to equate welfare in the sticky-price model with welfare in the analogous flexible price model. This is given by Λ such that

$$\log((1+\Lambda)\bar{C}) - \alpha\bar{L} = \log(\bar{C}_{flex}) - \alpha\bar{L}_{flex}$$

This gives the following expression for Λ in terms of model primitives:

$$\Lambda = \frac{\bar{C}_{flex}}{\bar{C}} \exp(\alpha\bar{L} - \alpha\bar{L}_{flex}) - 1$$

$$= \frac{\bar{A}}{\left[\int X_i^{1-\nu} A_i^{\nu-1} dF(i)\right]^{\frac{1}{\nu-1}}} \exp\left(\frac{\nu-1}{\nu} \frac{\int X_i^{-\nu} A_i^{\nu-1} dF(i)}{\int X_i^{1-\nu} A_i^{\nu-1} dF(i)} + \alpha L^{pc} - \frac{\nu-1}{\nu}\right) - 1$$

B Monte Carlo Simulations

B.1 The Identifying Power in the Stylized Model

In Section 2.2, I prove that parameters exist such that both the permanent shocks model and the transitory shocks model can exactly match an observed frequency and variance of price changes. I then assert that, given such parameters, it is better to use vintage moments to differentiate the two models than to use additional cross-sectional moments. I provide further evidence for this claim here via Monte Carlo simulations.

I first choose parameters σ and K in each model to match frequency of price change .088 and variance of price changes .025. Then for each model, I simulate panels of prices with 1,600 panel units over 60 periods. These data dimensions mimic the number of price quote observations within a typical product category in the UK CPI. I conduct 10,000 such simulations, and for each one I calculate the higher order even moments of the cross-sectional distribution of price changes, vintage conditional variances, and vintage conditional hazard rates. Figure B.1 plots averages of the simulated moments, along with bands marking the 5th and 95th percentiles.

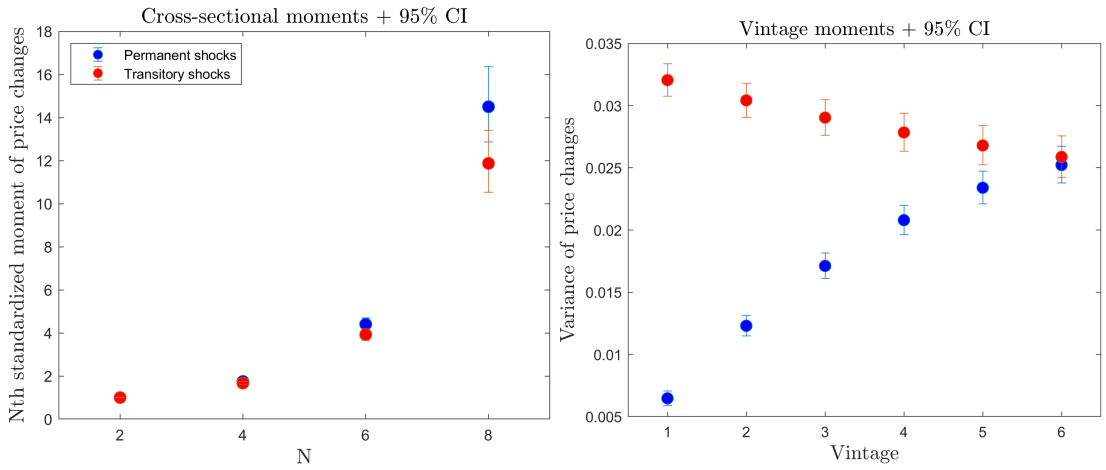


Figure B.1: Monte Carlo results: selected moments

The left panel shows the mean of the sampling distributions for higher order moments of the cross-section of price changes. These moments have been normalized using the observed variance of price changes for scaling. The right panel shows the mean of the sampling distributions of vintage conditional variances at various vintages. The bars delineate the 5th and 95th percentiles of the sampling distributions.

The separation between the percentiles of the sampling distributions clearly indicates that vintage moments have greater identifying power. In Figure B.2 I plot these sampling distributions for the fourth standardized moment (analogous to kurtosis) of the cross-section of price changes, as well as the variances of price changes conditional on a vintage of 1. The overlap between the distributions of kurtoses is such that, if the transitory shocks model were the true one, the hypothesis that the true

model were the permanent shocks model would only be rejected 48% of the time, with a 95% confidence level.

