Revisiting the Cost of Inflation: Identification Using Vintage Moments

Michael Jenuwine* and Joris Tielens†

November 2025

Updated regularly, please click HERE for the latest version.

Abstract

Measuring the welfare cost of inflation requires determining to what degree observed movements in relative prices reflect firms' idiosyncratic fundamentals. This paper introduces a new empirical approach to measuring the cost of inflation in sticky-price models: conditioning observed price changes on vintage, or the duration of price spell. Theoretically, we demonstrate that vintage-conditional data is more informative than typically used cross-sectional data. In particular, the vintage-conditional frequency and variance of price changes are informative about the persistence of idiosyncratic shocks to firms' fundamentals. We document vintage moments using the micro-price data underlying the UK CPI and Belgian PPI. Within narrow product categories, price changes become less frequent and less dispersed as more time passes since the previous price change. Estimating a quantitative sticky-price model to match these vintage moments indicates that idiosyncratic shocks are smaller and more transitory compared to a baseline estimation matching cross-sectional moments. In the model that matches vintage moments, the welfare cost of inflation is 1.3 percentage points higher compared to the baseline, and this welfare cost increases 4 times more quickly as inflation rises. We validate our estimation procedure using linked data on the output, costs and prices of Belgian manufacturing firms.

*Princeton University. E-mail: jenuwine@princeton.edu; † National Bank of Belgium. E-mail: joris.tielens@nbb.be. This research benefited from financial support from the International Economics Section at Princeton University. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the National Bank of Belgium, the Eurosystem, or any other institution with which the authors are affiliated.

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 firms' ideal prices – 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 their *duration of price spell*, which we refer to as vintage.[1] We demonstrate with an illustrative example that vintage-conditional moments are more informative for estimating sticky-price models than the typically used moments derived from a single cross-section of price changes. In particular, vintage moments better identify how firms' ideal prices evolve. Next, we document systematic empirical patterns in vintage moments that are not reproduced in previously estimated sticky-price models. Finally, we estimate a standard New Keynesian general equilibrium model to match vintage moments and calculate the associated welfare cost of inflation. We compare this welfare cost to that of a baseline model, which is estimated to match commonly used cross-sectional moments.

We begin in 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 estimation problem: 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-conditional moments assist model estimation.

We demonstrate the usefulness of vintage-conditional moments using a simple example, where menu costs are exponentially distributed and shocks follow an AR(1) process with persistence ρ. Comparing the opposing cases of $\rho = 1$ and $\rho = 0$, we show that both cases can exactly match cross-sectional data on the overall frequency and variance of price changes. Furthermore, the cross-sectional distributions of price

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

changes implied by each case are nearly identical. On the other hand, comparing moments across vintages provides a simple heuristic for distinguishing between the two cases. When $\rho = 1$, the vintage-conditional probability of price adjustment and vintage-conditional variance of price changes are monotonically increasing in vintage, while both are monotonically decreasing in vintage when $\rho = 0$. We demonstrate numerically that similar patterns also hold for $\rho \in (0, 1)$.

Cross-sectional data alone performs poorly because multiple underlying firm states may map to the same price change. Using vintage data is beneficial because it adds information about how long it takes firms to select into changing prices. Differences between the distribution of price changes at low vintages, when firms have been hit by few idiosyncratic shocks, and high vintages, when firms have been hit by many such shocks, are intuitively informative about the stochastic process by which shocks evolve.

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

Then, to demonstrate the impact of using vintage conditional moments in estimation, we move to a New Keynesian general equilibrium setting that more closely mirrors the types of quantitative models used for policymaking. Firms compete monopolistically à la CES 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.

We estimate this model twice: once to match cross-sectional moments that are typically targeted in such a procedure, and once to match vintage-conditional moments. When matching the vintage moments, we recover much less persistent idiosyncratic productivity shocks (a monthly autocorrelation of .24 as opposed to .94) with smaller innovations, as well as lower average menu costs compared to when we match cross-sectional moments.

For each set of estimated parameters, we then solve the quantitative model and compute the implied welfare loss at various rates of steady-state inflation. At the common inflation target of 2%, the model-implied welfare cost of inflation is 1.3 percentage points higher when using the vintage parameters compared to when using

the cross-section parameters. This welfare cost increases at a 4 times faster rate as inflation rises when using the vintage parameters compared to the cross-sectional parameters. The level of welfare loss is larger because firms react less to idiosyncratic productivity shocks, understanding that those movements are temporary.

Finally, we exploit a dataset linking Belgian manufacturing firms' costs, production, and prices to validate our estimation procedure. For each Belgian manufacturing sector, we estimate a process for ideal prices using cost and production data. We compare the parameters of these processes to the ones we estimate in the model using vintage moments, and validate that vintage moments recover reasonable values for the autocorrelation of firm productivity shocks.

Related Literature. Our paper primarily contributes to a literature concerning how to estimate sticky-price models using micro-price data. Models with menu costs are primarily disciplined by matching 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 observed price changes, provided that idiosyncratic shocks follow a random walk.

Our contribution is to introduce vintage-conditional moments as an important feature of micro-price data to match in model estimation. Matching these moments matters especially in the case of stationary idiosyncratic shocks.

Second, we 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).

We 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, we connect to a literature documenting dynamic features of micro-price data.[2] In both old surveys (Klenow and Malin, 2010) and recent papers de-

[2]We call vintage moments "dynamic" in these 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

scribing 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; Campbell and Eden, 2014). In the same UK CPI micro-price data we use, Bunn and Ellis (2012b,a) document facts concerning vintage-conditional hazard rates and sizes of price changes.

Our work adds to the evidence that hazard rates decline in vintage, and we document a new fact that the vintage-conditional variance of price changes declines in vintage. We also document these features in a new micro-price dataset that underlies the Belgian producer price index.

Finally, we contribute to literature that studies sticky-price models with linked firm-level cost and price data. These settings are few and far between, but are powerful as observing costs allows one to construct a proxy for firms' ideal prices.[3] Eichenbaum et al. (2011) use linked profit, price and quantity data for supermarkets to examine synchronicity between changes in prices and changes in costs. More recently, Gagliardone et al. (2023) and Gagliardone et al. (2025) use linked cost and price data for Belgian manufacturers to measure state-dependence in firm pricing policies and the slope of the Phillips curve.

Recognizing that panels of firms' prices are often observed but linked costs are rare, our contribution is to use cost data to inform what features in a panel of firm prices are important. Specifically, we validate that vintage-conditional moments are particularly informative about the how firms' costs stochastically evolve.

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-conditional moments. Section 3 documents the behavior of vintage-conditional moments in the data underlying the UK CPI. Section 4 presents the benefit of using vintage moments for estimation and its implications for the welfare cost of inflation, while Section 5 validates our estimation procedure using linked cost and price data from Belgium. Section 6 concludes.

2 Theoretical Framework

Our theoretical framework is a standard discrete-time, random menu cost model with idiosyncratic shocks in the tradition of Caballero and Engel (1993). This framework

evolves over the business cycle, as in Gagnon (2009), Vavra (2013) and Berger and Vavra (2018).

[3] Another way to proxy for ideal prices is to study either time periods in which policy actions alter costs, or industries in which a portion of marginal cost can be observed (Nakamura and Zerom, 2009; Gautier et al., 2023). While more common, in these approaches the observed shocks to ideal prices are the same across a set of firms, rather than idiosyncratic and uncorrelated across firms.

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 price setting frictions in the form of random menu costs.

Following Alvarez et al. (2016), we 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*.

We model ideal prices p_{it}^* as composed of a firm fixed effect, an aggregate component which follows a deterministic trend, and an idiosyncratic component which evolves either 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 to 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}. We 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. We do not permit aggregate shocks to ideal prices for expositional simplicity, but they can be accommodated.[4]

We 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 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.

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

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.

We define a firm's *price gap*, as $x_{it} \equiv p_{it-1} - p^*_{it}$. With this, we can write the firm's value function in recursive form

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

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

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

We cannot write a closed form expression for $\bar{x}$ at this level of generality. However, we can see that $\bar{x}$ is a function only of the current idiosyncratic shock, a. We 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:

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

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

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

This lets us write the adjustment hazard function, or the probability of a firm with state (x, a) changing its price, as

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

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

[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.

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 we 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 firms 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 when only panel data on the prices set by firms are observed. 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. We illustrate that there is a superior alternative to these cross-sectional data.

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]

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

To best highlight our key message, we focus on two extreme cases: a permanent shocks economy in which $\rho = 1$, and a transitory shocks economy in which $\rho = 0$. Given a panel of prices, we 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

[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 we derive. The assumption of myopia will later be relaxed.

price changes. We show that for each economy, we 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 extremely different processes for ideal prices, both can match these commonly used moments. If we restrict ourselves 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 we can analytically demonstrate 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 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 non-obvious how to harness the gaps between them.

