Identifying the Impact of Inflation Expectations

William A. Branch*
University of California, Irvine

February 4, 2026

Abstract

Individuals form inflation expectations differently based on their demographics and locations. Consequently, different demographic groups respond differently to sectoral price changes, a fact that we exploit to identify the inflationary impact of expectations. Our instrument combines national expectations of specific groups with these groups' share in regional populations. We find that a one-percentage-point rise in the expected rate of inflation increases (regional) inflation by 60 basis points. Interestingly, long-run expectations – say, over 5 to 10 years – don't seem to matter much. The estimates are most robust for a particular demographic: younger, married individuals holding at least a high school diploma. Their expectations mainly influence the prices of non-durable goods.

JEL Classification: D82; D83; E40; E50

Keywords: expectations, inflation, survey data.

1 Introduction

How do subjective inflation expectations impact inflation rates? When people expect prices to rise, they may act in ways that make prices rise. Businesses might set higher prices in anticipation of future cost increases, workers could demand higher wages, and consumers may choose to buy now rather than later. These behavioral channels are central to modern macroeconomic theory, yet the extent to which they operate in practice remains an open empirical question. Recent micro evidence provides mixed results, and it is still unclear how large a role inflation expectations actually play in driving realized inflation.

*Complete replication and robustness files available at https://github.com/William-Branch/groupexpects. I am grateful for comments and suggestions from Mary Burke, Michael Weber, Chris Gibbs, Yuriy Gorodnichenko, Stefan Nagel, Ivan Werning, Jang-Ting Guo, Mari Tanaka, and seminar participants at U.C. Riverside, the RBA, and the 2023 Workshop on Expectations in Dynamic Macro Models/Barcelona Summer Workshop.

One way to gauge the tie between expected and actual inflation is to measure the correlation between inflation rates and survey measures of inflation expectations. Since 1978, the University of Michigan Survey of Consumers has been eliciting monthly inflation expectations from a nationally representative sample of consumers. The full sample correlation between year-over-year changes in the consumer price index (CPI) and what the survey's mean response predicted is 0.18 after controlling for aggregate unemployment and lags of inflation, indicating a positive but modest relationship.

An economic interpretation of this correlation between inflation expectations and inflation outcomes is difficult because economic theory predicts that inflation expectations are endogenous. The rational expectations hypothesis holds that expectations are functions of the same state variables as the data-generating process, and forecast errors are unpredictable. Models of non-rational beliefs predict that expectations will move, partly, with the shocks driving inflation even if beliefs are biased and forecast errors predictable (Mavroeidis, Chevillon, and Massmann (2009)).

The standard empirical strategy to disentangle the causal effects of expectations employs the New Keynesian Phillips Curve, instrumenting for inflation expectations. Yet, this approach faces limitations. First, most derivations of the Phillips curve intrinsically depend on rational expectations, a strong theoretical and empirically contestable assumption. Second, identification challenges, coupled with the weak instrument problem, pose additional hurdles to econometric inference (Mavroeidis, Plagborg-Moller, and Stock (2014)).

This paper, instead, focuses on inflation outcomes by exploiting cross-sectional variation that provides a credible estimate of the impact from inflation expectations. Estimates derived from aggregate data may not capture dimensions of cross-sectional heterogeneity that are important in price-setting. For example, there is evidence that the extent of price stickiness varies across sectors (Cravino, Lan, and Levchenko (2020); Boivin, Giannoni, and Mihov (2009); Almas (2012)). Consumers experience inflation differently depending on their market baskets and varied prices for the same goods (Kaplan and Schulhofer-Wohl (2017), Hobijn and Lagakos (2005)). Differentiated market baskets lead to heterogeneity in inflation expectations (Angelico and Di Giacomo (2022)). Expectations are heavily influenced by the prices in a consumer's market basket, for instance, grocery store prices (D'Acunto, Malmendier, Ospina, and Weber (2021); Angelico and Di Giacomo (2022)). Different demographic groups have distinct market baskets and inflation expectations (Bryan and Venkatu (2001); D'Acunto, Malmendier, and Weber (2021); de Bruin, Van Der Klauuw, Downs, Fischhoff, Topa, and Armantier (2010); Das, Kuhnen, and Nagel (2020)). Similarly, firms' inflation expectations also reflect their exposure to sector-specific prices (Andrade, Coibion, Gautier, and Gorodnichenko (2021)). Using the University of Michigan Survey of Consumers' micro-data, this paper creates an extended panel to examine how inflation expectations across demographics and regions influence inflation rates.

A simple model of firm pricing motivates the empirical approach. In this model, consumers are not all the same: they prefer different baskets of goods, and these goods come from various sectors.[1] Within each sector, firms are monopolistic competitors and can only change their prices occasionally, as in Calvo (1983) and Woodford (2003). Now, the frequency a firm gets to change its price is not the same across all sectors.

What comes out of this setup? Three main points. First, because consumers have their own preferred baskets of goods they are going to have different expectations about where inflation is heading. Second, we observe consumer expectations as a weighted average of sector-specific inflation expectations. Third, aggregate inflation turns out to be an average of these different group

[1]Formally, heterogeneity enters the model through iid preference shocks across groups, leading to distinct sectoral expenditure weights. Empirically, similar cross-group differences in expectations could also arise from variation in attentiveness, literacy, or information acquisition. The analysis is agnostic as to the precise source of this heterogeneity.

expectations. For this analysis we are not making specific assumptions about how exactly people form these expectations. The only assumption is that group expectations reflect the underlying expected sectoral inflation rates.

The core focus of this paper is to estimate how inflation expectations influence inflation outcomes. Our identification strategy exploits heterogeneity across demographic groups, which shapes distinct inflation expectations and, in turn, regional differences in average expectations. The design exploits a quasi-experimental source of variation: regional demographic composition interacts with group-specific movements in expectations, yielding a shift-share ("Bartik-style") instrument that isolates plausibly exogenous differences in expected inflation across regions (following Bartik (1991), Blanchard and Katz (1992), Almas (2012), and Goldsmith-Pinkham, Sorkin, and Swift (2020)).

To clarify the intuition behind this design, consider two illustrative scenarios. First, suppose monetary policy unexpectedly tightens. Households that are young, college-educated, and preparing to buy homes tend to follow interest-rate news closely and may revise their inflation expectations downward more sharply than others. Regions with larger concentrations of such households therefore experience larger declines in average expectations – the first-stage relationship. If expectations causally influence inflation, these same regions should also exhibit weaker inflation outcomes—the second-stage effect.

Now contrast this with a cost-push disturbance, such as an oil price spike. The same households—often with longer commutes and higher fuel usage—may experience especially salient price increases and thus revise their expectations upward more strongly than others. Regions with higher shares of these groups again show larger movements in average expectations, but the inflation response may be muted or even negative if these households cut back on other spending to offset higher energy costs. Together, these examples illustrate that the 2SLS estimator does not isolate a single structural disturbance. Rather, it captures a weighted average of the causal pass-through coefficients across the full constellation of shocks that historically moved expectations.

Formally, the Bartik interaction combines group-level expectation shocks with regional demographic composition. The resulting estimate can be expressed as a weighted average of the underlying shock-specific pass-throughs, with weights determined by how strongly each shock moves expectations and by each region's demographic exposure (Koo, Lee, Seo, and Takano (2024)). While the precise mixture of underlying shocks cannot be separately identified, the timing of large expectation shifts and the decomposition of identifying weights indicate that monetary and demand shocks likely account for a substantial share of the variation. Thus, the estimate can be interpreted as an empirical average pass-through that primarily reflects policy-driven expectation shocks, though it may also contain smaller contributions from supply and other sources.

Our empirical results provide credible evidence that inflation expectations do pass through to regional inflation rates. This finding is robust across multiple controls, datasets, and specifications. The causal estimates reveal a strong but less than one-for-one pass-through: a one percentage-point rise in expected inflation leads to roughly a 60-basis-point increase in realized regional inflation.

In an extension to assess long-term inflation expectations—those spanning a 5 to 10-year horizon—the data indicate a significant effect from short-term forecasts but negligible impact from long-term expectations. While the first stage with long-horizon expectations remains statistically strong–monetary policy and other macroeconomic shocks can move long-run expectations somewhat–the 2SLS coefficient is small and statistically indistinguishable from zero. There are three reasons behind this null result: First, only short-run expectations are likely to matter directly for price- and wage-setting decisions, so a weaker link at long horizons is theoretically expected (Werning (2022)). Second, long-run expectations in the United States have been remarkably stable since the early 1990s. If the identifying variation in the short-run specification primarily reflects monetary policy shocks, the null result is consistent with an anchored inflation regime in which such shocks affect near-term but not long-term beliefs. Third, the Michigan survey only began consistently measuring 5-10 year expectations during the inflation-targeting era, leaving relatively little cross-sectional or time-series variation for identification. Taken together, these points indicate that the long-horizon null arises from a combination of limited variation, anchoring, and the inherently short-run nature of the expectations channel.

While the empirical framework employed does not account for spatial spillovers in expectations shocks across regions, it nonetheless is informative. Specifically, cross-regional analyses identify the inflationary response to changes in expectations, with time-fixed effects capturing general equilibrium phenomena. The results should be interpreted as providing a lower bound on the influence of consumer expectations on overall inflation rates. Utilizing a common effects methodology, as in Pesaran (2006) and Harding and Lamarche (2011), we disentangle the contributions of regional and national inflation expectations to observed inflation. Notably, aggregate expectations exert an influence approximately 2.5 times stronger than their regional counterparts. Applying this multiplier to our favored pass-through estimate suggests an aggregate pass-through coefficient of 1.5, aligning closely with the theoretical findings in Werning (2022).

Our analysis further reveals that inflation expectations of married, younger consumers, with a minimum of a high school education, predominantly drive inflation, most notably in the non-durable goods sector. These findings not only enrich our understanding of the transmission channels for monetary policy but also contribute valuable moments for the calibration of heterogeneous-agent macroeconomic models.

The estimates presented provide a credible and more precise estimate of the impact of inflation expectations on inflation outcomes. Relative to aggregate Phillips Curve estimates, we exploit cross-sectional variation to estimate the causal impact of inflation expectations. Relative to previous cross-sectional estimates of the slope of the Phillips Curve (c.f. Hazell, Herreno, Nakamura, and Steinsson (2022)), we find similar coefficient estimates while also estimating a significant independent role played by short-term inflation expectations. The results are robust to a variety of controls and specifications.[2]

The remainder of this paper is organized as follows: Section 2 examines the variation in inflation expectations data crucial for our identification strategy. Section 3 articulates the empirical model and outlines the identification approach, while Sections 4–6 provide the empirical results.

1.1 Related literature

This paper contributes to the empirical literature on the causal impact of inflation expectations on realized inflation using cross-sectional micro data. An extensive literature estimates the expectations–inflation pass-through in aggregate New Keynesian Phillips Curve frameworks (see Mavroeidis, Plagborg-Moller, and Stock (2014) for a survey). These aggregate approaches face identification challenges: common shocks affect both expectations and realized inflation, and available instruments are often weak or uninformative about the heterogeneity across agents that may drive the relationship. Our paper instead uses a micro-level design exploiting cross-sectional variation in demographic composition to isolate exogenous shifts in inflation expectations.