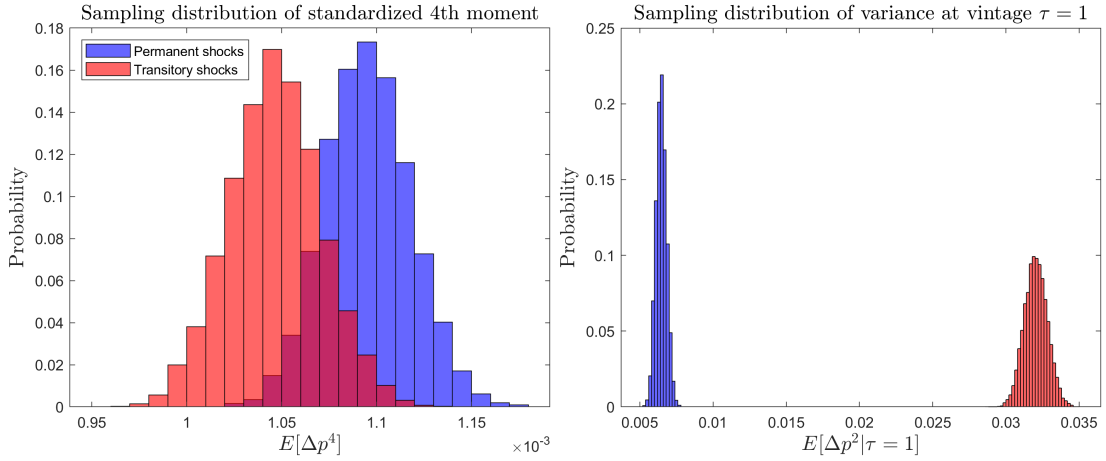


Figure B.2: Simulated distributions of selected moments

The left panel shows the sampling distributions of the fourth standardized moment of price changes in the simulated permanent shocks and transitory shocks case. The right panel shows the same, but for the variance of price changes conditional on a vintage of $\tau = 1$.

B.2 Composition Bias from Unobserved Heterogeneity

As stated in Section 3, the primary concern is composition bias across vintage with respect to unobserved types within product categories. In particular, the worry is that the observed empirical facts of hazard rates and variances of price changes declining with respect to vintage reflect this composition bias rather than firm-level shocks being temporary.

To investigate how much weight to place on this concern, I benchmark the degree of bias that would arise given observed heterogeneity *across* product categories, assuming that this is larger than whatever unobserved heterogeneity is present within product categories.

For each product category, I calculate the average frequency of price changes and variance of price changes. I then calibrate the permanent shocks economy of Section 2 twice, once matching 25th percentile of frequency and variance of price changes across product categories, and once matching the 75th percentiles of those statistics.[22] I then simulate a panel of prices with 800 panel units over 60 periods for each calibration, and calculate vintage moments for the two panels pooled together. I do this for 5,000 such simulations. Results for the vintage moments are plotted in Figure B.3, along with those for an analogous economy with homogeneous firms, calibrated

[22]Composition bias naturally arises for vintage-conditional hazard rates, but will only show up in vintage-conditional variances if unobserved types with high hazard rates also have more dispersed price changes. By matching the same percentile of each, I am being conservative by assuming maximal correlation. Empirically, there is a slight positive correlation between the average frequency of price change and the variance of price changes.

to match the average frequency and variance of price changes across the two cases. I also conduct the same exercise for the 10th and 90th percentiles.