A better method is to move beyond the cross-section of price changes. We 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. We 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, we 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.*
- *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

[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.

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. We 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 C.1). The difference between the two economies is stark, and clear to the eye. Vintage-conditional moments are highly informative for differentiating between these two economies.

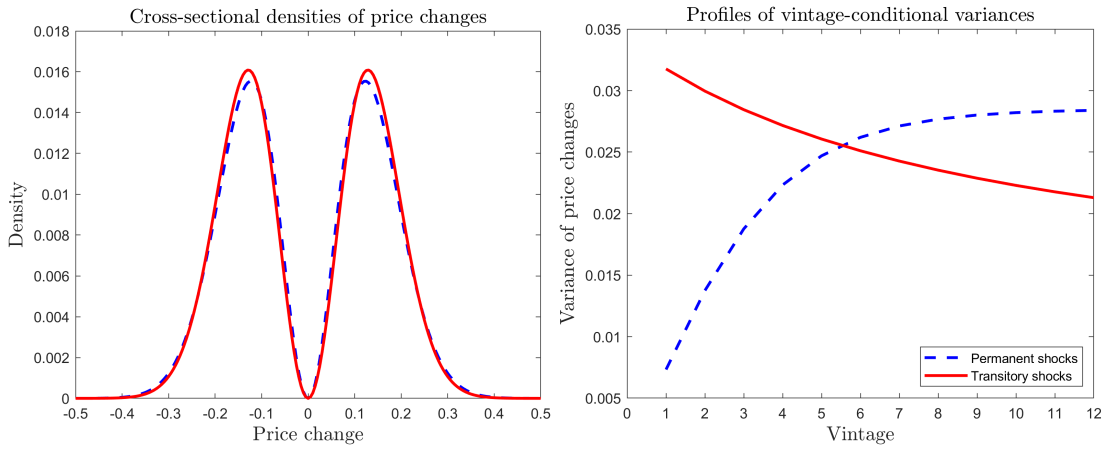

Figure 1: Cross-sections vs vintage-conditional variances

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.

We now show that vintage-conditional moments are also informative in a less styl-

ized setting. In particular, we relax the assumption that firms are myopic and allow for non-extreme values of the AR parameter, ρ. We 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 $\rho \in (0, 1)$, we select K and σ to match frequency of price changes $\lambda = .088$ and variance of price changes $Var(\Delta p) = .025$. 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, we 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.

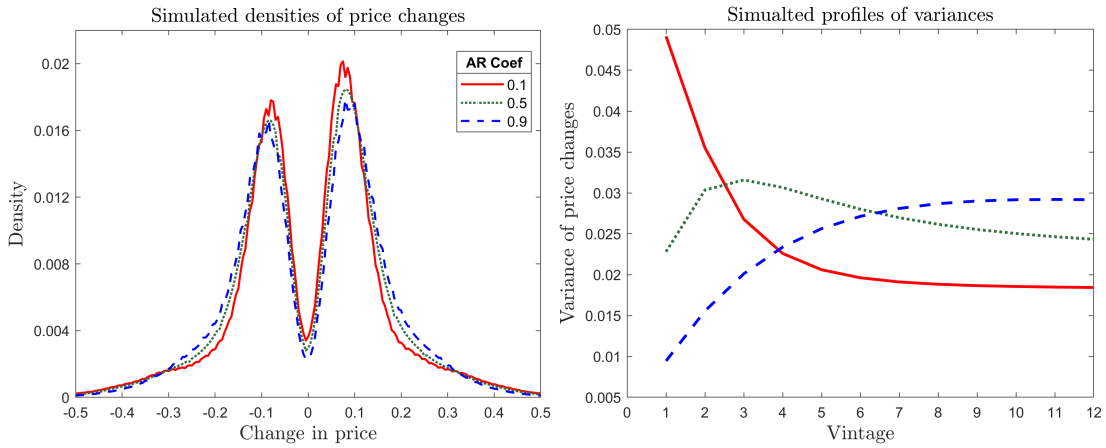

Figure 2: Simulated price change cross-sections

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 prices 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 C.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 be more effective than cross-sectional data at identifying the ideal price process. In some sense, this is not surprising. Vintage-conditional data contains information about both the size of price changes and the duration of price spells, whereas cross-sectional data only contains information about the size of price changes. Data that includes vintage intuitively is strictly more informative. 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.[8] 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

[8] 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.

their price τ periods ago. Vintage data is useful because this distribution of firms states changes in a predicable, recursive manner across vintage. At $\tau = 0$, price changers reset their price gaps according to policy function $x = \bar{x}(a)$. This provides us the 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 change price τ periods ago, but have not yet decided whether to change price today,

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

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 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

We 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). We are particularly interested in vintage-conditional hazard rates and variances, as these are the moments that are heuristically informative about the persistence of idiosyncratic shocks. We document two facts. First, in accord 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.[9] We consider data collected between January 1996 and December 2019. From these data, we want to extract price changes and their associated vintages.

In order to measure vintage, we need to follow the same product over time. ONS data collectors are instructed to record prices for the exact same product over time within a retailer, so we construct product identifiers from the data using product categories and retailer characteristics, and drop those 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, we assign all post-substitution prices to a new product identifier. We drop prices that fail to satisfy ONS-validation procedures, as well as those associated with missing and out of stock products. The remaining prices constitute our sample, which appears to be representative as the year-over-year change in CPI constructed from our sample closely tracks that of the official ONS release (See Figure B.1 in Appendix B).

We 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). We think that sales do not reflect changes in a retailers' ideal prices, so our theory is better suited to describe non-sale or "regular" prices. We identify sales using two approaches. First, we use an ONS variable that flags when a product is on sale or "recovering" from a sale. Second, we 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 rises back to the same level as prior to the initial decrease. We 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, our 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 B.

In this panel of regular prices, we determine vintage as follows. For each panel unit, we 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,

[9] 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 of National Statistics, 2024).

vintage is unknown prior to the first price change of each panel unit, and we do not impute over gaps in the panel. Vintage is unknown for 867,250 price changes, of which 736,208 are the first price change in a panel unit. Restricting our sample to prices and price changes for which vintage is observed may result in selection toward products that change prices more frequently. In Appendix B we 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, we do not think selection on observing vintage is a major concern.

3.2 Vintage-conditional hazard rates.

We 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. We see that the aggregate hazard rates decline in vintage, with spikes at regular four and twelve month intervals.

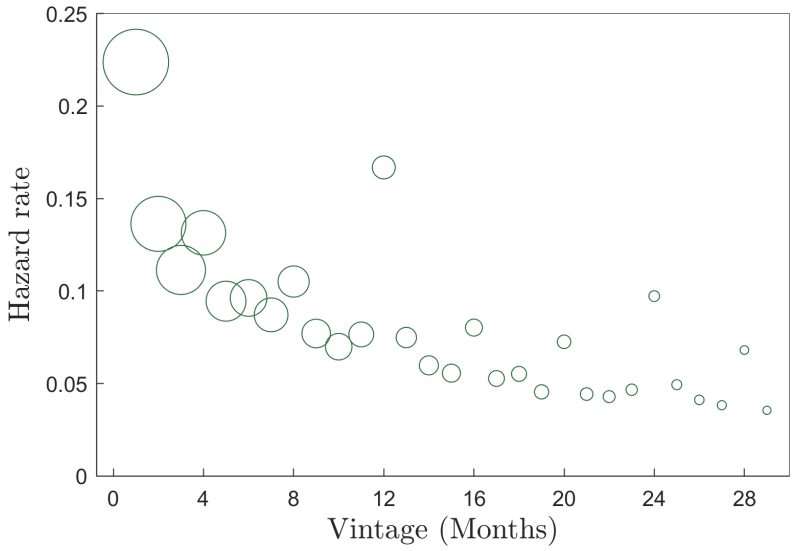

Figure 3: Vintage-conditional aggregate hazard rates

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, we calculate the vintage-conditional hazard rate within each vintage-product category pair[10] and estimate the following relation-

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

ship 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.

	(1) $\lambda_{j,\tau}$	(2) $\lambda_{j,\tau}$	(3) $\lambda_{j,\tau}$
δ_4	0.0105		0.0143
	(0.00138)		(0.00136)
δ_{12}	0.0671		0.0840
	(0.00326)		(0.00325)
β		-0.00192	-0.00274
		(0.0000999)	(0.000100)
Fixed Effects	Product Category	Product Category	Product Category
N	17172	17172	17172
R^2	0.730	0.725	0.742

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, we also observe that prices 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. As labor is the primary input in service industries, and most wage changes occur quarterly or annually, rather than monthly (Grigsby et al., 2021) we interpret this behavior as reflecting a longer time between firm-level cost shocks.[11] When estimating the same model as above with the time scale of vintage changed from every month to every 4 or 12 months, we find a similar declining relationship between hazard rate and vintage. These results are reported in Appendix C.2.