A growing micro literature investigates how expectations affect economic behavior and prices

[2] The robustness exercises address several concerns about the University of Michigan Survey. These reservations include the survey's ability to capture genuine expectations, its rotating panel structure that might present panel conditioning (Binder (2015), Kim and Binder (2023)), rising non-response rates, and the inflation-history-dependent prompting of 'unreasonable' expectations by respondents.

using survey or experimental data. D'Acunto, Malmendier, Ospina, and Weber (2021) and Bachmann, Berg, and Sims (2015) use household survey data to study the link between expectations and portfolio or spending decisions, while Burke and Ozdagli (2022) and Crump, Eusepi, Tambalotti, and Topa (2021) relate expected inflation to intertemporal substitution and durable consumption during the zero lower bound period. At the firm level, randomized information treatments in Coibion and Gorodnichenko (2015) and Coibion and Gorodnichenko (2012) (and later Coibion, Gorodnichenko, and Ropele (2020)) show that firms causally raise prices when induced to expect higher inflation. Similarly, Tanaka, Bloom, David, and Koga (2020) document that firm managers' macroeconomic expectations predict business investment and hiring. Together, these papers provide micro evidence that short-run inflation expectations matter for economic outcomes, though they generally focus on partial equilibrium decisions rather than realized inflation itself.

Closest to our study are recent shift–share designs that instrument for local inflation expectations using national variation across identifiable subgroups. Binder, Kamdar, and Ryngaert (2024) construct an instrument based on regional exposure to political identification—Republican- versus Democrat-leaning populations—interacted with national partisan differences in inflation expectations during the COVID-19 period. They find that partisan unanchoring of expectations can account for roughly 2–3 percentage points of the observed pandemic-era inflation surge. Our design is conceptually analogous but uses demographic rather than political composition—sex,age, education, marital and parental status—to construct the shift–share instrument. The two approaches are complementary: Binder et al. capture short-run variation arising from political polarization in expectations, whereas our instrument exploits long-run demographic heterogeneity in expectation formation across many decades. Moreover, in our data the education component of the instrument is particularly influential, and recent electoral sorting by education suggests a natural connection between the two identification strategies. Our estimates, based on a much longer panel, yield a smaller average pass-through (roughly 0.6) consistent with typical historical periods, while Binder et al. capture an exceptional high-inflation episode with stronger expectation effects.

Importantly, we also present event-study evidence for the onset of the COVID period that reveals pass-through rates of roughly 1.2–1.5 within one to two months—comparable to the magnitude reported by Binder et al.—before reverting toward our long-run average estimate of about 0.6. This pattern indicates that the two studies are broadly consistent: expectation shocks appear to have a much larger immediate effect when inflation is rapidly changing, but a smaller average impact in more typical regimes. Thus, our multi-decade estimates capture the long-run average elasticity of inflation to expectations, while Binder et al. highlight a high-pass-through episode in which expectations became temporarily unanchored.

Our paper also relates to work on heterogeneous inflation experiences and household-level inflation rates. Hobijn and Lagakos (2005) and Kaplan and Schulhofer-Wohl (2017) use consumption micro data to construct household-specific inflation indices and document substantial heterogeneity across demographic groups. We extend this line of research by combining group-specific inflation expectations with group shares in regional populations, allowing us to identify the causal effect of expected inflation on realized inflation across heterogeneous agents. Like Kaplan and Schulhofer-Wohl (2017), who emphasize non-durables, we find that inflation expectations predominantly affect non-durable prices and are most salient among younger, married, and better-educated households.

Finally, our study differs from the literature on the formation of expectations. Numerous papers document deviations from rational expectations in survey data—see Evans and Gulamani (1984), Carroll (2003), Branch (2004), and Branch (2007)—and models of inattention or adaptive learning in Coibion and Gorodnichenko (2015, 2012). Our focus, in contrast, is on the macroeconomic consequence of those expectations: we provide evidence that differences in subjective inflation beliefs across groups causally translate into differences in realized regional inflation.

2 Data

Regional inflation is computed from the BLS consumer price index series for all urban consumers and all items across the four census regions (west, midwest, northeast, and south).[3] The annual inflation rate is the annualized log difference expressed in percentage points. BLS coverage began in 1966, but the survey data series has been available monthly since 1978. This period covers six recessions, including the disinflation in the early 1980s, the Great Recession in 2007-2009, the COVID recession, inflationary periods in the late 1970s, and the post-pandemic era. Regional inflation rates covary, have different variances and can be quite different during specific periods.

Inflation expectations come from the University of Michigan Survey of Consumers ("Michigan survey"). The survey has been conducted monthly since 1978 and consists of 600 or more respondents. Roughly 60% of survey respondents are new to the survey, while the remaining 40% are re-interviewed for a second time six months after their first interview. A respondent appears in the survey no more than two times. The sampling procedure uses the universe of telephone numbers to obtain a nationally representative sample. Since 2015 this has been the universe of cellular telephone numbers. The re-interviews are randomly drawn from the numbers of respondents in the survey six months prior. So, the survey sample is not a balanced panel. The presence of the prior respondents does raise some concerns. First, telephone response rates are higher for the re-interviews. Second, there is the possibility of a sample selection bias if the probability of a respondent appearing a second time in the survey is correlated with their inflation expectations; for instance, if more accurate forecasters are more likely to participate in the survey again. Similarly, during the intervening six months, a respondent could become more attentive to inflation news (Kim and Binder (2023)). The analysis below considers alternative specifications to address this potential concern.

The Michigan survey asks consumers a wide variety of questions. The two questions of interest here relate to the expected evolution of prices. In particular, after soliciting a respondent's views on whether prices will increase or decrease, they are then asked,

PX1 *By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?*

PX5 *By about what percent per year do you expect prices to go (up/down) on the average, during the next 5 to 10 years?*

Responses are in percentage points.[4] An ongoing concern with the Michigan survey's measure of inflation expectations is the presence of outliers. Throughout the sample, some survey responses are unreasonable; i.e., deflation rates of 20% or more during the late 1970s and inflation rates above 50% during the 2010s. As a first principle, the analysis does not remove outliers. Relatedly, at various points, the survey created or revised prompts to deliver to those who reported unrealistic inflation expectations. Unfortunately, the prompt for unrealistic expectations, e.g., 5%, is a function of the recent history of realized inflation. The analysis will probe the estimates for robustness to outliers, though the endogenous prompting is a deficiency in survey design.

[3]The Michigan Survey of Consumers does make some studies available with more granular location identifiers. However, it is impossible to construct a long panel, given the paucity of these studies.

[4]The bunching of histogram responses around a few integer values could raise concerns about digit preferencing, however Branch (2007) presents evidence against digit preferencing in the Michigan survey.

Figure 1: Group categorization. Note: not all groups shown.

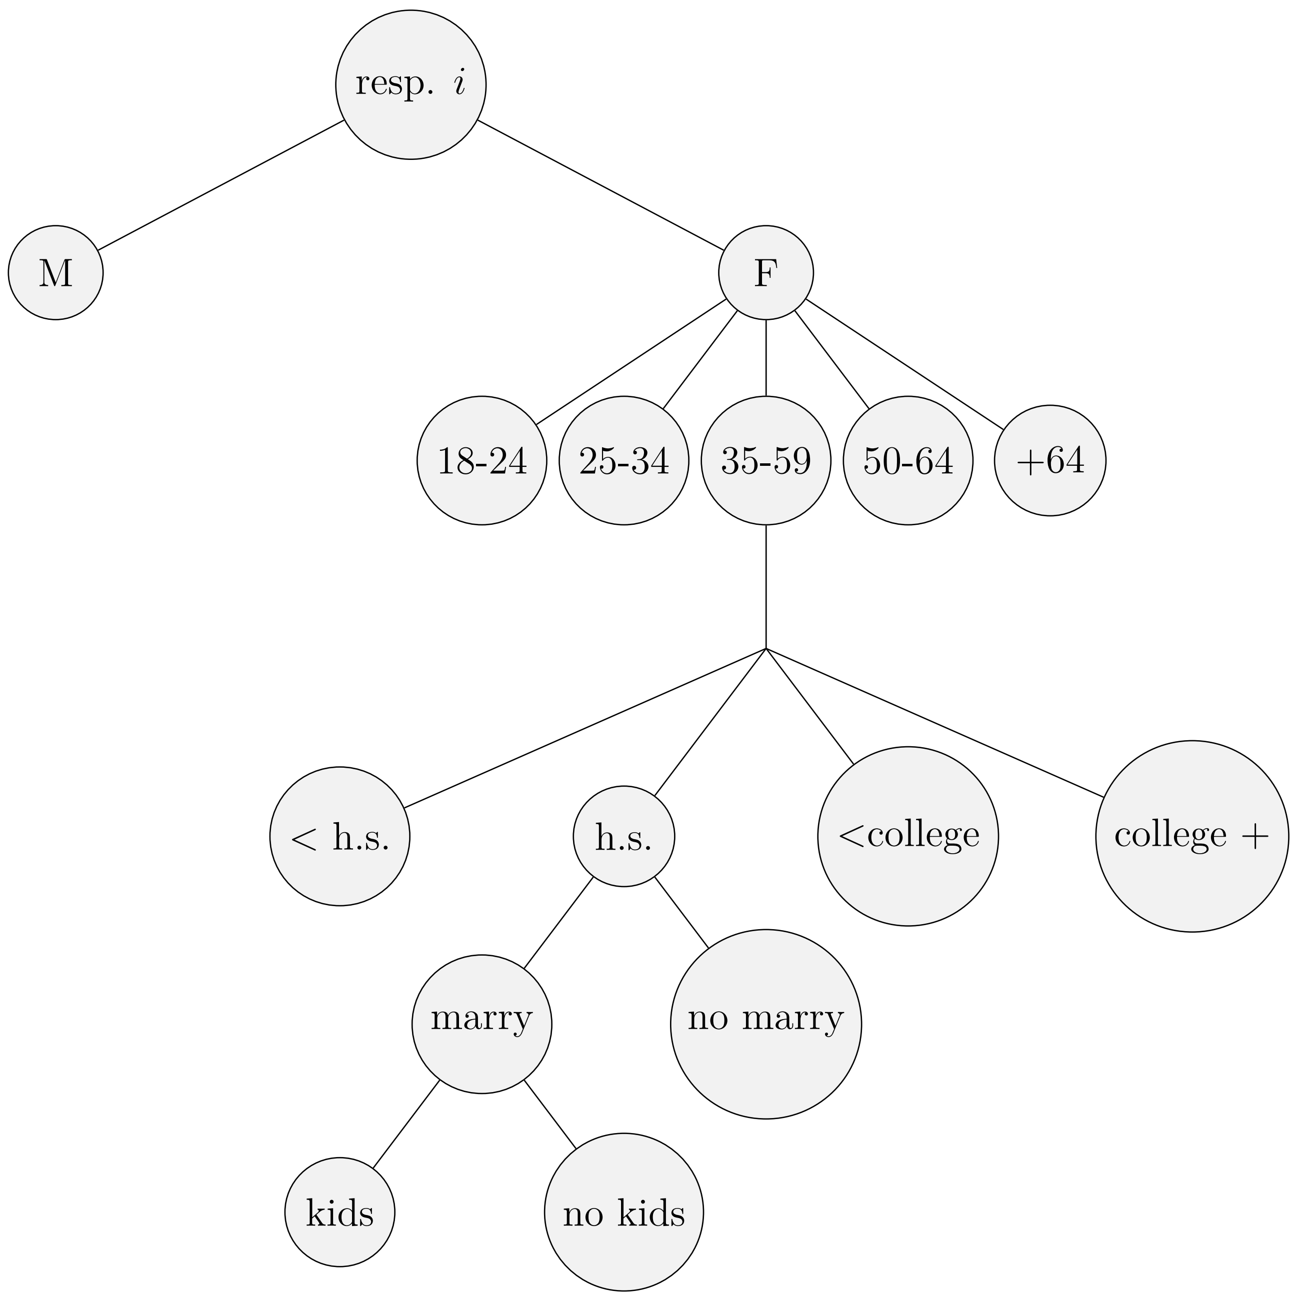

Motivation for the instrument follows from the heterogeneity in inflation expectations across demographic groups. While the diversity of Michigan survey expectations is well-known, the following summarizes the particular variation exploited by the current paper's research design. The remainder of this section documents the demographic and expectational variation.