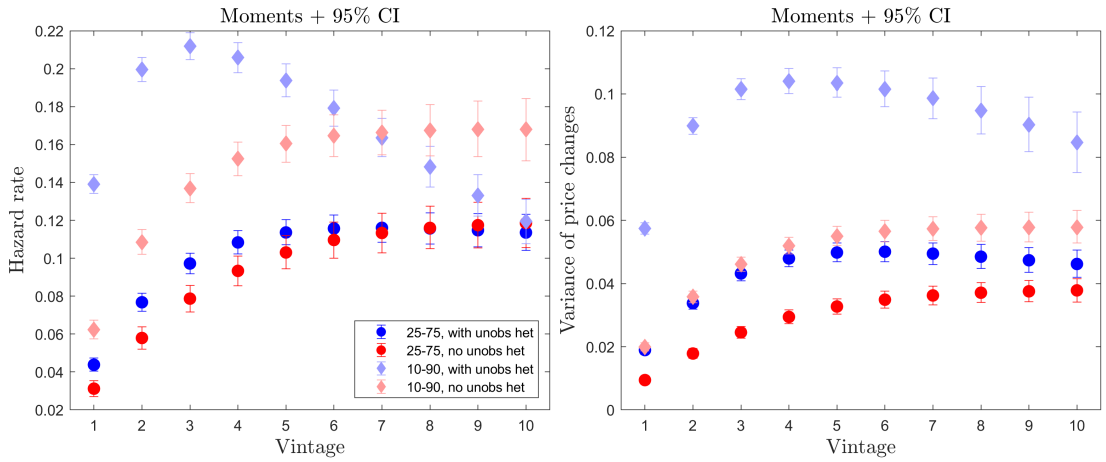


Figure B.3: Effect of unobs het on simulated vintage moments

The left panel shows profiles of the vintage conditional hazard rates for the simulated models described above. The right panel shows the same for vintage conditional variances of price changes. Blue markers are the simulations with unobserved heterogeneity, red markers are the simulations without unobserved heterogeneity.

Unobserved heterogeneity does clearly induce some composition bias, as seen by the differences between red and blue markers with the same shapes. Furthermore, as expected, larger heterogeneity increases the degree of the bias. This, however, remains a far cry from the pattern observed in the data presented in Section 3, in which not only are the profiles of vintage moments decreasing, but the majority of the decrease occurs at early vintages. Provided that within-category heterogeneity is smaller than across-category heterogeneity, I take this as evidence that, first, unobserved within-category heterogeneity alone is not sufficient to rationalize persistent firm-level shocks with the data, and second, vintage-conditional moments still have identifying power for the persistence of shocks.

C Data appendix

C.1 UK data: Additional details on sample construction

I begin with the sample of all locally collected price quotes from 1996 to 2023. This sample does not cover the entirety of prices used to construct the CPI, as some prices are "centrally collected." These are typically for products that are either priced nationally, primarily purchased online, or purchased through mail catalog. This includes products sold by large retailers with national presence. The price quotes are not linked over time, but I can reconstruct a panel identifier using provided information on the sample location and product type. The sample location can be identified

using variables 'region', 'shop type', and 'shop id'. This is combined with the product category variable, 'item id', to construct identifiers for panel units. For some shops, multiple items within the same product category are sampled. In these cases, I drop the observations which cannot be uniquely attributed to a panel unit.

Given this initial panel, I then create breaks in the data when there is a product substitution, weight change, or quality change, as recorded by the ONS during data collection. When such a substitution occurs, the ONS resamples the same, substituted product in subsequent periods. I therefore assign a new panel id to all periods following the break. This increases the number of panel units but does not change the overall number of observations.

Finally, I drop observations that the ONS classifies as invalid, as well as those that are missing, top coded, or have a price of less than 1 pence. The ONS validation procedure is intended to eliminate errors in the hand collection of prices. In broad strokes, it consists of double-checking price observations that are unexpectedly large or small, ensuring that sampled item is exactly the same month to month, and flagging internal inconsistencies or missing fields in collected data. More info on the data validation procedure can be found in the ONS Technical Manual (Office for National Statistics, 2024). The observations that remain are the final sample of price quotes. With these, I follow ONS procedure to reconstruct the CPI index. Figure C.1 compares the official and in-sample CPI. Table C.1 reports the number of observations remaining through each step.

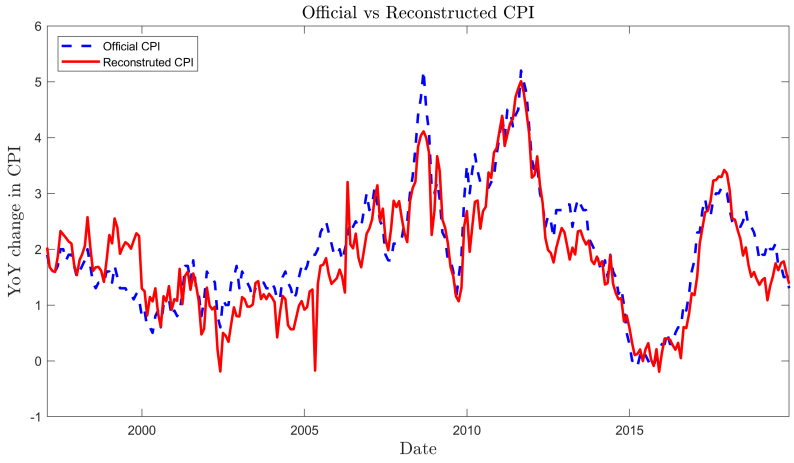


Figure C.1: UK CPI Inflation, Official vs Constructed

The blue line is the official year-over-year percent change in CPI, as published by the UK ONS. The red line is the year-over-year percent change of an analogous index constructed from the subsample of price observations.