[11] An alternative view is that firms price services using time-dependent rules, à la Taylor pricing.

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, we 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 C.2. To account for the latter, we estimate a specification with vintage-specific fixed effects 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 an overarching downward trend and spikes every 4 and 12 months, we also see see that much of the decline occurs at early vintages.

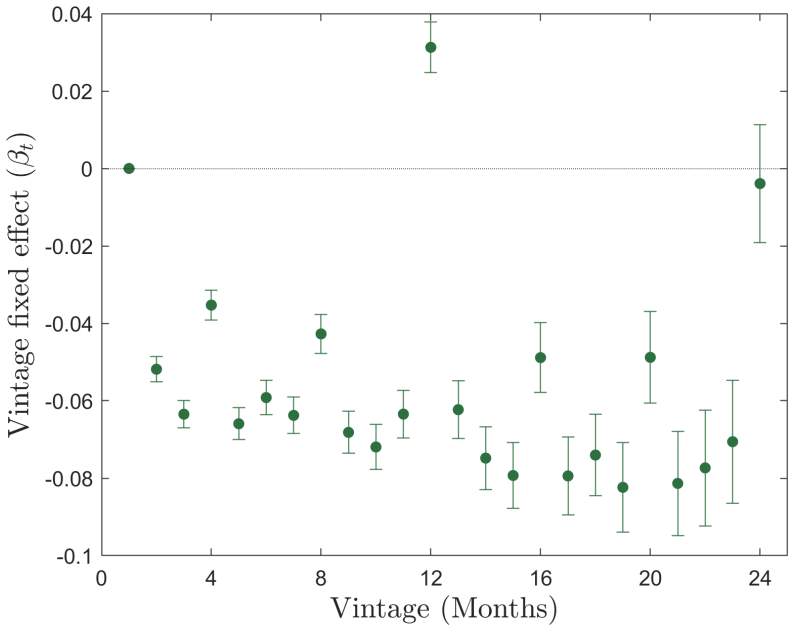

Figure 4: Coefficients from fixed effects regression

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 with relative to $\beta_1 = 0$. Bars represent a 95% confidence interval.

Overall, our 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). We also match the work of Bunn and Ellis (2012a), who measure hazard rates in these same data, although their work focuses on how vintage-conditional hazard rates differ across broad product groups, while we focus on what is common across product categories after removing heterogeneity.

3.3 Vintage-conditional variances of price changes.

We now document the vintage-conditional variances of price changes. To our knowledge, we are the first to give a rigorous examination to this statistic. To construct these variances, we 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 the aggregate variance changing scale across vintage as the composition of products changes. 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, we see that the variance of price changes declines in vintage before leveling off.

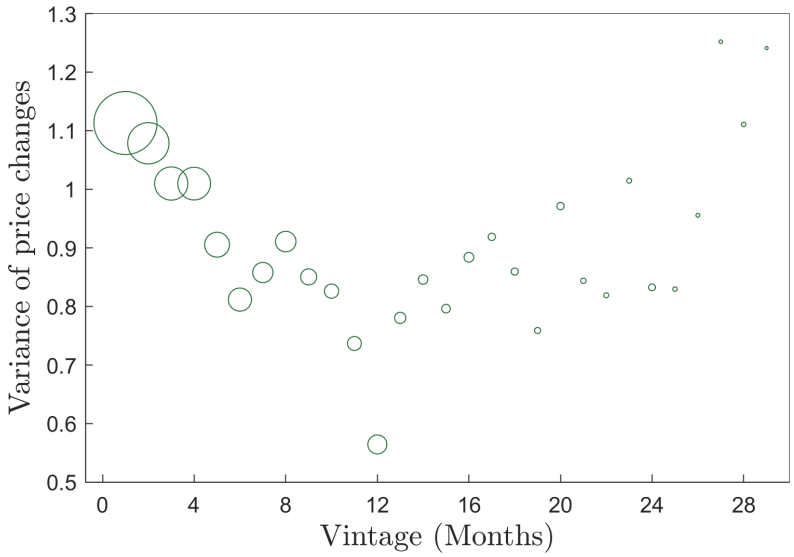

Figure 5: Vintage-conditional variances of price changes: aggregate

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, we want to control for heterogeneity in the relationship between the variance of price changes and vintage across product categories. To do so, we 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 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 than zero at the 10% level. Broadly speaking, the variance of price changes is greater at low vintages than at high vintages.

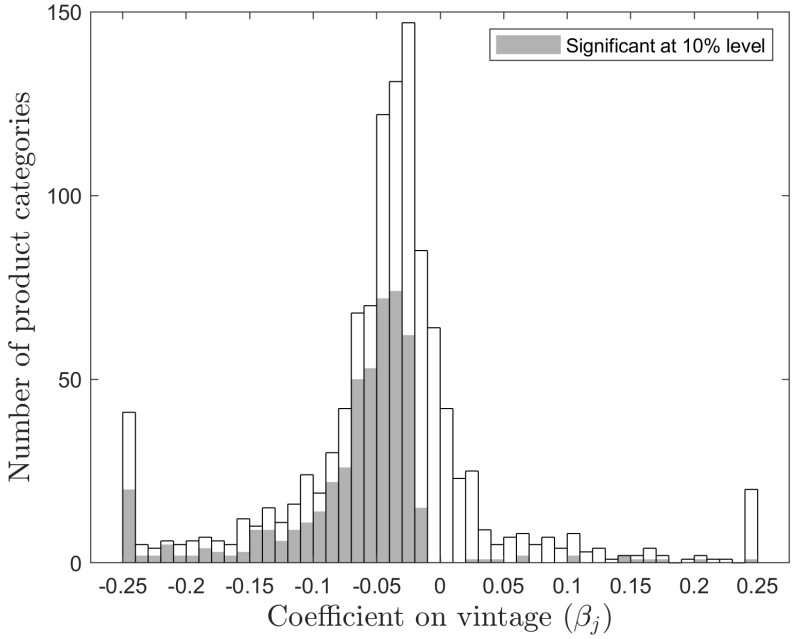

Figure 6: Estimated coefficients on vintage

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 $\bar{\tau}$ we 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, if the variance of price changes is lower among the older vintages than among the younger vintages, we say that the variance of price changes is "locally decreasing" at that cutoff vintage. We define "locally increasing" analogously.

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

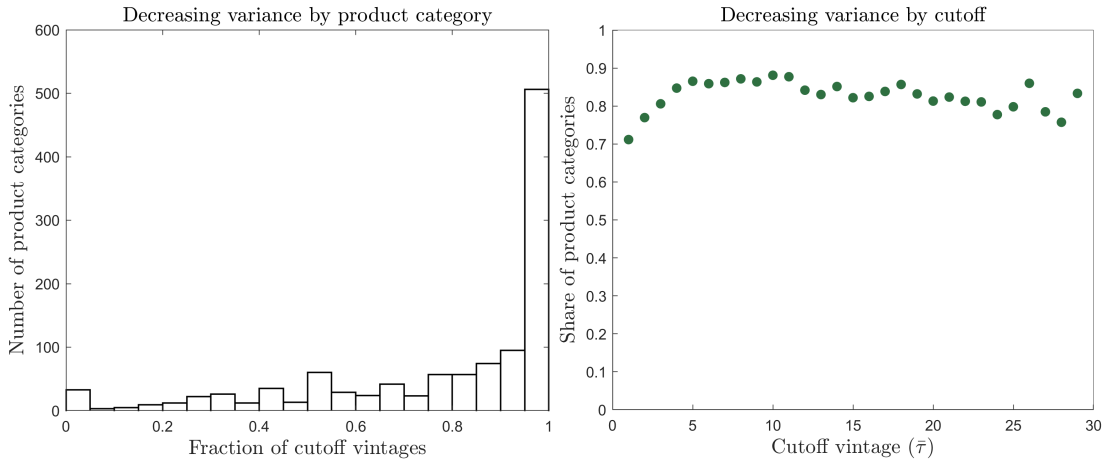

Figure 7: Prevalence of locally decreasing variance of price changes

The left panel is a histogram 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.

From the left panel, we see 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. From the right panel, we see that when the variance of price changes is locally increasing, it is typically at early threshold vintages of 2 to 5 months. All together, the variance of price changes is either monotonically decreasing in vintage, or rising 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 we fail to control for. As stated earlier, we are primarily concerned about unobserved firm types that generate composition bias across vintage.

For each panel unit i in product category j, we observe N_i price changes and vintages. This allows us 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.

We lack the observations to test whether vintage EDFs are identically distributed for any pair of panel units, so instead we test whether the EDF of each panel unit is distributed identically to the pooled EDF of vintage for all other panel units within a product category. In particular, we test whether the means of the two distributions

and the variances of the two distributions are the same.[12] We 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.