The Michigan survey records a variety of demographic factors as well as the consumer's Census region. A panel of demographic groups follows the categorization of each of the roughly 273,000 survey responses, from 1978.1-2022.5, into one of 160 categories based on sex, age, education, marital and parental status. Those categories shown in Figure 1, produce a panel that consists of $T = 528$ months, $N = 4$ regions, and $G = 160$ groups. For each group, the vector of controls includes, among other variables, additional survey questions about consumer perceptions of the economy, unemployment, personal income, expected future income, current consumer financial status, and expected household financial status. For each period, each region, and each group, the average inflation expectation is treated as the unit of observation.[5]

Tables 1-2 summarize the demographic variation in inflation expectations that underpins our quasi-experimental design. Each table reports the mean and standard deviation of inflation expectations by demographic group, with Table 1 presenting results for single individuals and Table 2 for married respondents. For each group, data are shown separately for men and women, and are further broken down by education and age. The columns capture the "national shock" component in the shift-share IV design. Our identification strategy leverages the covariance between this national shock component, regional exposure shares, and inflation outcomes. Groups with higher standard deviations (such as married individuals) exhibit more variation in group-level expectations over time. The following sections provide further detail on the construction of exposure shares and discuss identification.

The tables show clear differences in both the level and variability of inflation expectations across groups. Women generally report higher and more variable inflation expectations than men. Younger and less educated respondents also display higher and more volatile expectations. Later analysis shows that identification is mainly driven by younger, married, and more educated groups, as indicated by Rotemberg weights. Figure 2 presents these cross-sectional differences using a heat map, where each square represents a demographic group's average expectation.

3 Empirical model

3.1 Motivating theory

Our goal is to estimate the causal effect of regional *inflation expectations* on regional inflation. We adopt a minimal multi-sector, multi-group New Keynesian benchmark with Calvo pricing (e.g., Cravino, Lan, and Levchenko, 2020) and subjective (possibly non-rational) expectations, and use only three ingredients: (i) groups differ in their expenditure weights across sectors and in their subjective inflation expectations; (ii) regions differ in their group composition; (iii) Calvo pricing implies a locally linear (around steady state) relation between inflation and expected inflation, with intercepts absorbed by fixed effects. These ingredients imply a *linear* mapping from group-specific expectations to regional inflation and motivate our reduced form and instruments (shift-share interactions between region-specific group shares and common (leave-one-out) innovations in group expectations). We remain agnostic about which demographic attribute generates exposure.

[5] The panel dimensions fit a "small N, big T" setting that raises potential finite sample bias concerns that the estimation will address.

Figure 2: Mean inflation expectations by demographics. Each square is particular group's sample average inflation expectations. There is evident expectational heterogeneity across groups. For example, the highest average expectations are in groups of younger women with less than a high school degree. While the lowest average expectation are older, more educated consumers.

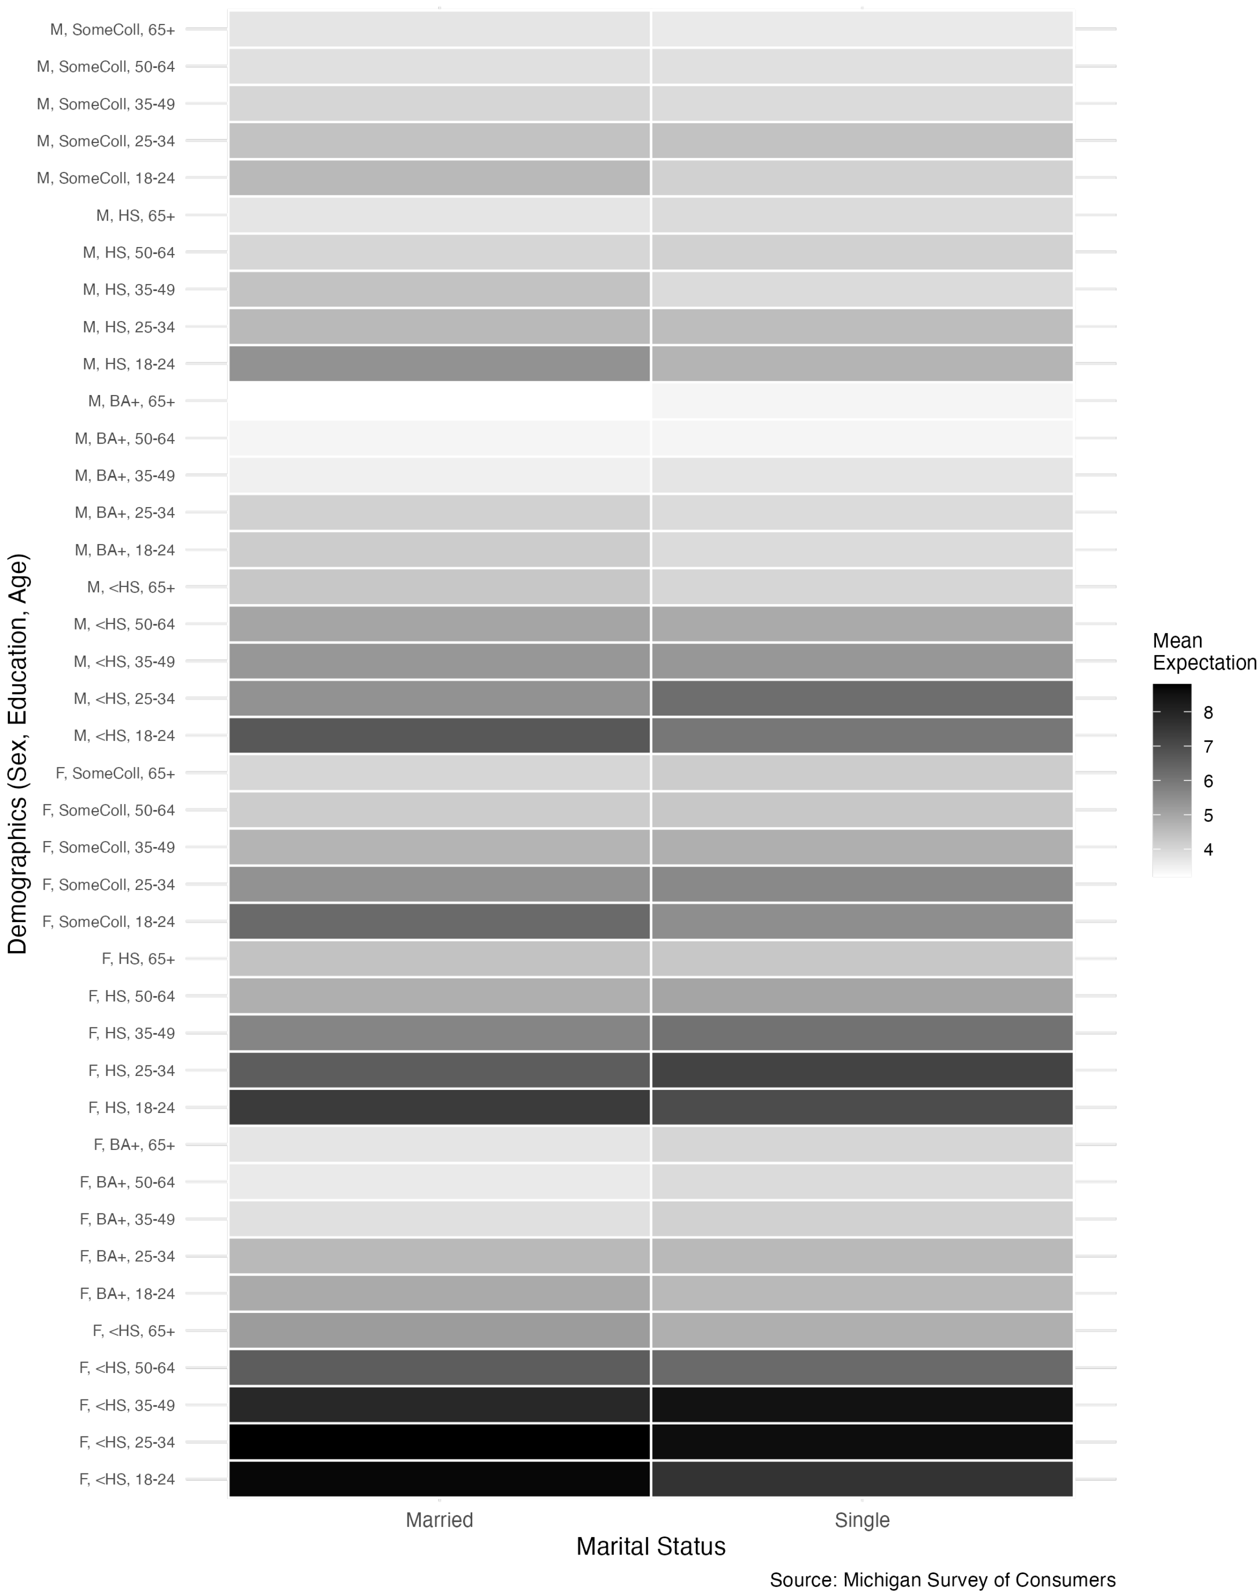

Table 1: Inflation Expectations with Cross-Regional Dispersion: Single

		Male				Female			
Education	Age	Mean	SD	σ_{reg}	N	Mean	SD	σ_{reg}	N
Less than HS	18-24	5.9	3.4	2.82	139	6.8	5.1	3.58	100
Less than HS	25-34	5.6	4.3	2.42	38	7.5	4.6	5.56	87
Less than HS	35-49	4.2	5.2	4.48	76	7.5	4.4	4.33	121
Less than HS	50-64	4.1	4.0	2.04	110	5.6	3.7	3.94	234
Less than HS	65+	3.7	3.5	2.97	212	4.2	2.7	2.93	403
HS	18-24	4.0	2.7	2.63	432	5.5	3.4	3.26	309
HS	25-34	4.1	3.1	2.34	353	5.7	3.4	3.27	303
HS	35-49	3.7	2.9	2.02	415	5.2	2.9	2.90	417
HS	50-64	4.0	2.6	2.24	378	4.5	2.9	2.56	477
HS	65+	3.9	3.0	2.12	313	4.1	2.2	2.64	512
SomeColl	18-24	3.7	2.5	2.16	490	4.8	2.6	2.73	443
SomeColl	25-34	3.9	3.0	2.06	393	5.0	3.0	2.89	371
SomeColl	35-49	3.9	2.8	2.00	431	4.5	2.9	2.43	477
SomeColl	50-64	3.8	2.5	2.01	305	4.1	2.8	2.53	455
SomeColl	65+	3.6	2.5	1.99	236	4.1	2.5	2.43	482
BA+	18-24	3.6	2.8	1.98	283	4.3	3.6	2.82	251
BA+	25-34	3.4	2.3	1.66	471	4.0	2.9	2.10	460
BA+	35-49	3.7	2.5	1.75	512	4.1	2.6	2.14	511
BA+	50-64	3.4	2.5	1.73	408	3.9	2.2	2.06	461
BA+	65+	3.2	2.0	1.66	290	3.8	2.9	2.15	460

Note:
Each statistic is computed from monthly group means. Mean and SD compute full-sample moments. σ_{reg} is the average within-month cross-regional standard deviation of those means. N is the number of months with coverage across Census regions.