Initial number of price observations...	39,526,450
...removing without unique identifier:	33,221,389
...removing invalid and missing prices:	31,431,675
...removing missing price changes:	23,099,634
...removing missing vintages:	9,992,628

Table C.1: Number of observations

There are further refinements for the sample with price changes and sample with price changes and vintage. In the sample with price changes, I drop observations where the price change is missing, most of which are the first observations of panel units, and I drop price changes that are greater than a factor of 10 on the grounds that these are implausibly large. I also drop observations from after 2019. In the sample with vintage, I simply drop observations for which vintage is unknown, which correspond to the first price spell for each panel unit. Dimensions of the main samples are reported in Table C.2.

	Sample with price changes	Sample with vintage
Number of price quotes	23,099,634	9,992,628
Number of price changes	2,085,261	1,215,488
Number of panel units	2,060,402	624,003
Number of product categories	1,285	1,285

Table C.2: Dimensions of UK data

C.2 Summary statistics: price change and vintage

	All price changes	With vintage
Frequency of price changes	.0903	.122
Average price change	.00844	.00955
Average abs price change	.129	.117
Variance of price changes	.0391	.0328
Kurtosis of price changes	8.66	9.21
Observations	2,085,261	1,215,488

Table C.3: UK Data Summary Statistics

D Supplemental figures and tables

D.1 Theory and simulation: supplemental figures

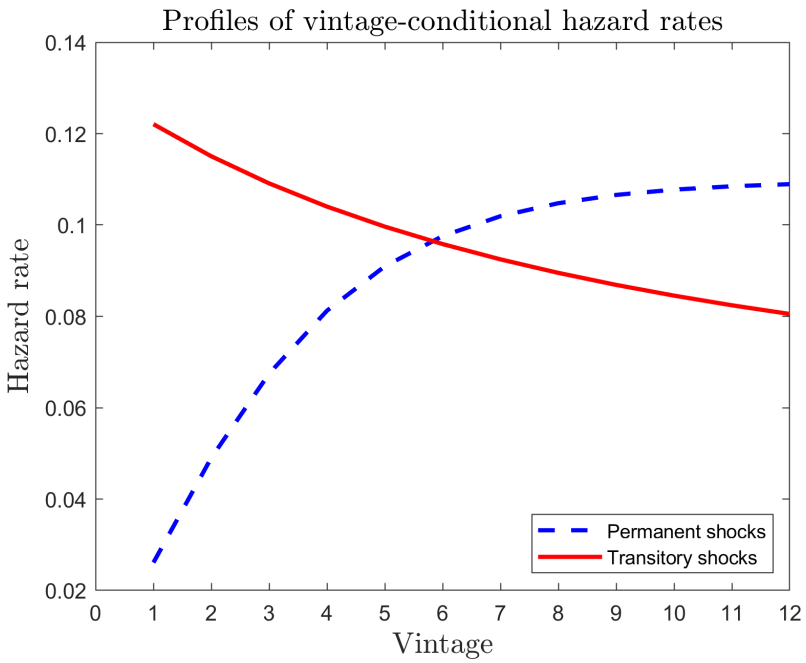


Figure D.1: Profiles of vintage-conditional hazard rates

This figure plots vintage-conditional hazard rates in the permanent and transitory shocks economies described in Section 2.2.

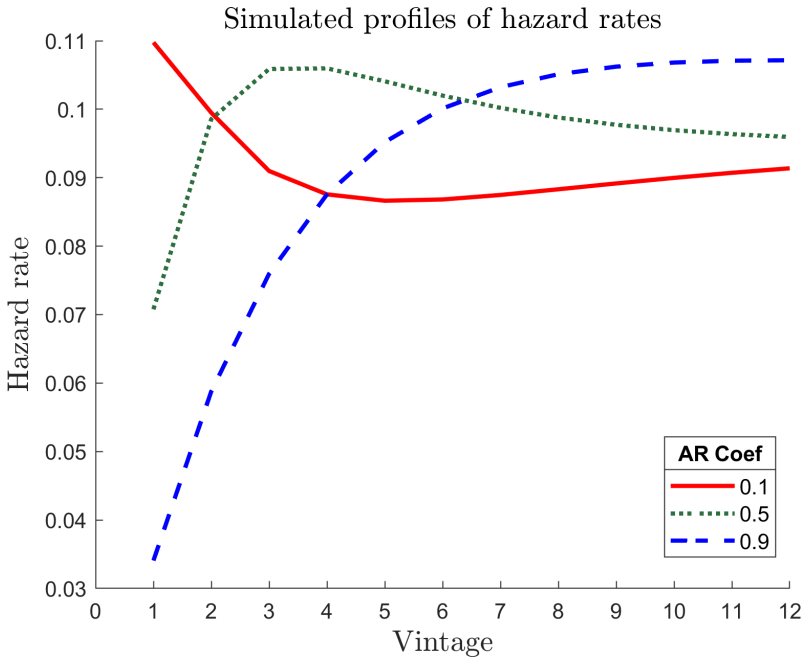


Figure D.2: Hazard rates

This figure plots simulated vintage-conditional hazard rates for selected AR coefficients in the simulated economy described in Section 2.2.