We repeat the exercises above on this limited sample and report results in Appendix C.3. Our findings are mostly similar. The one major difference is that, with the linear regression specification, we estimate that hazard rates are very slightly increasing vintage. However, in the specification regressing hazard rates on vintage fixed effects, we still find that the estimated fixed effects decline markedly in the early vintages. We 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. We 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, we now assess the quantitative importance of using them for model estimation. To do so, we move to a simple general equilibrium setting. This is for for two reasons. First, we 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, we want examine what our 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 we use is essentially the canonical New Keynesian model (Galí, 2015), but with nominal frictions in the form of random menu costs.[13]

[12] We 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.

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

An infinitely-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 an 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 are

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

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_t / P_t discounted at rate β. In the absence of menu costs, the firm's static, profit maximizing ideal price would be $P_{it}^* = \frac{\nu}{\nu-1} \frac{W_t}{A_{it}}$. We can thus write the firm profit maximization problem in terms of 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.

We 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.[14]

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

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

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

4.2 Calibration.

We calibrate this model to match features of the aggregate price data in the United Kingdom. The firm pricing decision is determined by seven parameters. The first two, $\{\nu, \beta\}$ are governed by preferences and set to standard values of $\nu = 4$ and $\beta = 0.96^{\frac{1}{12}}$. We set $\pi = .0017$ to equal the observed average monthly inflation rate in the UK from 1996 to 2019. The final four parameters, $\{\rho, \sigma, \lambda, \kappa\}$ are estimated within the model by matching moments.

We compare 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, we 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 and Nakov, 2011).[15]

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

[15] 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).

For the vintage-conditional calibration, we target the path of vintage-conditional hazard rates and variances of price changes for vintages of 1 to 10 months.[16] 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 are presented in Table 3 and Figure 8, respectively.

Parameter	Description	Cross-Section Calibration	Vintage Calibration
ν	Elasticity of Substitution	4	
β	Discount factor (monthly)	$0.96^{\frac{1}{12}}$	
π	Monthly trend inflation	0.0017	
ρ	Autocorrelation of productivity shocks	.94	.24
σ	S.d. of innovation to productivity	.18	.13
λ	Average non-zero menu cost	.52	.25
κ	Probability of free price adjustment	.003	.00

Table 2: Calibrated Parameters

The cross-sectional calibration recovers parameters that are 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, as well as smaller shock innovations and menu costs.

The low autocorrelation of shocks is needed to match the downward sloping profiles of the vintage moments, and the other parameters values follow from there. When shocks are more transitory, it becomes less costly for a firm to keep the same price in the face of 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 price in the face of a shock, so menu costs must be lower on average. Similarly, smaller shock innovations are needed to match the vintage-conditional variances.

Moment	Data	Cross-Section Calibration	Vintage Calibration
Freq of price changes	.088	.089	.099
Share of price decreases	.437	.425	.449
Avg absolute price change	.116	.137	.108
Var of price changes	.025	.023	.016

Table 3: Cross-Sectional Moments

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

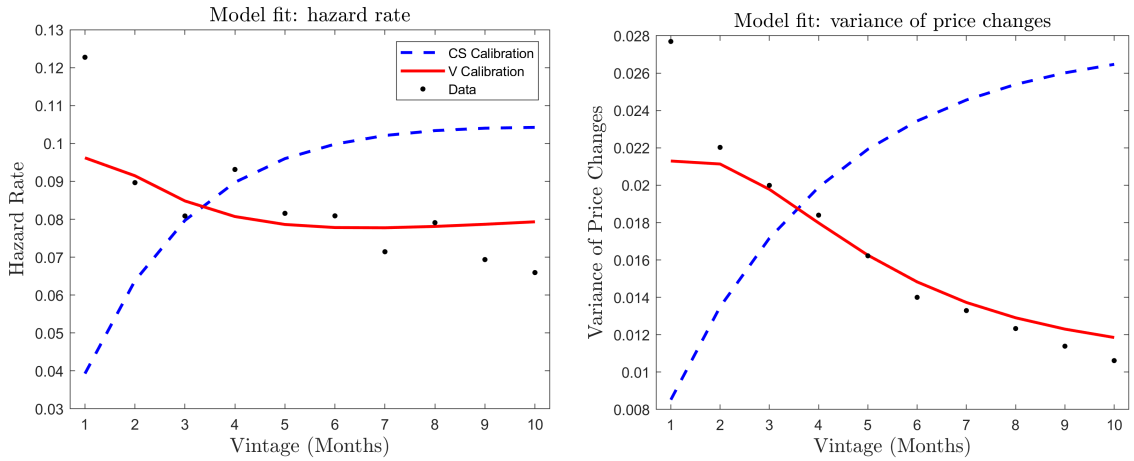

Figure 8: Vintage-Conditional Moments

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. From the simulations in Section 2.2, we saw 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.

On the other hand, it is surprising how well we are able to 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 the variance of idiosyncratic shocks becomes unreasonably large, but perform this analysis assuming that of those idiosyncratic shocks are fairly persistent.[17] We observe that if the mean-reversion of idiosyncratic shocks is sufficiently low, we 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. Our measure of welfare loss is the percentage change in consumption needed equate the welfare of the

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

household in the sticky-price economy to 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.[18] Note the lack of an expectation operator, as there is no aggregate uncertainty.

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

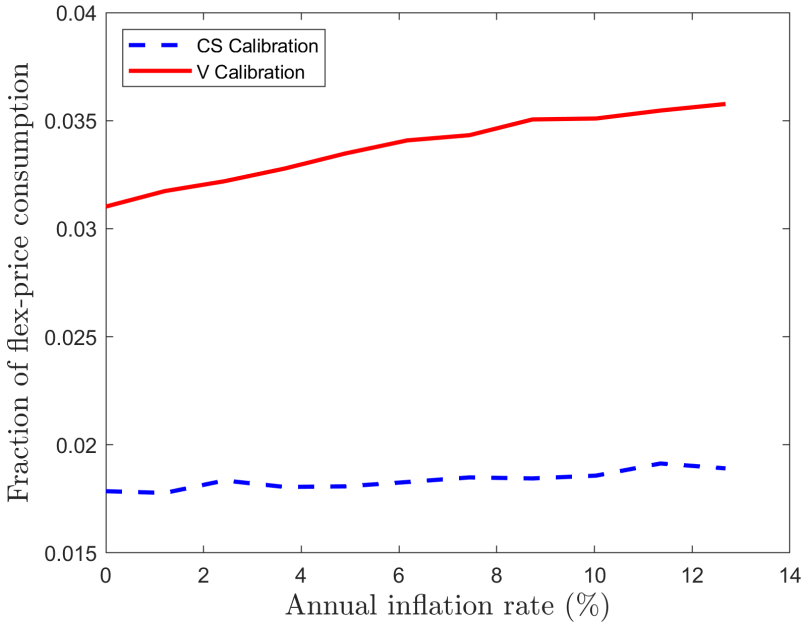

Figure 9: Welfare loss

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.

We focus on two elements of this figure: the levels of welfare loss between the two calibrations, and the slopes of welfare loss with respect to steady state inflation. It is clear that when using parameters form the vintage calibration, the level of welfare loss is both higher and increasing more quickly in steady state inflation.

At the common inflation target of 2%, the model-implied welfare cost of inflation is 1.3 percentage points higher in the vintage calibration than in the cross-section calibration. This is because idiosyncratic shocks are more temporary in the vintage calibration. The incentive to change prices in response to temporary shocks is dampened – firms are simply willing to tolerate larger price gaps, understanding that their current state is likely to revert to the mean in the next period. As a result, price

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

gaps are on average larger in the model calibrated to vintage, and the welfare cost of inflation is greater.

This welfare cost increases about 4 times faster as steady state inflation rises in the vintage calibration compared to the cross-section calibration. To contextualize this difference in slopes, we consider a change in steady-state inflation from 2% to 8%, roughly the same magnitude as the temporary inflation surge from 2018 to 2022. In the vintage calibration, this change would imply an increase in welfare loss that is 0.28 percentage points greater than the increase in welfare loss implied by the cross-section 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. The 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 C.5 we decompose welfare loss into the contribution made by each of these channels. The majority of the difference in our two calibrations stems from the allocative efficiency channel, rather than from the labor paid to menu costs.

5 Validation with Cost Data

In the previous section, we show that estimating models using vintage conditional moments implies a low autocorrelation of idiosyncratic shocks, which in turn leads to higher welfare cost of inflation. This low persistence of shocks is somewhat surprising, although not unprecedented.[19] In this section, we exploit a unique dataset which links firms' costs, production, and prices to support this finding and validate the use of vintage moments in estimation.