3.1.1 Environment and preferences

Time is discrete. A region r is a closed economy populated by a continuum of households and a continuum of monopolistically competitive firms producing differentiated goods in each of S consumption sectors, indexed by $s = 1, \ldots, S$. Households are partitioned into G observable groups $g = 1, \ldots, G$ (e.g., by sex, age, marital status, children, and/or education). Group g consumes a Dixit–Stiglitz basket of sectoral goods with elasticity of substitution $\theta > 1$, time-invariant budget shares $\{\omega_s^g\}_{s=1}^S$ over sectors, and a group-specific price index[6]:

$$P_t(g) = \left(\sum_{s=1}^{S} \omega_s^g \, P_{s,t}^{1-\theta} \right)^{\frac{1}{1-\theta}}, \qquad \sum_{s=1}^{S} \omega_s^g = 1. \tag{1}$$

[6]For motivation, the model abstracts from gradual shifts in group preferences or consumption patterns. The empirical model allows for varying group shares.

Table 2: Inflation Expectations with Cross-Regional Dispersion: Married

		Male				Female			
Education	Age	Mean	SD	σ_{reg}	N	Mean	SD	σ_{reg}	N
Less than HS	18-24	8.3	5.3	3.06	26	8.1	4.7	3.36	57
Less than HS	25-34	4.4	4.7	3.50	138	7.5	4.0	4.03	145
Less than HS	35-49	4.6	3.3	2.76	246	6.4	4.0	3.60	210
Less than HS	50-64	4.1	2.9	2.43	283	5.3	3.2	3.01	206
Less than HS	65+	3.8	3.1	2.65	315	4.5	3.7	3.61	194
HS	18-24	4.9	3.5	3.11	163	5.8	3.9	3.13	236
HS	25-34	3.8	2.5	2.04	436	5.2	2.7	2.37	365
HS	35-49	3.9	2.1	1.83	523	5.0	2.3	2.25	468
HS	50-64	3.7	2.0	1.82	527	4.4	2.1	2.22	518
HS	65+	3.6	2.2	2.01	509	4.3	2.7	2.37	454
SomeColl	18-24	4.7	3.9	3.60	109	5.7	3.9	3.23	170
SomeColl	25-34	3.8	2.5	1.78	458	4.6	2.7	2.52	464
SomeColl	35-49	3.8	2.2	1.82	530	4.5	2.3	2.25	528
SomeColl	50-64	3.7	2.2	1.74	511	3.9	2.4	2.13	510
SomeColl	65+	3.7	2.5	1.94	425	3.8	2.6	2.19	378
BA+	18-24	5.1	3.6	1.84	59	4.8	3.7	2.98	108
BA+	25-34	3.4	2.2	1.43	510	4.0	2.3	1.90	526
BA+	35-49	3.5	2.0	1.22	531	3.9	2.0	1.67	531
BA+	50-64	3.5	2.1	1.32	531	3.7	2.2	1.76	521
BA+	65+	3.2	1.9	1.39	481	3.5	2.8	1.88	364

Note:
Each statistic is computed from monthly group means. Mean and SD compute full-sample moments. σ_{reg} is the average within-month cross-regional standard deviation of those means. N is the number of months with coverage across Census regions.

Let $\mu_{r,g}$ denote the (time-invariant) population expenditure share of group g in region r, $\sum_g \mu_{r,g} = 1$. The region-level price index is

$$P_{r,t} = \left(\sum_{s=1}^{S} \omega_{r,s} \, P_{s,t}^{1-\theta} \right)^{\frac{1}{1-\theta}}, \qquad \omega_{r,s} \equiv \sum_{g=1}^{G} \mu_{r,g} \, \omega_s^g. \tag{2}$$

3.1.2 Price setting and subjective expectations

In sector s, a firm faces a Calvo friction: with probability $1 - \lambda_s$ it cannot reset its price.[7] Marginal cost (in logs) is w_t and sectoral inflation is $\pi_{s,t} \equiv \Delta p_{s,t}$, where $p_{s,t} \equiv \log P_{s,t}$. We assume *anticipated-utility* and *steady-state learning* expectations: agents may learn over time, but for current decisions they treat future inflation in sector s as constant at its current expected rate $\pi^e_{s,t}$; see Woodford (2013); Evans and Honkapohja (2001).

[7]Heterogeneous stickiness across sectors follows Cravino, Lan, and Levchenko (2020).

Under these assumptions, the sectoral New Keynesian Phillips relation collapses to a static linear form:

$$\pi_{s,t} \;=\; \kappa_s\,\pi^e_{s,t} \;+\; x_{s,t}, \qquad \kappa_s \equiv \frac{1}{1-\beta\lambda_s}, \tag{3}$$

where $x_{s,t}$ is the discounted present value of marginal cost factors.[8]

3.1.3 From sectoral to group and regional inflation

Group-g expected inflation is the sector-weighted average

$$\pi^e_{g,t} \;=\; \sum_{s=1}^{S} \omega^g_s\,\pi^e_{s,t}. \tag{4}$$

Stacking across groups and sectors, define $\mathbf{\Pi}^e_t \equiv (\pi^e_{g,t})_{g=1}^G \in \mathbb{R}^G$ and $\mathbf{\Pi}^e_{s,t} \equiv (\pi^e_{s,t})_{s=1}^S \in \mathbb{R}^S$. Let $W \in \mathbb{R}^{G\times S}$ collect the budget shares with entries $W_{gs} = \omega^g_s$. Then (4) can be written succinctly as

$$\mathbf{\Pi}^e_t \;=\; W\,\mathbf{\Pi}^e_{s,t}. \tag{5}$$

Under the identifying assumption that the S sector share vectors are linearly independent (i.e., W has full column rank with $S \le G$), let $B \in \mathbb{R}^{S\times G}$ denote any right inverse of W satisfying $WB = I_S$. Then sectoral inflation expectations are linearly recovered from group expectations: $\mathbf{\Pi}^e_{s,t} = B\mathbf{\Pi}^e_t$.

Region-r inflation aggregates sectoral inflation using $\omega_{r,s}$:

$$\pi_{r,t} \;=\; \sum_{s=1}^{S}\omega_{r,s}\,\pi_{s,t} \;=\; \sum_{s=1}^{S}\omega_{r,s}\left(\kappa_s\pi^e_{s,t} + x_{s,t}\right) \;=\; \underbrace{\boldsymbol{\omega}_r'\,\mathrm{diag}(\boldsymbol{\kappa})\,\mathbf{\Pi}^e_{s,t}}_{\text{expectation component}} \;+\; x_{r,t}, \tag{6}$$

where $\boldsymbol{\omega}_r \equiv (\omega_{r,1},\ldots,\omega_{r,S})'$, $\boldsymbol{\kappa} \equiv (\kappa_1,\ldots,\kappa_S)'$, and $x_{r,t} \equiv \sum_s \omega_{r,s}x_{s,t}$. Substituting $\mathbf{\Pi}^e_{s,t} = B\mathbf{\Pi}^e_t$ yields the key linear mapping

$$\pi^R_{r,t} \;=\; \Gamma_r\,\mathbf{\Pi}^e_t \;+\; x^R_{r,t}, \qquad \Gamma_r \;\equiv\; \boldsymbol{\omega}_r'\,\mathrm{diag}(\boldsymbol{\kappa})\,B \;\in\; \mathbb{R}^{1\times G}. \tag{7}$$

Equation (7) shows that regional inflation is a linear function of the vector of group expectations. Any partition of the G groups into, say, two disjoint sets (e.g., education bins versus other demographics) induces an additive decomposition of the expectation component in (7).[9]

3.1.4 Markups and demographic composition

The shift-share instrument interacts regional group shares $\mu_{r,g} \in \Gamma_r$ with national group expectations. A natural concern is that group-specific markups might vary across regions due to differences in $\mu_{r,g}$. This could occur if firms account for different willingness to pay across groups. In such a case, regional variations in mark-ups, as well as pass-through, might become endogenous to the group distribution.

[8] This is a reduced form version of the forward-looking Calvo equation under steady-state learning: see also Werning (2022). We do not require rational expectations and only linear dependence on $\pi^e_{s,t}$. Instead of steady-state learning one can alternatively allow for, say, more dynamic learning models. Later sections allow for dependence on long-run expectations. The Appendix explores lagged-expectation regressions to assess robustness.

[9] This is a reduced-form, local linear approximation. Non-linear generalizations or more dynamic expectations formation would complicate the mapping but are unlikely to overturn the core cross-sectional logic.

To frame the issue, consider the following extension to Anderson, Rebelo, and Wong (2020). There exist $s = 1, \ldots, S$ sectors, $g = 1, \ldots, G$ groups, and a continuum $[0, n]$ of retail firms within each sector. While each firm sells the identical good, they offer a range of amenities (e.g., ease of parking, self-checkout), resulting in imperfect substitution among the firms' sectoral products. Let's denote $C_{s,g,t}$ as group g's consumption of sector s's goods, formulated as:

$$C_{s,g,t} = v_g^{\gamma} \left[\int_0^n x_{s,i,g,t}^{1/v_g} di \right]^{v_g} \tag{8}$$

The elasticity of substitution across firms becomes:

$$-\frac{v_g}{1 - v_g} \tag{9}$$

A consumer of type g, in sector s, will face prices calculated as:

$$P_{s,i,g,t} = v_g^{\gamma/v_g} P_t C_{g,t}^{v_g - 1/v_g} x_{s,i,g,t}^{1 - v_g/v_g} \tag{10}$$

Consequently, the optimal price will have a constant mark-up v_g above the marginal cost. With group-specific elasticities of substitution across goods, it is conceivable that mark-ups and expectation pass-through could correlate with the distribution of groups across regions.