D.2 UK CPI: supplemental figures and tables

Results of regression

$$\lambda_{j,\tau} = \alpha_j + \beta\tau + u_{j,\tau}$$

where vintage τ is calculated at the quarterly or annual frequency.

	(1) $\lambda_{j,\tau}$	(2) $\lambda_{j,\tau}$
β_q	-.00584	
	(.000307)	
β_a		-.0134
		(.000944)
Fixed Effects	Product Category	Product Category
Aggregation	Quarterly	Annual
Observations	6808	2266
R-Squared	.854	.855

Standard errors in parentheses

Table D.1: Vintage vs hazard rate, alternate time aggregation

Results of regression

$$\lambda_{j,\tau} = \alpha_j + \beta\tau + \delta_4\tau_4 + \delta_1 2\tau_1 2 + u_{j,\tau}$$

where τ is monthly and the sample is only market goods.

	(1) $\lambda_{j,\tau}$	(2) $\lambda_{j,\tau}$	(3) $\lambda_{j,\tau}$
δ_4	-.0185		-.0114
	(.00153)		(.00143)
δ_{12}	.0172		.0501
	(.00386)		(.00366)
β		-.00482	-.00505
		(.000112)	(.000116)
Fixed Effects	Product Category	Product Category	Product Category
Observations	12836	12836	12836
R-Squared	.746	.777	.781

Standard errors in parentheses

Table D.2: Vintage vs hazard rate, market goods

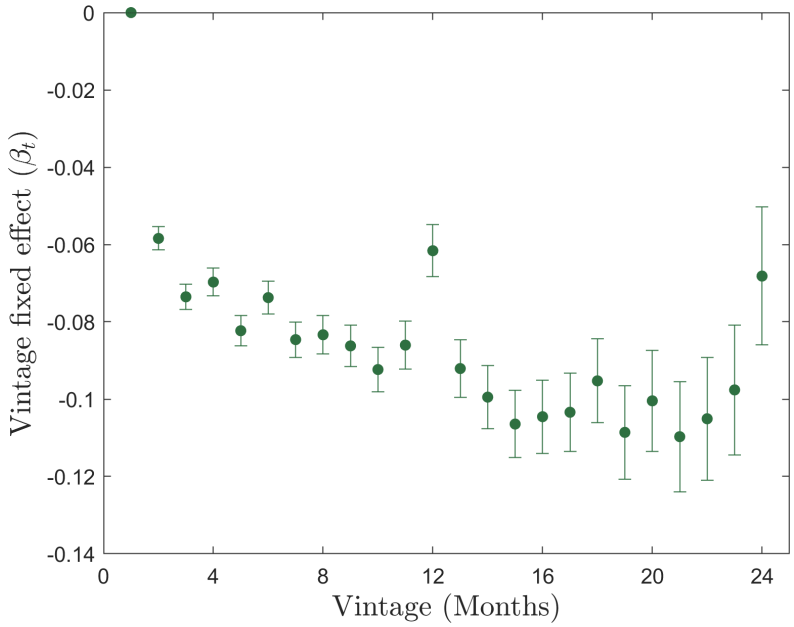


Figure D.3: Coefficients from fixed effects regression, market goods

This figure plots the estimated vintage fixed effects from the regression $\lambda_{j,\tau} = \alpha_j + \sum_{t=1}^{24} \beta_t \mathbb{1}\{\tau = t\} + u_{j,\tau}$ on the sample of market goods. Estimates are plotted relative to $\beta_1 = 0$. Bars represent a 95% confidence interval.