The data we use for this exercise is the price micro-data underlying the Belgian producer price index (PPI), collected by Statbel, and linked with the Belgian PRODCOM database, VAT declarations, and social security wages. This creates a monthly panel of manufacturers' prices, production, revenues, and costs. Although the PRODCOM dataset has been studied previously in a similar context to ours,[20] we are the first to link it with PPI prices. This allows us to estimate a stochastic process for ideal prices using firm characteristics and cost data, and then compare the parameters of that process to the ones recovered by our calibration procedure using only price data.

[19] Eichenbaum et al. (2011) recover a Markov process for supermarket costs with a monthly autocorrelation of under .2.

[20] See Gagliardone et al. (2023, 2025).

5.1 Data description and cleaning.

The PPI micro-price data contains monthly price quotes from February 2016 to January 2023, collected from Belgian manufacturing firms selling to the domestic market. Price quotes are collected from wholesaler catalogs, online postings, and surveys sent to firms. The sampling procedure is designed to be representative of total production in each manufacturing product category. Products are followed at the barcode level, and changes in product definition are flagged. We restrict our analysis to unchanged products.

We construct price change and vintage variables in the same manner as for the UK CPI data, excepting the lack of a sale filter, as sales are uncommon in business-facing producer prices.[21] In total, the sample consists of 8,687 products in 706 categories, with 414,828 price quote observations and 90,358 price change observations. Of these, there are 298,810 price quote observations and 80,480 price change observations where vintage is observed.

Price quotes are linked to firms, identified by VAT declaration number. This allows us to link prices with detailed information on firms' real activity (sales and quantities of production) from PRODCOM, and with firms' costs from VAT declarations and social security filings. Appendix B.3 provides more information about each of these data sets and harmonization between them. All together, this gives us a monthly panel of prices at the firm-product level, quantities and revenues at the firm-product category level,[22] and costs at the firm level.

5.2 Vintage conditional moments.

In this section, we document vintage-conditional moments in the PPI data. Figure 10 plots the aggregate hazard rates and variances of price changes against vintage. Their profiles mirror the CPI data. Hazard rates decline in vintage, and variances of price changes peak at a vintage of 2 before subsequently declining.

[21] Applying a V-shaped sales filter to the PPI changes fewer than 1% of price quotes, compared to over 8% of price quotes in the CPI data.

[22] Product categories are defined by 8-digit product codes, where the first 4 digits are the industry. For example, industry 15.20 is "Manufacture of footwear," and contains products codes for " Town footwear with rubber or plastic uppers" (15.20.12.31), "Men's sandals with leather uppers" (15.20.13.62), "Women's sandals with leather uppers" (15.20.13.62), etc.

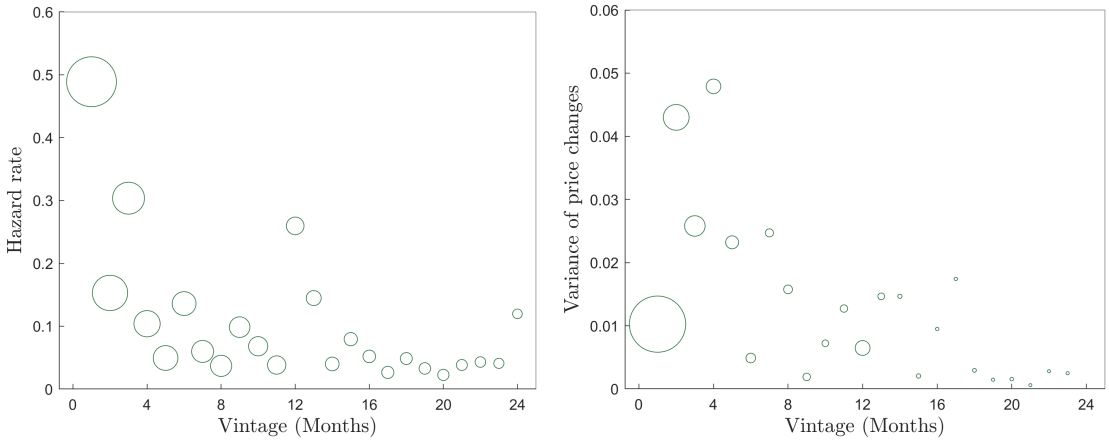

Figure 10: Belgian aggregate vintage-conditional statistics

The left panel plots the vintage-conditional hazard rate against vintage The right panel 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.

We cannot replicate all the results presented in Section 3 because there are far fewer observations and panel units in the PPI price data than in the CPI data. However, because panel units remain in the PPI data for a long time, we can instead perform exercises that fully control for unit-level heterogeneity. This completely eliminates concerns about unobserved heterogeneity.

For panel units with at least 20 price change observations, we estimate regressions akin to those of Section 3, but with two differences. First, we estimate fixed effects and random effects at the panel unit level rather than the product category level. Second, we include an indicator variable for vintage $\tau = 1$. This is because some panel units lack sufficient observations across all vintages, which makes estimating a linear trend in vintage difficult. Results are reported in Appendix C.4, and comport well with what we find in the CPI data. Declining hazard rates and variances of price changes are robust across datasets. We therefore believe there is validity in comparing the outcome of our model calibration in Section 4 to statistics estimated from the Belgian cost data.

5.3 Estimating a process for costs.

As noted previously, in these data we have the rare benefit of observing firm costs, production, and revenues in addition to prices. This makes it feasible to directly estimate a stochastic process for firms' nominal marginal costs.

We assume all products i sold by firm f in product category j are produced using a constant returns to scale technology with a common marginal cost. Marginal costs are thus equal to average variable costs $mc_{if,t} = avc_{jf,t}$. We measure average variable cost within a product category by assigning shares of observed firm level total variable costs to product categories by revenue shares, and dividing by observed product

category level quantity produced.[23]

This assumption stipulates that multi-product firms face the same supply fundamentals among all products in their portfolios. This is sensible insofar as a firm's production processes are common to all products it produces. As a robustness check, we perform the estimation procedure below on a subsample of firms that manufacture products in the same PRODCOM product category, and on a subsample of single product firms.

We further assume nominal marginal cost can be written as

$$\begin{aligned} mc_{if,t} &= \bar{\alpha}_{if} + \pi t + \delta_t + \alpha_{if,t}, \\ \alpha_{if,t} &= \rho\alpha_{if,t-k} + \varepsilon_{if,t}, \quad \varepsilon_{if,t} \sim \mathcal{N}(0,\sigma), \end{aligned}$$

where $\bar{\alpha}_{if}$ is product i's steady state real marginal cost, δ_t is an economy-wide cost (supply) shock, π is the trend inflation rate, and $\alpha_{if,t}$ is a product-level productivity shock that follows an AR(1) process.

Substituting average variable costs for marginal cost, we can rewrite this as a single equation in terms of observables.

$$avc_{if,t} = \underbrace{(1-\rho)\bar{\alpha}_{if} + \rho\pi}_{\hat{\alpha}_{if}} + \underbrace{(1-\rho)\pi}_{\hat{\pi}} t + \underbrace{\delta_t - \rho\delta_{t-1}}_{\hat{\delta}_t} + \rho avc_{it-1} + \varepsilon_{it}$$

Estimating this equation is more complex than it appears, due to our panel data containing many more panel units than time periods. Rather than take a stand about which identifying assumptions produce an estimator with finite sample properties best suited to our particular setting, we instead take two simple OLS estimates of ρ with known directional bias to produce bounds on ρ. We do so for each NACE 4-digit manufacturing industry.

The lower bound is given by

$$avc_{it} = \alpha_i + \hat{\rho}_{lb} avc_{it-1} + \delta_t + \varepsilon_{it}.$$

This estimate for ρ is biased downward due to correlation between the estimated fixed effect α_i and the shock realization ε_{it}. The upper bound is given by

$$avc_{it} = \alpha + \hat{\rho}_{ub} avc_{it-1} + \delta_t + \varepsilon_{it}.$$

This estimate for ρ is biased upward due to the exclusion of firm fixed effects. Histograms of these bounds are plotted in Figure 11. There is variation in both the level

[23] This is, in logs, $avc_{jf,t} = tvc_{f,t} - r_{f,t} + r_{jf,t} - q_{jf,t}$, where $tvc_{f,t}$ is total variable cost, $r_{jf,t}$ is revenues in product category j, and $q_{jf,t}$ is quantity produced in product category j.

and the tightness of bounds across industries. The median lower bound across industries is .27, while the median upper bound is .85 – we take these as benchmark bounds for the aggregate economy.