This concern, however, doesn't necessarily violate the exclusion restriction. The crucial point is that the pure cross-group differences in levels of markups (time invariant group shares) impact price-levels but not the inflation rate. Further, an acyclical, or mostly time-invariant mark-up, would not threaten the exclusion restriction, as shown by Anderson, Rebelo, and Wong (2020). Nevertheless, we investigate this possibility in the data and find that mark-ups do not significantly correlate with group shares.

3.2 Empirical model

The object of interest is inflation expectations' impact on regional inflation rates. An instrument is constructed that exploits that demographic groups may have different expected inflation, e.g. because of distinct preferences for the goods that make up a consumption basket, varying attentiveness or economic literacy, etc. Interacting regional group shares with aggregate group inflation expectations instruments for the endogenous expectations. The empirical strategy is a differential exposure design: we identify the impact of expectations by measuring how a region's exposure to aggregate shocks leads to a differentiated inflation response. Each region has differential exposure to the shocks because of different population distributions. The identification strategy is valid so long as the demographic shares satisfy a relevance and an exogeneity condition.

The coefficient of interest is β in the equation

$$\pi_{r,t} = \delta_r + \beta \pi_{r,t}^e + \kappa U_{r,t} + \gamma' x_{r,t} + \mu_t + \varepsilon_{r,t} \tag{11}$$

where $\pi_{r,t}$ is the inflation rate in region r at time t, $\pi_{r,t}^e$ is the expectation of 12-month ahead regional inflation, $U_{r,t}$ is the regional unemployment rate, $x_{r,t}$ is a vector of controls (that includes lags of inflation, financial and non-financial income, regional business conditions), $\varepsilon_{r,t}$ is the structural disturbance, δ_r, μ_t are region and time fixed effects, respectively. The endogeneity concern is that estimating (11) via ordinary least squares will produce biased estimates of β because $\pi_{r,t}^e$ is endogenous. In particular, endogeneity will arise if expectations respond, after controlling for exogenous covariates

$x_{r,t}$, to the factors driving $\varepsilon_{r,t}$. While estimates of β are not sensitive to bias from the endogeneity of unemployment, we follow Hazell, Herreno, Nakamura, and Steinsson (2022) and instrument for $U_{r,t}$ using its 12-month lag.

A shift-share instrument ("Bartik instrument") addresses the key endogeneity. Building on the aggregation relationships developed in Section 3.1.3, a region's expected inflation can be written as the population-weighted average of its groups' expectations:

$$\pi^e_{r,t} = \sum_g \mu_{r,g} \pi^e_{r,g,t}$$

Furthermore, if each region group's inflation expectation decomposes into aggregate and idiosyncratic components, then

$$\pi^e_{r,g,t} = \pi^e_{g,t} + u_{r,g,t}$$

The empirical strategy uses exogenous variation in $\mu_{r,g}$ to generate differential exposure to the group-specific aggregate component $\pi^e_{g,t}$. The instrument is

$$z_{r,t} = \sum_g \mu_{r,g} \pi^e_{g,t}$$

Identification requires the standard IV assumptions:

$$\text{Relevance:} \quad \text{Cov}\left(z_{r,t},\, \pi^e_{r,t} \mid x_{r,t},\, \delta_r,\, \mu_t\right) \neq 0 \tag{12}$$

$$\text{Exclusion:} \quad \text{Cov}\left(z_{r,t},\, \varepsilon_{r,t} \mid x_{r,t},\, \delta_r,\, \mu_t\right) = 0 \tag{13}$$

One advantage of the shift-share instrument is that we can remain agnostic about the underlying mechanism that satisfies the relevance condition. Either exposure heterogeneity, of the type emphasized here, or differences in the common component of group expectations, can generate a non-zero covariance in equation (12). Hence, relevance does not hinge on whether cross-group differences reflect consumption basket preferences or, say, attentiveness and economic literacy.

The key identifying assumption is that, conditional on observables and region/time fixed effects, the common group-specific inflation expectations when interacted with predetermined regional group exposure shares, shifts $\pi^e_{r,t}$ but has no independent effect on regional inflation $\pi_{r,t}$ except through expectations. Operationally, (13) is satisfied when (i.) we include observable controls and time fixed effects that capture economy-wide macroeconomic factors, and (ii.) the instrument $z_{r,t}$ is constructed using a leave-one-out procedure that removes mechanical correlation by omitting the region's own inflation expectations. Intuitively, cross-regional demographic composition differences expose regions differentially to the group-specific aggregate component of expectations, and those shocks are as-good-as-random once aggregate macroeconomic factors are controlled for.

The main potential concern which threatens identification is whether group shares predict regional inflation rates through channels other than those posited here. A reasonable conjecture is that the distribution of groups is endogenous to a region's price level, a cost-of-living measure. However, for identification, it is sufficient that the group shares in a region are exogenous to the change in prices, i.e., inflation (Goldsmith-Pinkham, Sorkin, and Swift (2020)). This exogeneity assumption is plausible. However, to make the case convincing, the empirical analysis measures shares using either the beginning of sample population distribution or time-varying survey shares. In the former, those shares are not predictive of the exogenous covariates.

The empirical approach is flexible along several key dimensions. Rather than requiring rational expectations, the theory simply motivates a locally linear relationship between regional inflation and expected inflation. The identification strategy is also robust to demographic partitioning—while different partitions (education vs. other demographics) yield alternative group share measures, they target the same coefficient of interest. Finally, exact recovery of sectoral expectations from group expectations is unnecessary; the reduced form (11) and IV moments (12)–(13) provide sufficient identification.

Using the beginning of period shares, and probing the predictive power of those shares, helps allay concerns over whether regional demographics are endogenous to the inflation rate. One plausible story could be that younger and more educated groups tend to live in regions with more dynamic or concentrated industries that experience increasing rates of price changes. In that case, those beginning-of-period shares would help predict the other variables that also predict inflation.

Standard macroeconomic identification strategies center around measuring the economic response to a particular structural shock. The 2sls estimator can be written compactly as $\hat{\beta} = {}^{\mathrm{Cov}_t(\pi_{r,t},\, z_{r,t})}/_{\mathrm{Cov}_t\left(\pi^e_{r,t},\, z_{r,t}\right)}$.[10] If group expectations are impacted by a composite of underlying shocks (e.g. demand, mark-ups, etc.), then the same mixture affects both the numerator and denominator. Under a standard shock decomposition, $\hat{\beta}$ equals a Rotemberg-weighted average of shock-specific pass-through rates, with larger weights on shocks that move regional expectations strongly and to which the shift-share instrument is more exposed. The estimate, therefore, provides the average causal response of inflation to inflation expectations across the realized mixture of shocks that move expectations.

Besides using predetermined shares, and probing the predictive power of those shares, the empirical strategy also presents results from a variety of diagnostic tests. A stacked-IV approach is used with two shift-share instruments that differ by the group partitions (e.g. education vs. non-education demographics). A Hansen J test does not reject the overidentifying restrictions, indicating that the different exposure instruments identify a common β. This particular finding aligns with the mechanism-agnostic nature of the relevance condition. A placebo analysis is also conducted to ensure that group demographics are a key source of identifying variation.

3.3 The shift-share instrument

We exploit cross-regional differences in demographic composition to obtain differential exposure to common innovations in group inflation expectations. Let $\mu_{r,g}$ denote the (predetermined) population share of group g in region r measured in a pre-sample cross-section (January 1978 CPS), and let $\pi^e_{-r,g,t}$ be the (leave–one–out) innovation in group g's expected inflation at time t constructed from the Michigan microdata. The benchmark instrument is

$$z_{r,t} \;=\; \sum_g \mu_{r,g}\, \pi^e_{-r,g,t}, \tag{14}$$

a standard shift-share (Bartik) combination of exposure $\mu_{r,g}$ and shifts $\pi^e_{-r,g,t}$. As a check on the use of CPS shares, we also report estimates based on time-varying shares $\xi_{r,g,t}$ built from the Michigan sample:

$$z^{\mathrm{MI}}_{r,t} \;=\; \sum_g \xi_{r,g,t}\, \pi^e_{-r,g,t}.$$

The identifying assumption is that, conditional on region/time fixed effects and observables, $z_{r,t}$ shifts regional expected inflation but has no independent effect on regional inflation except through

[10]These issues are studied in a local projection-IV framework by Koo, Lee, Seo, and Takano (2024). See the Appendix for a formal discussion of these points.

expectations.[11]

How shares are constructed. The CPS shares $\mu_{r,g}$ are computed once at the start of the sample (January 1978), by Census region and the 160 demographic groups used in the instrument. These shares are plausibly exogenous with respect to subsequent inflation shocks and, empirically, do not predict the included controls. The Michigan-based shares $\xi_{r,g,t}$ are calculated at the same group level by region and month directly from the surveys. They are used as an alternative exposure measure.

Coverage and representativeness. Two potential issues arise. First, does the Michigan survey provide enough observations to form reliable group-time aggregate expectation measures $\pi^e_{g,t}$? Second, are regional share measures representative? Table 3 summarizes the dispersion of shares across broad groups using (i) Michigan survey shares averaged over time and (ii) CPS January 1978 shares across regions. With roughly 600 monthly survey participants, and 160 groups, not every group is observed every month, creating a potential coverage concern. Table 4 reports coverage for the 160 groups and for the most identification-relevant subsets (ranked by Rotemberg weight α): the average number of respondents per group–month and the fraction of months with at least 2 or 5 respondents. Coverage is adequate for the top groups that carry most instrument weight, and we show below that results are stable when we (i) use coarser groupings and (ii) restrict to high-coverage groups (Section 5).

Finally, the Michigan survey's national sampling frame implies that sampling error in $\pi^e_{g,t}$ is not systematically related to region–time outcomes.[12] Averaging at the group–time level and the leave–one–out construction further mitigate noise and mechanical correlation. In Section 4.3 we verify that alternative partitions of groups (education vs. other demographics) identify a common effect and, in Section 5, we show robustness to coarser groupings and to restricting to high-coverage groups.

4 Impact of one-year ahead inflation expectations

4.1 Preliminary results

Figure 3 previews and visualizes the empirical estimates to follow. The left panels report results using the Michigan survey shift-share, while the right is for the CPS shift-share instrument. The top panels (3a-3b) plot the first stage, and the bottom panels (3c-3d) visualize the reduced-form regression of regional inflation on the Bartik instrument. In each panel, the solid line is the linear regression equation. The shift-share instrument correlates strongly with the Michigan survey expectations. The Michigan survey-based shift-share instrument correlates more closely with expectations than the instrument calculated with CPS78 shares. In both cases, the first stage relationship is as expected: regions exposed to groups with high national expectations have high regional expectations.

[11]The Michigan survey monthly shares introduce two potential concerns. First, the re-interviewing of previous respondents and the potential endogeneity of the response rate. Second, it is plausible that an individual may move through groups over time and in response to a region's economy and inflation rate. Since the panel consists of groups rather than individuals, evolving group composition is not a particularly great concern except for those repeating respondents. Subsequent sections alleviate these concerns by demonstrating the robustness of the coefficient estimates to a specification with first-time respondents.

[12]Our main estimates form group-time means from the raw Michigan responses. Recomputing the group-time shocks with Survey of Consumers sampling weights yields instruments and 2SLS estimates that are robust. We favor the unweighted construction because the weights do not target region-by-month representativeness and can add noise in small group-month cells.