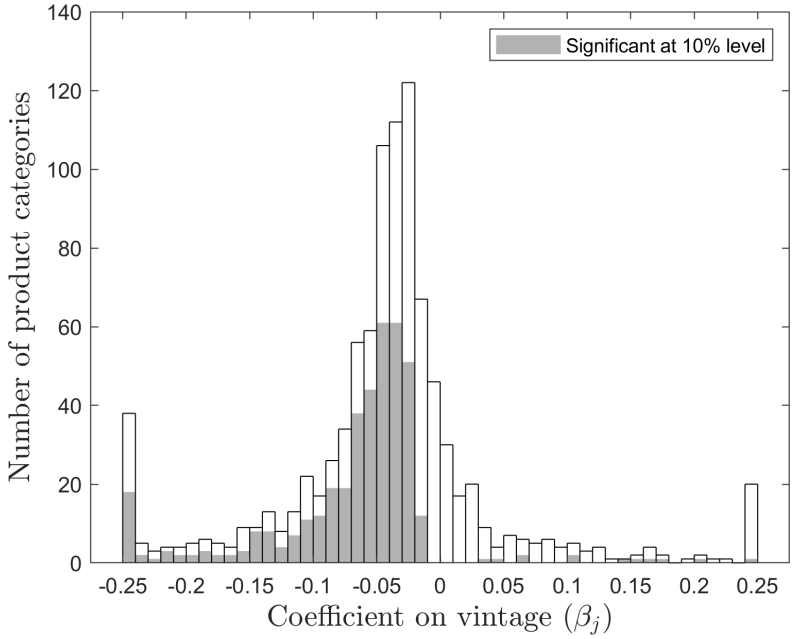


Figure D.4: Coefficients of variance vs vintage, market goods

This is a histogram of the estimated β_j coefficients from the regression $Var(\Delta\tilde{p}_{it}|j,\tau) = \alpha_j + \beta_j\tau + u_{j,\tau}$, on the subsample of market goods. Coefficients that are significant at the 10% level are overlaid in grey.

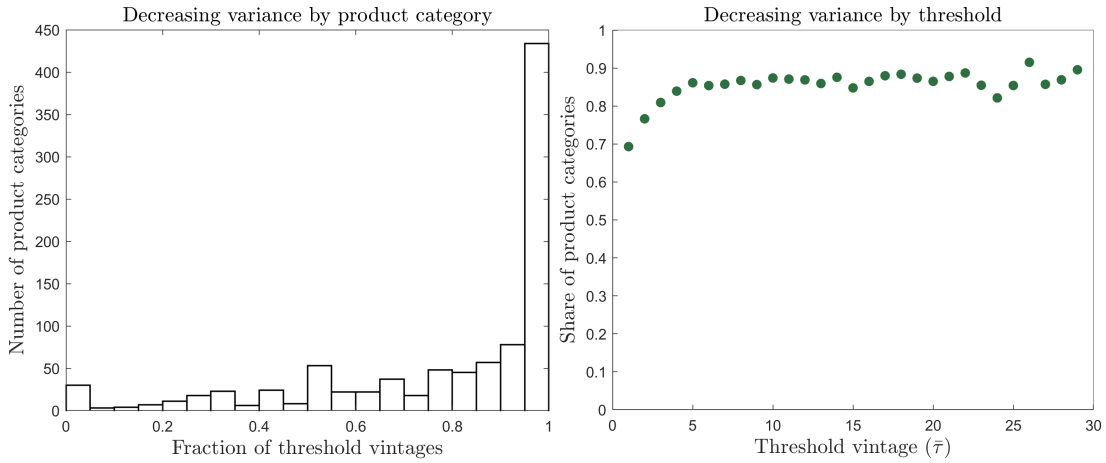


Figure D.5: Prevalence of locally decreasing variance of price changes, market goods

The left panel is a histogram of the number of product categories by the share of cutoff vintages $\bar{\tau}$ for which the variance of price changes is locally decreasing in vintage. The right panel shows the fraction of product categories for which the variance of price changes is higher among vintages younger than $\bar{\tau}$ than among vintages older than $\bar{\tau}$, where $\bar{\tau}$ is vintage on the x-axis. Both are constructed using the subsample of market goods.

D.3 UK CPI: controlling for unobserved heterogeneity

	Original sample (with vintage)	Adjusted for heterogeneity
Average price change	.00955	.0119
Average abs price change	.117	.106
Variance of price changes	.0328	.0285
Kurtosis of price changes	9.21	10.2
Observations	1,215,488	555,416

Table D.3: UK Data Summary Statistics - no unobs heterogeneity

Results of regression

$$\lambda_{j,\tau} = \alpha_j + \beta\tau + \delta_4\tau_4 + \delta_{12}\tau_{12} + u_{j,\tau}$$

on the subsample of panel units that are not rejected for unobserved heterogeneity.