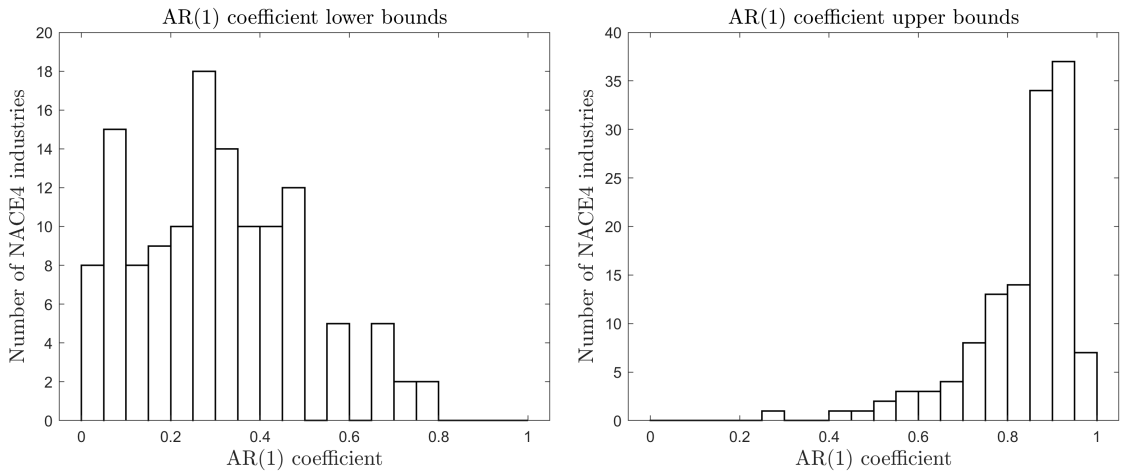

Figure 11: Estimated AR(1) coefficient bounds

The left panel plots a histogram of the estimated lower bounds for the AR coefficient of each Belgian NACE 4 manufacturing industry. The right panel plots a histogram of the estimated upper bounds for those same AR coefficients.

For both calibrations of Section 4, the estimated autocorrelations of firm shocks fall outside of these bounds. However, our estimate of $\rho = .24$ from the vintage calibration is closer to the lower bound than our estimate of $\rho = .95$ from the cross-section calibration is to the upper bound.

How do the autocorrelations of ideal prices implied by other sticky price models compare to these bounds? In many cases, ideal prices are simply assumed to follow a random walk. This is not reproduced in our data – a unit root is rejected in all industries. Among previously estimated firm pricing models that admit an AR(1) process for ideal prices, estimates tend to be closer to the upper bounds. Golosov and Lucas (2007) and Nakamura and Steinsson (2008) each calibrate a fixed menu cost model to match features of cross-sectional distribution of price changes that underlie the US CPI. They find AR coefficients of .58 and .7, respectively.[24] Midrigan (2011) develops a model designed to reproduce the large, temporary price changes associated with sales. This is achieved by assuming a two-state ("normal" and "low cost") Markov process for marginal costs. His benchmark calibration results in an autocorrelation of .40 for ideal prices, which toward the center of our bounds. Notably, though, his setting is expressly concerned with matching the behavior of sales, whereas these

[24]Differences in estimates reflect the different moments used in the calibration strategies. Golosov and Lucas (2007) match the frequency of price changes, average size of price increases, and variance of new prices; Nakamura and Steinsson (2008) match the frequency of price changes, average absolute size of price changes, and the share of price changes that are increases. The sensitivity of these estimates the set of cross-sectional moments used further supports our position that cross-sectional moments do not well identify the ideal price process.

find evidence of temporary ideal price shocks even when examining regular prices. Eichenbaum et al. (2011) use data on retailers' prices, revenues, and costs to recover a 3-state Markov process for marginal costs and a 12-state Markov process for ideal prices in a fixed menu cost model with price plans. They recover a serial correlation of .17 in ideal prices and of .14 in marginal costs when calibrating to match the frequency and volatility of both cost changes and price changes.

Overall, even though the autocorrelation of idiosyncratic shocks recovered by our vintage estimation seems low at first blush, it close to the bounds we estimate from the firm cost data. At a minimum, our estimate is no less reasonable than assuming ideal prices follow a random walk. Furthermore, we must note that our bounds are on the autocorrelation of idiosyncratic shocks to firms' costs. Fluctuations in firms' ideal prices which determine, their price setting behavior, are likely also determined by demand side factors that are not captured here. If shocks to firm-level costs are in actuality quite persistent, then vintage-conditional moments suggest that other components of the ideal price process are transitory, and play a large role in driving firms' pricing behavior.

6 Conclusion

This paper argues that vintage-conditional moments should be matched when estimating sticky-price models. We demonstrate in a simple example that these moments identify the persistence of idiosyncratic shocks to firms' ideal prices, and derive a heuristic to determine whether those shocks are permanent or transitory. Using the price micro-data underlying the UK CPI and Belgian PPI, we document that hazard rates and variances of price changes decline in vintage, both of which indicate more transitory idiosyncratic shocks. Compared to a random menu-cost model estimated in the typical manner, 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 significantly more quickly as inflation rises. We validate the output of our estimation procedure using linked price and cost data and confirm that the low persistence of idiosyncratic shocks is empirically plausible.

7 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," techreport 18088, CEPR Press April 2023.

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*, nov 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.

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*, jan 2020, *61* (1), 159–188.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo, "Reference Prices, Costs, and Nominal Rigidities," *American Economic Review*, February 2011, *101* (1), 234–262.

Gagliardone, Luca, Mark Gertler, Simone Lenzu, and Joris Tielens, *Anatomy of the Phillips Curve: Micro Evidence and Macro Implications* June 2023.

__ , __ , __ , and __ , *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*, aug 2009, *124* (3), 1221–1263.

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

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.

__ , **Magali Marx, and Paul Vertier**, "How Do Gasoline Prices Respond to a Cost Shock?," *Journal of Political Economy Macroeconomics*, December 2023, *1* (4), 707–741.

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.

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 Disaggregat 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 Dawit Zerom, "Accounting for Incomplete Pass-Through," *Review of Economic Studies*, December 2009, *77* (3), 1192–1230.

__ **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 of National Statistics, "Consumer Prices Indices Technical Manual, 2019," Technical Report, United Kingdom Office of 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, "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 presents proofs underlying the propositions given in Section 2.2.

Setup. Starting with the model presented in Section 2, we 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, we assume that we are either in 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 we can then derive hazard function

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

Proposition 3. *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.*

Proof. 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 $k_0 = 0$, we have that sequence $\{k_i\}$ is monotonically increasing and bounded above by $\bar{k} = \frac{\sqrt{\sigma^4 + 4K\sigma^2} - \sigma^2}{2}$, and that 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

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

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

Taking these together, in order to match data $\hat{\lambda}, \widehat{Var}(\Delta p)$ we 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

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

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 $\{\sigma_p, K_p\}$ that rationalize the data in the permanent shocks economy.

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

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

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

conditions:

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

$$\hat{f}_t(x,a) = \frac{1}{2\pi\sigma^2}\exp\left(-\frac{(x+a)^2+a^2}{2\sigma^2}\right)$$

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 $\hat{\lambda}$, $\widehat{Var}(\Delta p)$ we can set $\sigma_t^2 = \frac{Var(\Delta p)}{2(\lambda^2-3\lambda+3)}$ and $K_t = \frac{2(1-\lambda)^2\sigma_t^2}{\lambda(2-\lambda)}$. □

Proposition 4.

- *In the permanent shocks model, the vintage-conditional hazard rate and the vintage-*

conditional variance of price changes are both monotonically increasing in vintage.

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

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:

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

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

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

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

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

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

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

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

with

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

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

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

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

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

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

By differentiating $\lambda_{t,\tau}$ and $Var_{t,\tau}(\Delta p)$ with respect to τ, we can algebraically show that both derivatives are negative for $\tau \geq 1$. Therefore, $\lambda_{t,\tau}$ and $Var_{t,\tau}(\Delta p)$ are both decreasing in τ.

A.2 GE Model

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

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

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, we 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 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 we 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

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

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

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

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_t/P_t discounted at rate β. Defining price deviations $X_{it} = \frac{P_{it}}{P_{it}^{*}}$, 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 $B = \bar{C}^{2-\nu}(\frac{\nu}{\nu-1})^{1-\nu}$ and $\hat{\kappa} = (\frac{\nu}{\nu-1}\bar{C})^{\nu-1}\kappa$, and $\bar{C}$ is steady-state consumption.

The stick-price equilibrium is givne 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 stationary distribution over firms of $F = (X_i, A_i)$. We will use the notation $dF(i)$ to denote integrating over firms with respect to this distribution.

To contrast with the flexible price economy, we 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

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

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

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

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

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

B Data appendix

B.1 UK data: Additional details on sample construction

We 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 prices are linked over time panel, but we 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, we drop the observations which cannot be uniquely attributed to a panel unit.

Given this initial panel, we 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. We 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, we 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, assuring that sampled item is exactly the same month to month, and flagging internal inconsistencies or missing fields in collected data. More info in the ONS Technical Manual (Office of National Statistics (2024)). The observations that remain are the final sample of price quotes. With these, we follow ONS procedure to recon-

struct the CPI index. Figure B.1 compares the official and and in-sample CPI. Table B.1 reports the number of observations remaining through each step.