Table 3: Detailed Demographic Shares by Region

Panel A: Michigan Survey Shares

Region	Male (%)	Female (%)	<HS (%)	HS (%)	College (%)	<35 (%)	35-49 (%)	50+ (%)	Married (%)	Single (%)
Northeast	49.3	50.7	8.0	22.8	39.9	28.5	28.0	43.6	59.3	40.7
Midwest	47.7	52.3	10.2	32.4	32.6	28.1	27.2	44.7	61.6	38.4
South	47.7	52.3	10.0	29.7	40.1	28.5	28.3	43.2	59.1	40.9
West	47.9	52.1	11.9	28.0	35.4	27.2	27.2	45.6	60.8	39.2
Std. Dev.	0.781	0.781	1.602	4.020	3.653	0.624	0.570	1.127	1.232	1.232

Panel B: CPS Population Shares

Region	Male (%)	Female (%)	<HS (%)	HS (%)	College (%)	<35 (%)	35-49 (%)	50+ (%)	Married (%)	Single (%)
Northeast	47.7	52.3	77.7	2.0	3.9	56.3	16.1	27.5	44.3	55.7
Midwest	48.2	51.8	78.3	2.4	3.4	58.3	15.8	26.0	46.5	53.5
South	47.9	52.1	79.1	2.3	3.3	58.3	16.5	25.3	46.1	53.9
West	49.4	50.6	72.5	3.2	4.3	62.6	16.2	21.3	45.1	54.9
Std. Dev.	0.744	0.744	3.002	0.529	0.451	2.634	0.291	2.663	1.020	1.020

Note: Panel A shows demographic shares computed from the Michigan Survey of Consumers. Panel B shows population shares from the 1978 Current Population Survey. Education categories: <HS (less than high school), HS (high school graduate), College (bachelor's degree or higher). Age groups: <35 (under 35), 35-49 (ages 35-49), 50+ (age 50 and older). Marriage status: Married vs Single. Standard deviations measure variation across the four Census regions (Northeast, Midwest, South, West).

Table 4: Coverage Statistics by Group Importance

Group Category	Groups	Avg Monthly Obs	% Months $\geq$2 Obs	% Months $\geq$5 Obs
All Groups (160)	160	3.8	51	23
Top 40 by α	40	5.1	74	40
Top 20 by α	20	6.0	81	51

Note:
Coverage statistics show data availability across demographic groups. 'Avg Monthly Obs' shows mean respondents per month per group. '

Panels (3c-3d) plot the results from a reduced-form panel regression of the inflation rate on the predicted inflation expectations. Regardless of the instrument, there is a strong positive relationship between predicted inflation expectations and inflation.

The rest of the analysis probes the interpretation of the results in Figure 3.

4.2 2sls estimates

Tables 5-6 report two-stage least squares estimates of the impact of consumer inflation expectations on regional inflation rates. Panel A details the coefficient estimate with the CPS 1978.1 population shares. Panel B encompasses the Michigan survey measure of group shares. In both instances, constructing the instrument uses a regional leave-one-out approach. All specifications include a set of control variables (two lags of inflation, regional unemployment, and survey measures for current and anticipated financial/business conditions) and region and time-fixed effects—the latter controls for aggregate business cycle factors common to all regions.[13] Following Hazell, Herreno, Nakamura, and Steinsson (2022), we also instrument $U_{r,t}$ with its 12-month lag, though, the estimated pass-through is not sensitive to instrumenting. Each numbered column corresponds to a model with or without fixed effects.

First stage. Table 5 provides first stage estimates. The Bartik instrument is relevant and has significant predictive power for inflation expectations. The Durbin-Wu-Hausman test statistic for endogeneity is 8.074, rejecting the consistency of OLS at a 1% significance level. The first-stage F-statistic is 16.4 (CPS shares) and 24.4 (Michigan shares), comfortably above common thresholds. While the identifying variation in the shares is plausibly exogenous to regional inflation rates, in the case of CPS78, it is possible to probe the exogeneity assumption by examining whether those 1978 population shares are predictive of the regression covariates. In each case, there is no significant correlation between the covariates and the CPS group shares, and the regression coefficients are economically close to zero.

Second stage. Table 6 finds a positive, statistically significant pass-through from expectations to inflation. Using CPS shares, $\hat{\beta} = 0.609$ (1% level), i.e. pass-through a bit above one-half. Using the Michigan survey shares, the estimated coefficient is 0.339, so a 1% increase in the average expected inflation rate in a region would lead to a 34 basis point increase in that region's inflation rate. The

[13] As a robustness check, we also include gas price expectations as a control variable. Including gas price expectations leads to a sightly higher estimate of pass-through. However, the estimate is less precise since gas price expectations are available less frequently than inflation expectations.

Figure 3: Inflation expectations and inflation: reduced-form estimates. Top panels show first-stage relationships between regional inflation expectations and shift-share instruments constructed using Michigan survey shares (left) and CPS 1978.1 demographic shares (right). Bottom panels show reduced-form relationships between regional inflation rates and the shift-share instruments. Solid lines show OLS fits from unconditional bivariate regressions. These plots illustrate the unconditional correlations, while Tables 5–6 report the full 2SLS estimates with controls for lagged inflation, unemployment, survey covariates, and fixed effects.

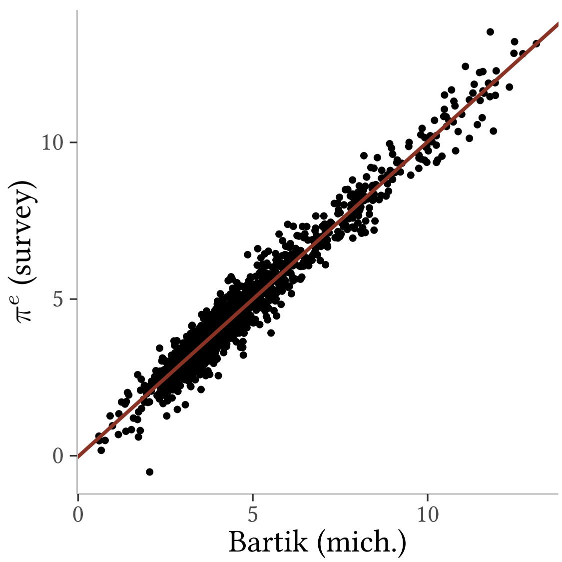

(a) First stage: shares calculated from Michigan survey.

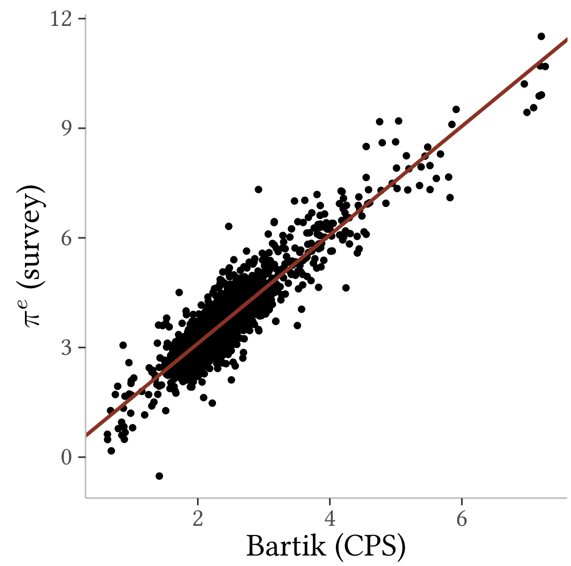

(b) First stage: shares calculated from 1978.1 CPS.

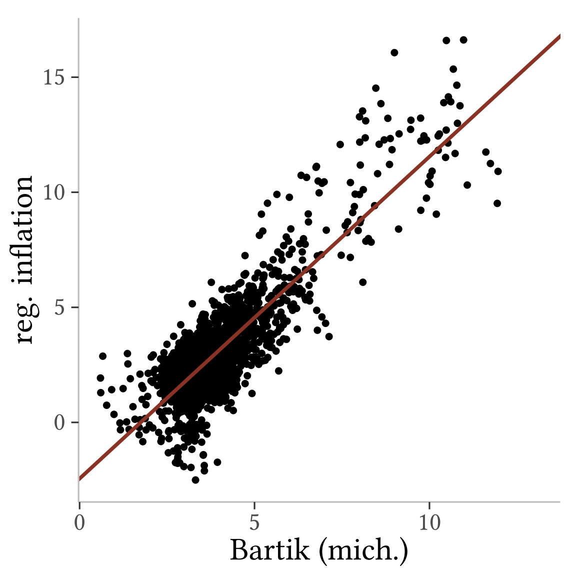

(c) OLS regression: Michigan shares.

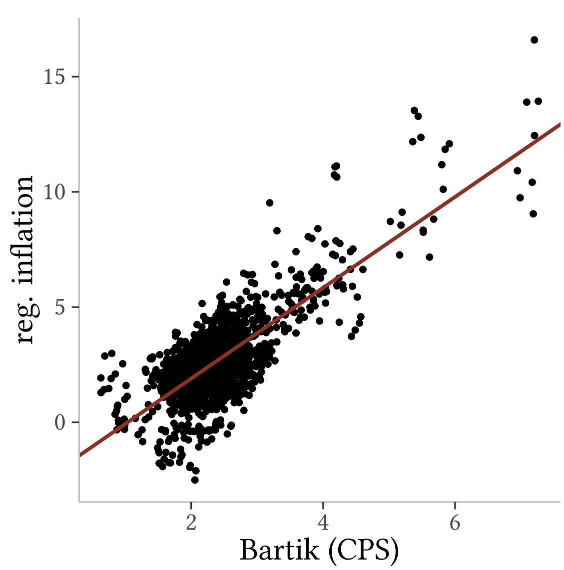

(d) OLS regression: CPS78 shares.

Table 5: 2SLS: first stage

Panel A: CPS shares	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	1.09 (0.05)	1.11 (0.07)	0.23 (0.03)	0.32 (0.05)
u^r	0.52 (0.06)	0.46 (0.04)	0.74 (0.09)	0.53 (0.04)
R^2	0.685	0.715	0.766	0.781
KP-F stat	514.2	574.4	9.3	16.4

Panel B: Michigan shares	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	0.75 (0.04)	0.74 (0.04)	0.22 (0.01)	0.23 (0.01)
u^r	0.52 (0.06)	0.46 (0.04)	0.74 (0.09)	0.53 (0.04)
R^2	0.694	0.713	0.770	0.784
KP-F stat	551.0	566.5	21.7	24.4

Note: Driscoll–Kraay standard errors in parentheses. KP-F is the Kleibergen–Paap rank-Wald weak-instrument statistic for the first-stage regression of one-year inflation expectations p^e on the Bartik shift–share instrument. Instruments: Bartik shift–share for the one-year horizon together with the 12-month lag of the unemployment rate u^r_{t-12}. All specifications include two lags of regional inflation, the full set of survey-covariate controls, and the fixed effects indicated at the top of each column.