	(1) $\lambda_{j,\tau}$	(2) $\lambda_{j,\tau}$	(3) $\lambda_{j,\tau}$
δ_4	.0172		.016
	(.00201)		(.00202)
δ_{12}	.108		.103
	(.00499)		(.0051)
β		.00224	.000931
		(.000181)	(.000184)
Fixed Effects	Product Category	Product Category	Product Category
Observations	10607	10607	10607
R-Squared	.702	.684	.702

Standard errors in parentheses

Table D.4: Vintage and hazard rate - no unobs heterogeneity

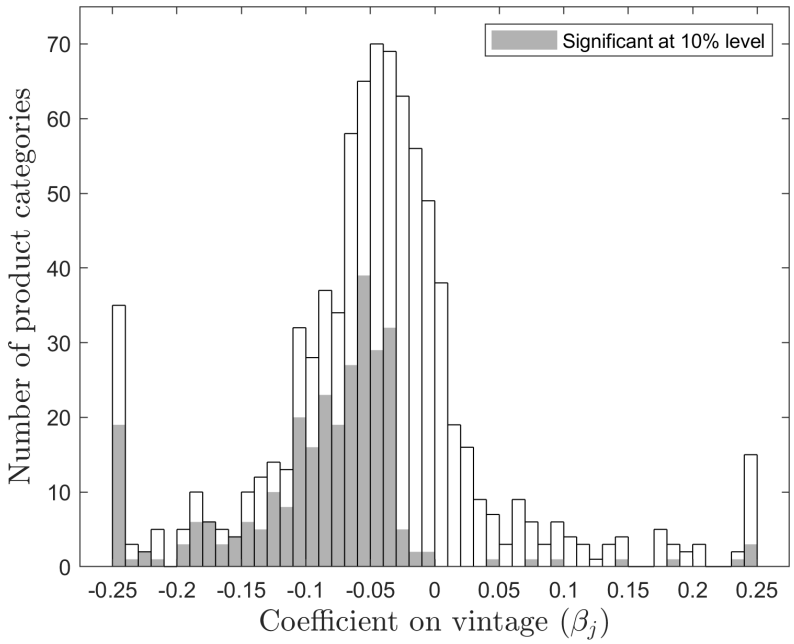


Figure D.6: Coefficients of variance vs vintage - no unobs het

This is a histogram of the estimated β_j coefficients from the regression $Var(\Delta \tilde{p}_{it}|j,\tau) = \alpha_j + \beta_j \tau + u_{j,\tau}$, on the subsample of panel units that are not rejected for unobserved heterogeneity, as described in Section 3.4. Coefficients that are significant at the 10% level are overlaid in grey.

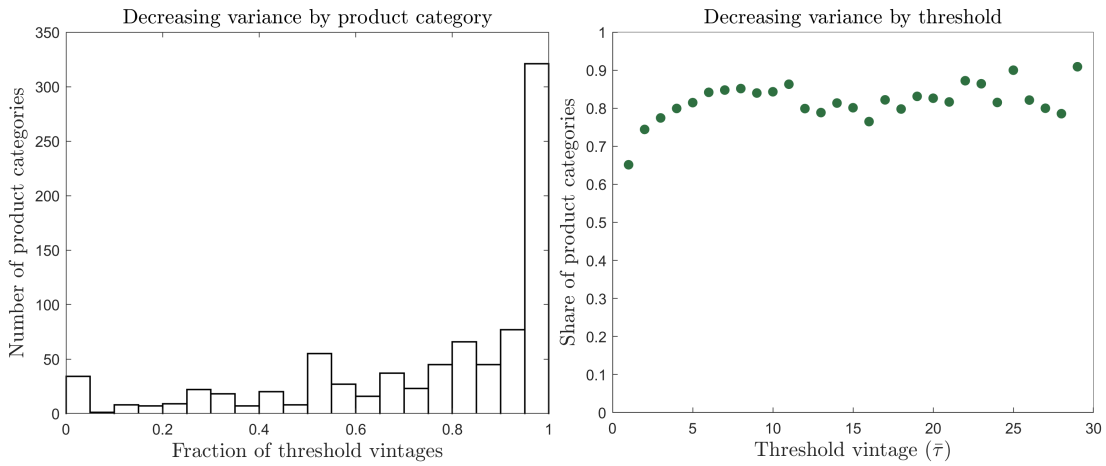


Figure D.7: Prevalence of locally decreasing variance of price changes - no unobs het

The left panel is a histogram of the number of product categories by the share of cutoff vintages $\bar{\tau}$ for which the variance of price changes is locally decreasing in vintage. The right panel shows the fraction of product categories for which the variance of price changes is higher among vintages younger than $\bar{\tau}$ than among vintages older than $\bar{\tau}$, where $\bar{\tau}$ is vintage on the x-axis. Both are constructed using the limited sample of price quotes that are not rejected for unobserved heterogeneity, as described in Section 3.4.