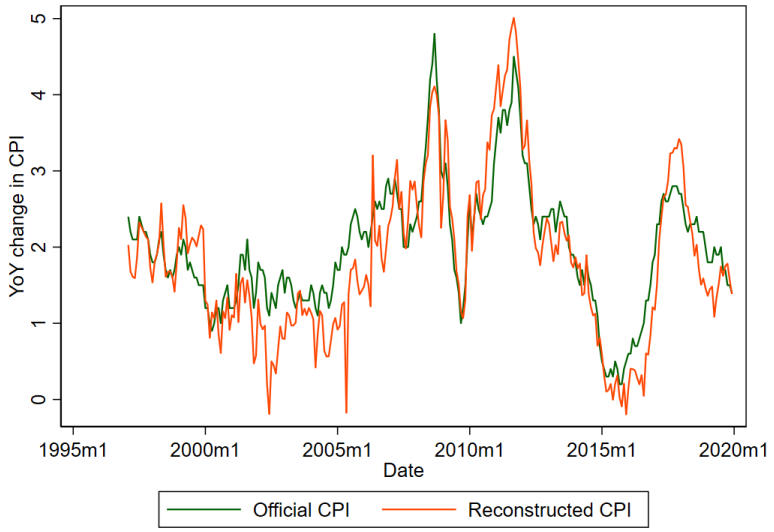

Figure B.1: UK CPI Inflation, Official vs Constructed

Initial number of price observations...	39,526,450
...removing without unique panel unit:	33,221,389
...removing invalid and missing:	31,431,675
...removing with missing price change:	23,099,634
...removing with missing vintage:	9,992,628

Table B.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, we drop observations where the price change is missing, most of which are the first observations of panel units, and we drop price changes that are greater than a factor of 10 on the grounds that these are implausibly large. We also drop observations from after 2019. In the sample with vintage, we 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 B.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 B.2: Dimensions of UK data

B.2 Summary statistics: price change and vintage

	All price changes	With vintage
Frequency of price changes	.0903	.121
Average price change	.00844	.00955
Average abs price change	.129	.117
Variance of price changes	.0391	.0328
Kurtosis of price changes	8.65	9.21
Observations	2,085,261	1,215,488

Table B.3: UK Data Summary Statistics

B.3 Belgian data: Additional description of data sources

PRODCOM Data The PRODCOM survey, commissioned by Eurostat and administered in Belgium by the National Statistical Agency, is designed to cover at least 90% of production value within each NACE 4-digit manufacturing industry by surveying all firms operating in the country with either at least 20 employees or total revenue above 4.5 million euros. The sampling design assures that large firms (with size over some industry-specific threshold) are always surveyed, while smaller firms are selected at random.

Firms are required to disclose, on a monthly basis, product-specific physical quantities of production sold (in volume, kg., m2, etc.) and the value of production sold (in euros) for all their manufacturing products, which are defined by an 8-digit PRODCOM product code.

We drop firms for which reported revenues in PRODCOM are less than 70% of those disclosed to the VAT authority. This assures that our sample encompasses firms whose primary activity is manufacturing, and therefore that the costs we measure are primarily attributed to manufacturing activity.

PPI Data The PPI surveys firms on a monthly basis about the price set for the products sold in the last month. Its sampling strategy is similar to that of the PRODCOM survey. The population of firms producing a good is stratified by size, such that larger firms are more likely to be sampled than smaller ones. Unlike PRODCOM, though, the largest firms are not certain to be sampled, and conversely, smaller firms are more likely to be sampled, compared to the PRODCOM survey. Also, firms are surveyed about all products that are produced, so multi-product firms may be over-sampled. Products are defined at the barcode level.

Prices are collected for products attributed to one of three markets: the domestic market (Belgium), the non-domestic market (exports to the Euro Area), and the global market (exports outside the Euro Area). We retain only prices for products sold to

the domestic market. This is not costly in terms of observations, as relatively few firms persistently export, and allows us to avoid concerns about differential demand shocks and exchange rate variation driving price dynamics. In the lens of the model, firms selling to the domestic market more plausibly have the same stochastic process for ideal prices.

In cases where there is a continuous spell of missing price data, but the price observations surrounding that spell take the same value, we impute the missing price spell with that value. Also, as in the UK data, we exclude price changes that exceed a factor of 10.

Wages and VAT declarations. To measure firm-level costs, we supplement our main datasets with two additional sources: Social Security declarations filed with the Department of Social Security of Belgium, and VAT declarations filed with the Federal Public Service.

Social Security declarations contain information on employment and labor costs. For our measure of variable labor costs, we use a wage concept devised by the social security office that is intended to tightly follow production. This measure is equal to total wage bill minus components that do not necessarily co-move with production activity, such as holiday pay, waiting salary, reimbursements, and other pecuniary benefits.

These data are quarterly, so to find monthly wages we divide the quarterly wage bill equally across the three months of the quarter and assume an equal headcount in each month. We view this assumption as innocuous in light of tight labor market regulations in Belgium, such as restrictions on hiring or firing. Furthermore, manufacturing firms typically adjust short run labor force through contracts with temporary employment agencies, rather than through direct hiring. In such a contracting setup, expenses on temporary workers' wages will be measured as a component of the monthly VAT declaration of the manufacturing firm as a payment to the employment agency.

VAT declarations contain firm-level information on total revenues and on purchases of raw materials and other goods and services that are VAT-liable. These include both domestic and international transactions. Firms are required to submit this information on a monthly basis, except for those with an annual turnover of under 2.5 million euros (excluding VAT), which are permitted to submit on a quarterly basis.

Both datasets cover the universe of Belgian firms, and the only restriction we make is to exclude the firms which file VAT on a quarterly rather than a monthly basis. Given the the stratified sampling strategy of PRODCOM, this has little impact on the composition of our sample. We sum the wage bill with intermediate purchases

to generate total variable costs, and adjust this value by the ratio of reported revenues in VAT and PRODCOM.

C Supplemental figures and tables

C.1 Theory and simulation: supplemental figures

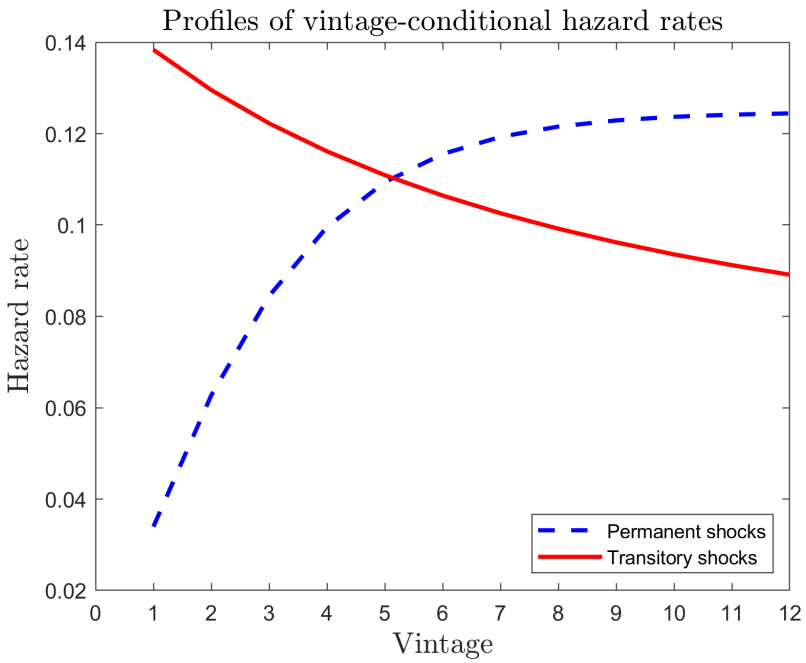

Figure C.1: Profiles of vintage-conditional hazard rates

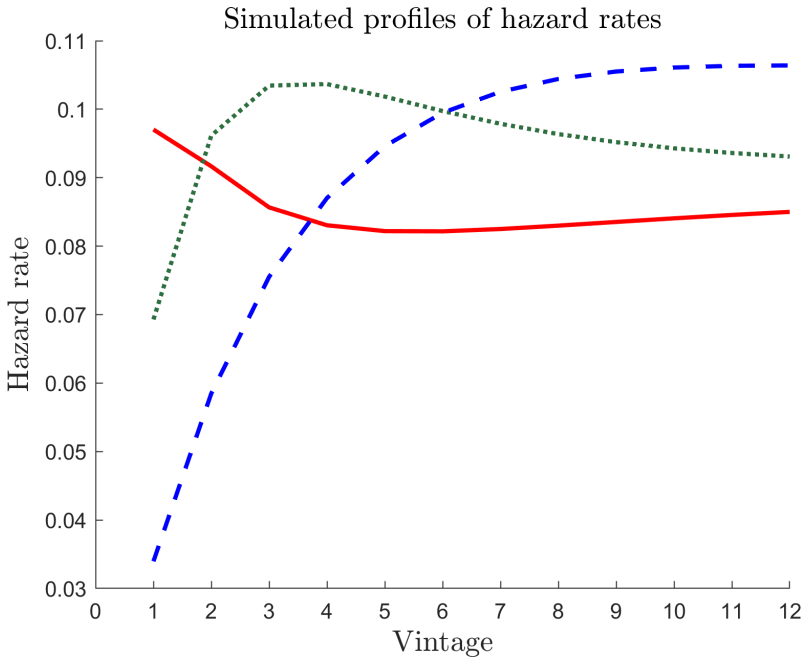

Figure C.2: Hazard rates

C.2 UK CPI: supplemental figures and tables

	(1) $\lambda_{j,\tau}$	(2) $\lambda_{j,\tau}$
τ	-0.00584	-0.0134
	(0.000307)	(0.000944)
Fixed Effects	Product Category	Product Category
Aggregation	Quarterly	Annual
N	6808	2266
R^2	0.854	0.855