Table 6: 2SLS estimates

Panel A: CPS shares				
	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	0.279 (0.038)***	0.275 (0.038)***	0.748 (0.399)*	0.609 (0.247)**
u^r	0.011 (0.032)	0.019 (0.041)	0.026 (0.045)	-0.134 (0.083)
R^2	0.928	0.929	0.912	0.932
Within R^2	0.928	0.928	0.190	0.340
N	1,531	1,531	1,531	1,531

Panel B: Michigan shares				
	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	0.249 (0.035)***	0.254 (0.036)***	0.391 (0.155)**	0.373 (0.138)***
u^r	0.016 (0.033)	0.023 (0.042)	0.010 (0.028)	-0.137 (0.064)**
R^2	0.929	0.929	0.948	0.952
Within R^2	0.929	0.929	0.526	0.529
N	1,531	1,531	1,531	1,531

Note: Driscoll–Kraay standard errors in parentheses. Instruments: Bartik shift–share for the one-year horizon shown together with the 12-month lag of unemployment u^r_{t-12}. All regressions include two lags of regional inflation, survey-covariate controls, and the fixed effects indicated above each column.

estimated coefficient is significant at the 1% level. OLS estimates with the same controls produce an estimate around 0.06, which is biased downwards.

While the empirical model does not account for spatial spillovers, a back-of-the-envelope calculation suggests that the aggregate impact of expectations is stronger than the regional effect. Much like the literature on regional fiscal multipliers, after accounting for cross-regional spillovers– an increase in inflation expectations in one region also impacts the demand for goods produced in a different region– the aggregate impact of expectations is stronger than the regional effect. After controlling for lags, the correlation between median aggregate inflation expectations and US aggregate inflation rates is 0.18, compared to a correlation of 0.06 for regional measures. A similar magnitude in the 2sls estimates would imply a pass-through of approximately $1.0 - 1.6$.

The coefficient estimates are in line with the theoretical analysis in Werning (2022), which studies pass-through in a variety of conventional pricing models with time-dependent price rigidities while not making *a priori* assumptions about the formation of subjective inflation expectations. While the pass-through from expectations to inflation can take any positive value, Werning (2022) shows that for the Calvo and Taylor models of price stickiness, that pass-through should be in the range $[1/2, 1]$, and possibly above one in a more general framework. While the regional estimates in Table 6 are on the lower end, or slightly below, of this range, after accounting for regional spillovers, the estimates are in line with the theoretical predictions.

The panel data is susceptible to a potential finite sample bias. Of particular concern are the data's "small N/big T" dimensions. Table 7 applies a split-sample jackknife bias correction (Fernandez-Val and Weidner (2018)). After correcting for finite sample bias in both the cross-section and time dimensions, the estimated effect of inflation expectations decreases slightly. Table 7 provides the preferred

Table 7: Bias Correction

	survey shares	CPS78 shares
π^e	0.3680	0.5491
se	0.1345	0.2006

Note:
Applies the split-sample jacknife bias correction to the 2sls coefficient estimates. The column "survey shares" computes shares from Michigan survey,"CPS78 shares" uses the CPS 1978.1 shares.

estimate: a 1% increase in inflation expectations leads to a 55 basis point increase in (regional) inflation.

Dynamic response. Another way to measure the impact of expectations is by estimating an impulse response function. The empirical model is not a vector autoregression, but a local projections approach can estimate the impulse response function.[14]. The impulse response function comes from running 2sls regressions of the form

$$E_t \pi_{r,t+h} = \delta_{r,h} + \beta_h \pi^e_{r,t} + \gamma'_h x_{r,h} + \mu_{t,h} + \varepsilon_{r,h,t}$$

for horizons $h = 1, ..., H$. The impulse response function then is given by $(\beta_h)_{1 \le h \le H}$. The estimates are provided in Figure 4.

On impact, inflation expectations have a significant positive effect on inflation. There is a cyclical dampening process that is mean-reverting within 12 months. However, the confidence bands are wide in future quarters and cannot rule out any lingering impacts. Thus, the estimates suggest a moderate contemporaneous response to inflation from a shock to inflation expectations. The lack of a strongly persistent expectation effect could reflect the specific US inflation history. It would be interesting to extend the analysis to countries with persistent or volatile inflation.

4.3 Identification and interpretation

The 2SLS estimate captures how regional inflation responds to shifts in inflation expectations, but this aggregate coefficient masks important heterogeneity. Different demographic groups may have systematically different pass-through rates from expectations to actual inflation—some groups' expectations may be more predictive of realized inflation than others. At the same time, our identification strategy exploits the fact that regions differ in their exposure to national group-level expectation shocks, and some group shocks provide more informative variation for identification than others. It is important to distinguish these two concepts: group heterogeneity in true causal effects, and group heterogeneity in identifying power.

To understand both dimensions, we decompose our 2SLS estimate using Rotemberg weights following Goldsmith-Pinkham, Sorkin, and Swift (2020). These weights diagnose (1) which demographic groups' expectation shocks provide the most identifying variation for the instrument, and (2) how the overall 2SLS coefficient can be expressed as a weighted average of group-specific pass-through

[14] See Jorda (2005).

Figure 4: Impulse responses

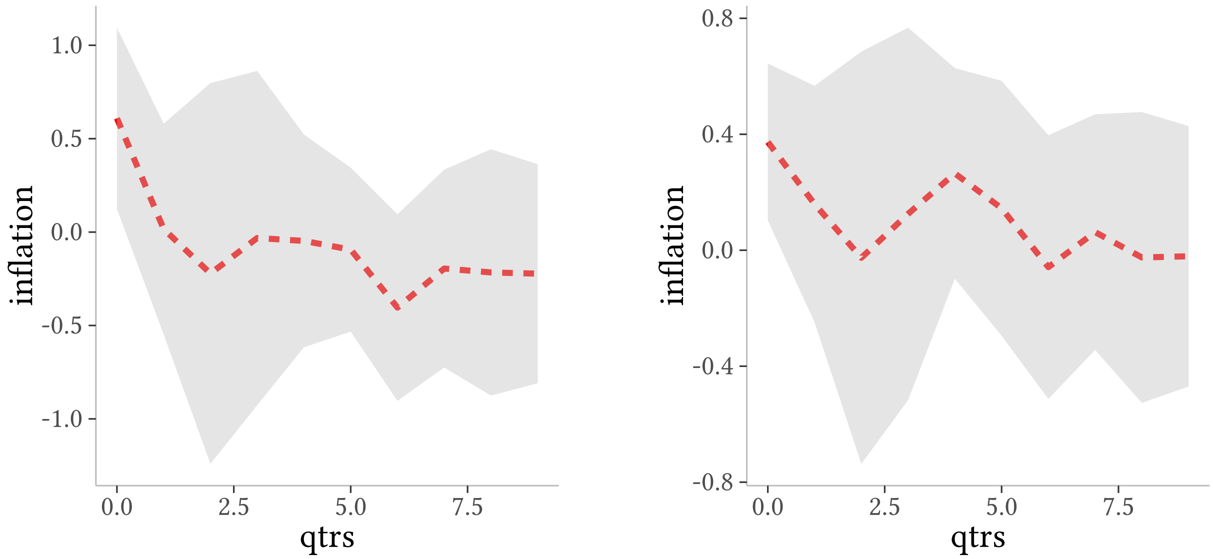

(a) cps shares

(b) michigan shares

estimates. The former speaks to sources of econometric leverage, while the latter shows which heterogeneous group effects are being aggregated in our estimate.

The Rotemberg decomposition. With 160 demographic groups defined by age, education, marital status, and parental status, the shift-share instrument combines variation from 160 just-identified instruments. Goldsmith-Pinkham, Sorkin, and Swift (2020) exploit that the 2SLS estimate can be decomposed as:

$$\hat{\beta} = \sum_g \alpha_g \beta_g$$

where β_g denotes the just-identified estimate based solely on group g's variation. The Rotemberg weight α_g reflects the sensitivity of the overall estimate to shocks originating from group g. These weights sum to one but can be either positive or negative. In cases where most weights are positive–as is true here–$\hat{\beta}$ may be interpreted as an approximate weighted average of group-specific effects. However, the presence of negative weights indicates that certain groups' variation counteracts the main identifying variation, complicating this interpretation.[15]

The weight α_g is high when group g provides strong identifying variation: large expectation shocks combined with substantial cross-regional differences in population shares $\mu_{r,g}$ that together generate a powerful first-stage relationship. For this reason, ranking groups by their Rotemberg weights assesses which groups drive identification. Formally, α_g also represents the sensitivity of $\hat{\beta}$ to potential misspecification in group g's instrument, making these the groups where the exclusion restriction matters most.

Which groups drive identification? Tables 8-9 reveal a striking pattern: identification comes disproportionately from younger, educated households. Using CPS shares (Table 8), the top 10 groups are aged 25-49, with 8 of 10 having at least some college, 5 having college degrees (or more), and 7

[15]Substantial negative weights mean the estimator differentially weights positive and negative effects rather than simply averaging them; see Goldsmith-Pinkham, Sorkin, and Swift (2020) for more details.

Table 8: Top 10 weighted groups: CPS shares

group id	α_g	β_g
Ml,35-49,s.c.,M,NK	0.0241	1.4313
Ml,25-34,s.c.,M,K	0.0215	1.3205
Ml,18-24,c+,M,NK	0.0207	1.6029
Ml,35-49,c+,M,K	0.0203	1.3793
Ml,25-34,c+,NM,NK	0.0194	1.6651
Ml,35-49,c+,NM,NK	0.0176	1.3630
Fl,25-34,h.s.,M,NK	0.0168	1.7446
Fl,35-49,s.c.,M,NK	0.0164	1.2675
Fl,25-34,h.s.,NM,NK	0.0161	1.4837
Fl,25-34,c+,M,K	0.0157	1.8796

Note:
Group labels ordered: sex, age, educ., marital, children. Top-10 weighted groups according to the "Rotemberg" weights as in Goldsmith-Pinkham, et al (2020).

being married. These groups alone account for a substantial portion of the identifying variation. The Michigan shares (Table 9) paint an even younger picture—most high-weight groups are aged 18-24, though they are more evenly split between high school and college education.

Two features stand out. First, all heavily-weighted groups have positive weights and group-specific estimates β_g significantly above 1, suggesting their expectations have strong effects on realized inflation. This could reflect that younger, educated households are more economically active, update expectations more frequently, or have better information about economic conditions. Second, the difference between CPS and Michigan instruments is economically meaningful: CPS shares exploit fixed 1978 demographic differences across regions, while Michigan shares capture evolving demographic patterns. That both identify similar groups, suggests our results are robust to the source of demographic variation, though the specific weighted average differs.

Table 10 aggregates group weights for broader demographic groupings. Regardless of instrument choice, groups with ages above 50 receive minimal weight, while prime-age dominate. The CPS instrument weights college-educated groups most heavily, consistent with an information-based or literacy story, while the Michigan instrument spreads weight more evenly across education levels, potentially capturing broader heterogeneity.