D.4 Additional figures from Calibrated Model

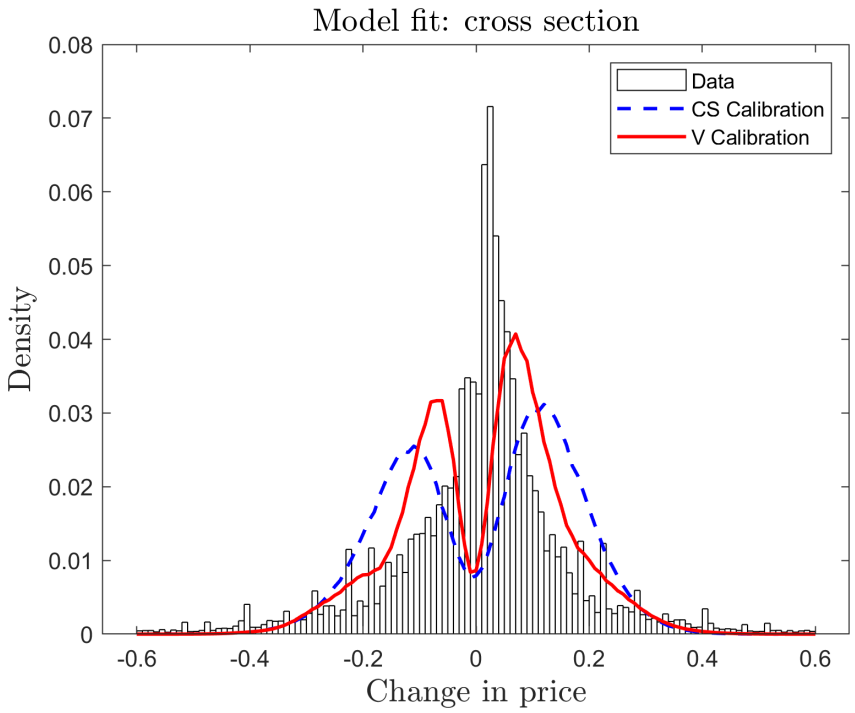


Figure D.8: Calibrated model fit - density of price changes

This figure compares the model-implied invariant distribution of price changes to the observed distribution of price changes in the UK from 1996 to 2019. The model is that of Section 4, calibrated to match cross-section (CS) moments or vintage (V) moments as described in Section 4.2.

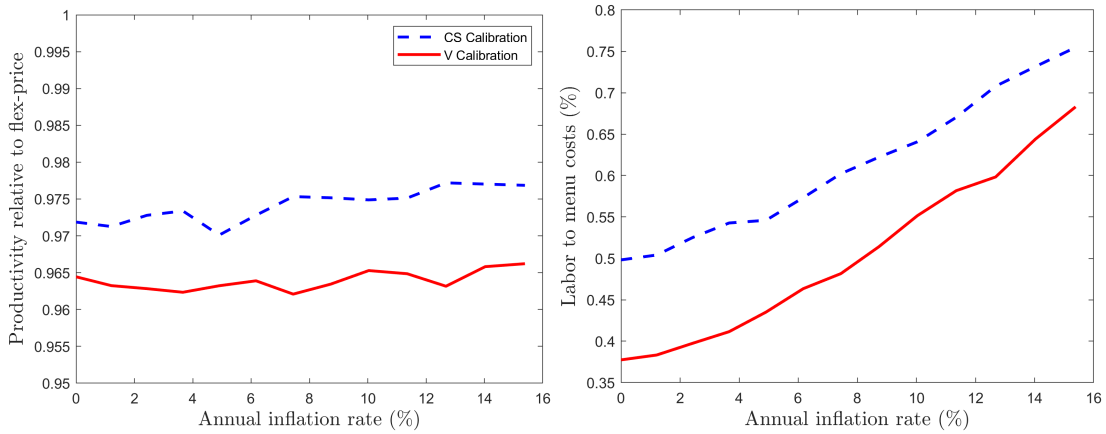


Figure D.9: Decomposition - sources of welfare loss

The left panel plots the model-implied allocative efficiency as a percentage of allocative efficiency in the frictionless equilibrium, at varying levels of steady-state inflation. The right panel plots the percentage of total labor used to pay menu costs at varying levels of steady-state inflation. Both panels show results for the model of Section 4, calibrated to match cross-section (CS) moments or vintage (V) moments as described in Section 4.2.