Standard errors in parentheses

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

C.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	.0284
Kurtosis of price changes	9.21	10.22
Observations	1,215,488	555,416

Table C.2: Summary Statistics - removing unobserved heterogeneity

	(1) $\lambda_{j,\tau}$	(2) $\lambda_{j,\tau}$	(3) $\lambda_{j,\tau}$
τ_4	0.0177		0.0162
	(0.00200)		(0.00201)
τ_{12}	0.107		0.0994
	(0.00496)		(0.00506)
τ		0.00240	0.00121
		(0.000174)	(0.000176)
N	10822	10822	10822
R^2	0.699	0.682	0.700

Standard errors in parentheses

Table C.3: Vintage and hazard rate - no heterogeneity

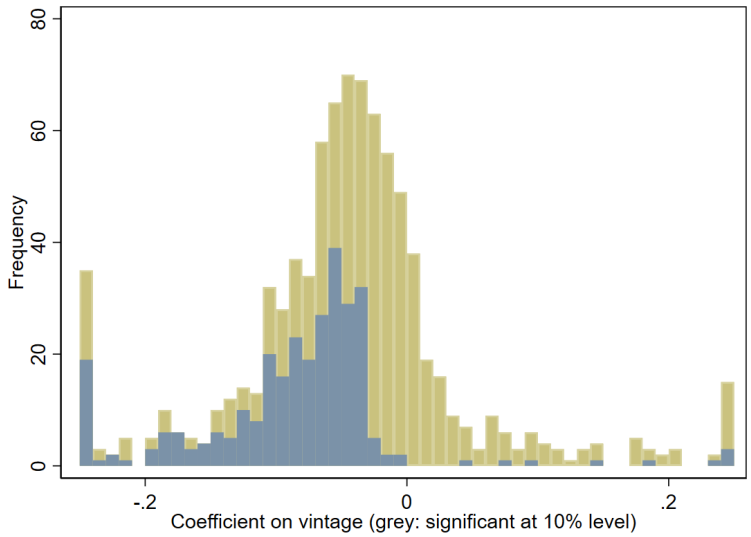

Figure C.3: Coefficients of variance vs vintage

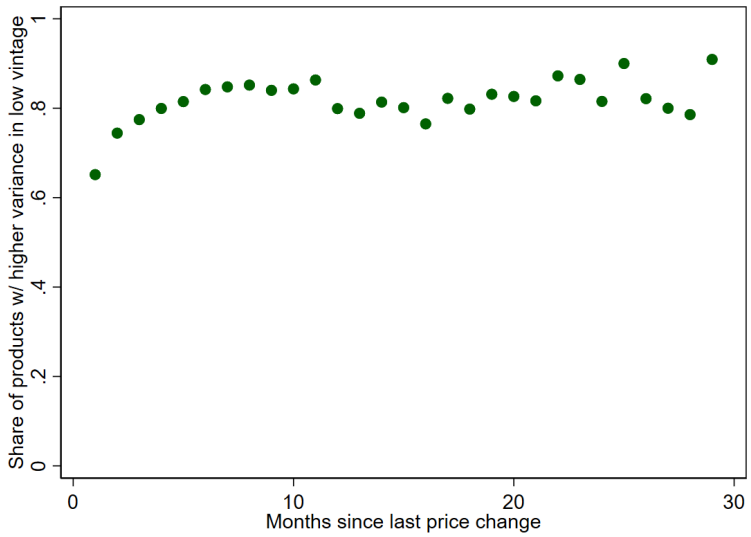

Figure C.4: Prevalence of variance decreasing in vintage

C.4 Belgian PPI: supplemental figures and tables

	(1) $\lambda_{i,\tau}$	(2) $\lambda_{i,\tau}$	(3) $\lambda_{i,\tau}$	(4) $\lambda_{i,\tau}$
τ_3	-0.0599		-0.0578	
	(0.00421)		(0.00416)	
τ_{12}	-0.0295		0.0278	
	(0.0282)		(0.0285)	
τ		-0.00841	-0.00778	
		(0.000822)	(0.000823)	
$\mathbb{1}\{\tau = 1\}$				0.0941
				(0.00667)
N	4865	4865	4865	2308
R^2	0.908	0.905	0.910	0.845

Standard errors in parentheses

Table C.4: Vintage and hazard rate - within panel unit

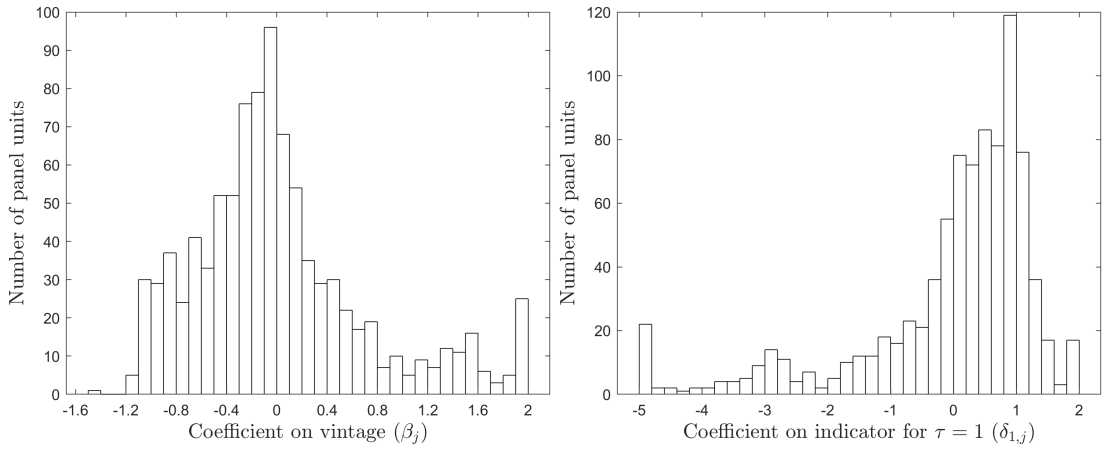

Figure C.5: Vintage and variance of price changes - within panel unit

C.5 Additional figures from Calibrated Model

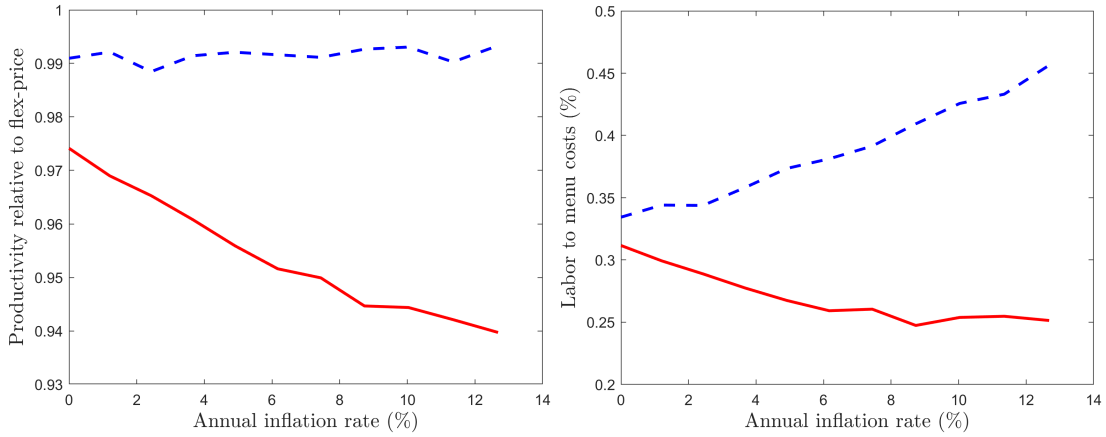

Figure C.6: Decomposition - sources of welfare loss