Graphical evidence of heterogeneous effects. Figure 5 visualizes the heterogeneity in estimated effects. Each point represents a demographic group, with circle size proportional to its Rotemberg weight. The horizontal axis shows the group's first-stage F-statistic (instrument strength), while the vertical axis shows the just-identified β_g. Three patterns emerge. First, the scarcity of negative weights (triangles) confirms we are identifying a weighted average of positive effects rather than

Table 9: Top 10 weighted groups: Michigan shares

group id	α_g	β_g
Ml,18-24,h.s.,M,NK	0.0330	1.5574
Ml,25-34,h.s.,M,NK	0.0177	1.7646
Ml,18-24,c+,M,NK	0.0174	1.5920
Fl,18-24,s.h.s.,M,NK	0.0160	1.4804
Ml,18-24,s.c.,M,NK	0.0155	1.4664
Ml,18-24,h.s.,M,K	0.0149	1.6477
Fl,18-24,h.s.,M,NK	0.0146	1.6288
Fl,18-24,s.c.,M,K	0.0146	1.6878
Ml,35-49,s.c.,M,NK	0.0146	1.4254
Fl,25-34,h.s.,M,NK	0.0141	1.8965

Note:
Group labels ordered: sex, age, educ., marital, children. Top-10 weighted groups according to the "Rotemberg" weights as in Goldsmith-Pinkham, et al (2020).

Table 10: Rotemberg weights for broad demographic groups.

shares	Women	Ages 18-24	Ages 25-34	Ages 35-49	Ages 50-64	Ages 65+	w/h.s.	college +
CPS	0.450	0.245	0.340	0.255	0.118	0.041	0.247	0.329
Michigan	0.413	0.333	0.296	0.148	0.113	0.064	0.279	0.261

Note:
Table reports fraction of overall Rotemberg weights attributable to a broader categorization of demographic groups. This table gives indication of which broad groups are key to identification.

offsetting positive and negative relationships. Second, the most heavily-weighted groups (largest circles) cluster relatively close to the overall estimate (dashed line), suggesting the aggregate effect is not driven by outliers. Third, the positive correlation between first-stage strength and effect size hints that groups whose expectations we can measure most precisely also have the strongest pass-through to inflation.

Figure 5: Estimation weights across groups: Michigan shares. Scatter plot of each group's first stage F-statistic against its just-identified effect. The size of the circle represents the Rotemberg weight. The largest circles are the top-10 groups. Triangles denote negative weights.

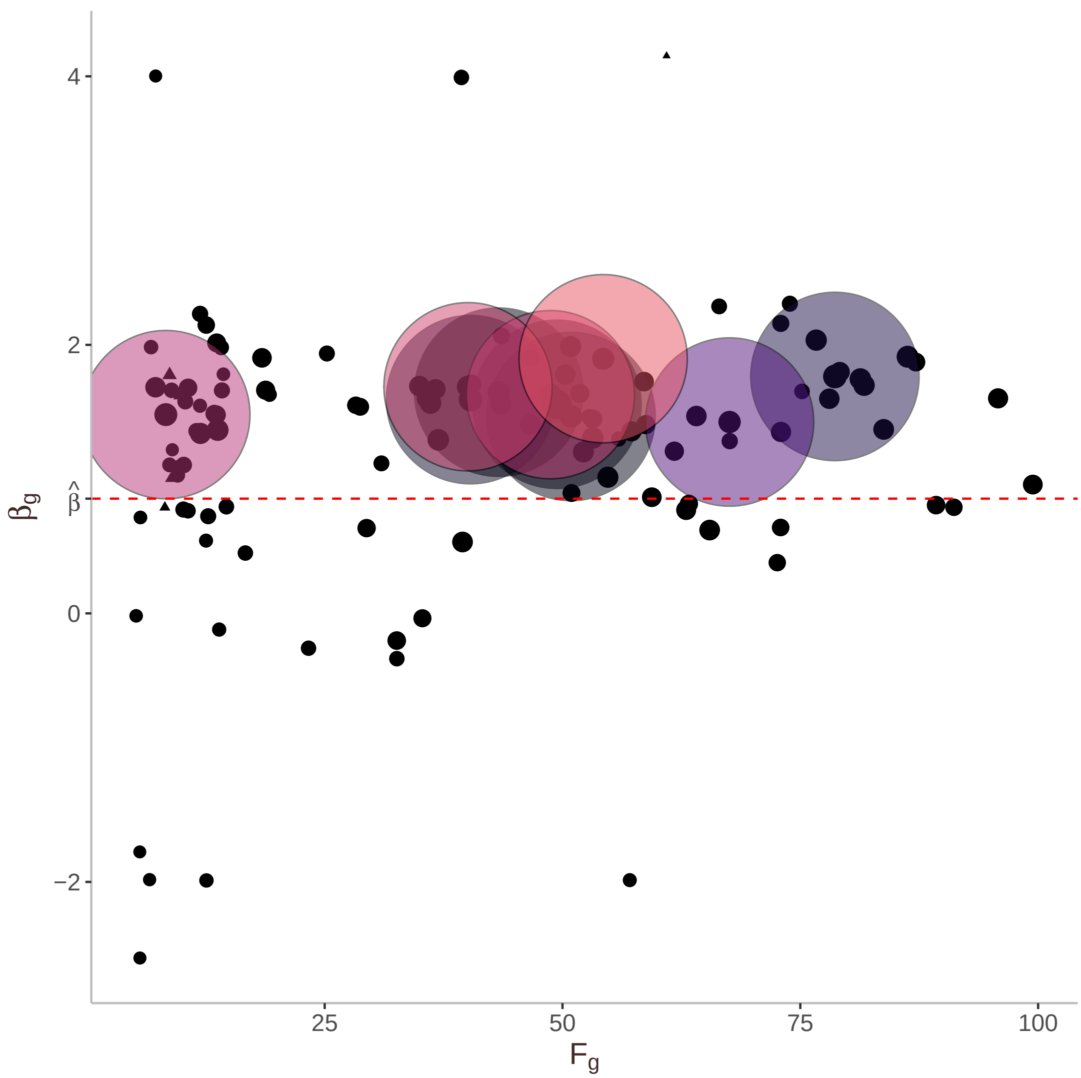

When does identification occur? Figure 6 reveals the time periods when identifying variation occurs. Rather than being spread evenly across the sample, identification concentrates in specific periods with greater inflation volatility: the end of the 1970s Great Inflation, the Volcker disinflation, the Great Recession (2007-2009), and the post-pandemic inflation surge.

The time frame concentration has two implications. First, it reinforces confidence in the empirical strategy—identification comes precisely from periods when expectational heterogeneity across groups was likely highest and when the expectations-inflation pass-through was most salient. Second,

Figure 6: Estimation weights across time.

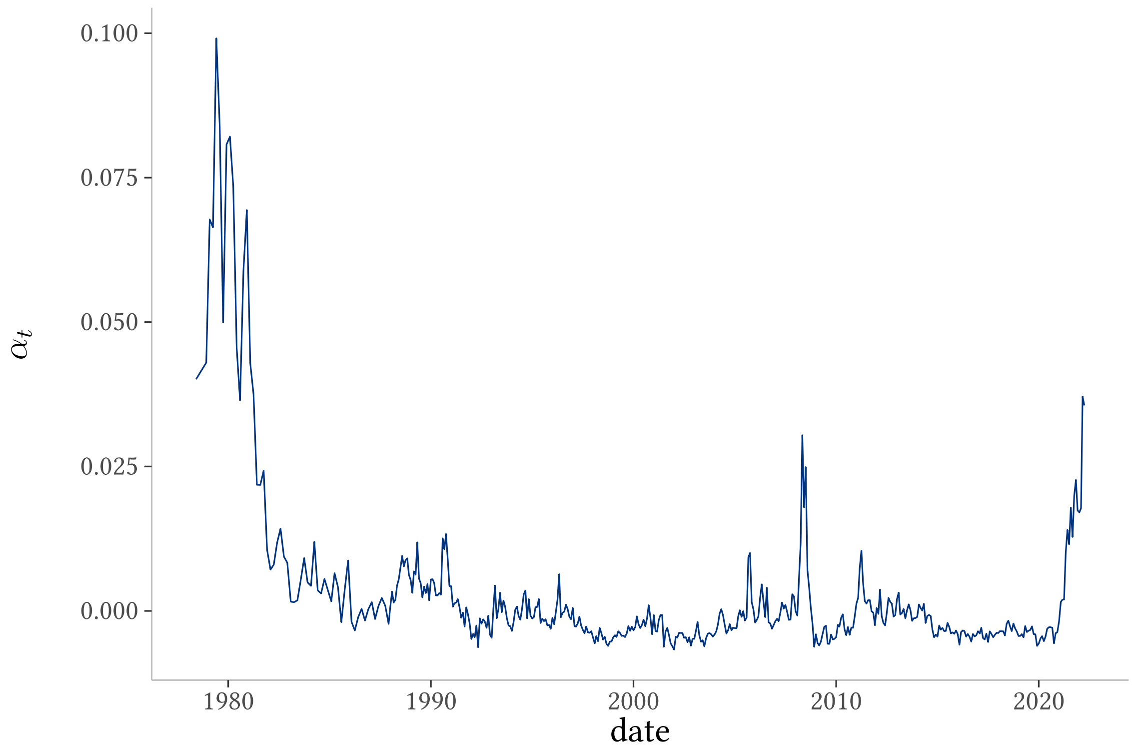

it implies the estimate captures how inflation responds to an "average" expectation shock during volatile times rather than normal times. This aligns with state-dependent models where expectations matter more when uncertainty is high Branch (2004); Brock and Hommes (1997), and suggests the pass-through estimate may be an upper bound for calmer periods. A related observation is made by Binder, Kamdar, and Ryngaert (2024), who use a shift–share instrument based on political affiliation and find a larger pass-through during politically polarized periods. Their design, like ours, identifies the causal effect of expectations in states of high dispersion, suggesting that both estimates likely represent upper bounds relative to calmer periods.

Distinguishing between information and consumption heterogeneity The empirical results naturally lead to the question of why certain groups appear so prominently in the Rotemberg decomposition. Do educated groups contribute more to identification because they are better informed (an information/attention channel), or because different demographic groups consume different baskets and thus experience different inflation rates (a consumption channel)? While the Michigan survey data do not allow us to perfectly separate these mechanisms, we can probe the issue further by constructing separate instruments and testing whether they identify the same underlying parameter.

To this end, Table 11 presents an overidentification test using stacked instruments. The stacked IV approach constructs multiple shift-share instruments using different subsets of groups, then combines them in a single 2SLS regression. This creates an overidentified system where it is possible to test whether different sources of variation identify the same parameter. If they do not—that is, if the Hansen J test rejects—it suggests the instruments capture fundamentally different economic mechanisms rather than different sources of variation for the same underlying relationship.

Here it is possible to decompose the 160 demographic groups into education-based variation and other demographic variation. Specifically, we construct two instruments:

1. A "Pure Education" component using variation across education categories.
2. A "Demographics Component" capturing variation from age, marital status, and family structure, orthogonalized from the education component through a Gram-Schmidt decomposition.

Table 11: Decomposing Bartik Instruments: Education vs. Demographics

Instrument Components	Price Expectations	(Std. Error)
Demographics Component	1.175	(0.827)
Pure Education	$0.875^{\dagger}$	(0.461)
Both Components	0.942^{*}	(0.430)
Orthogonalized Components	0.942^{*}	(0.430)
Diagnostic Statistics:		
Kleibergen-Paap F-Statistic	6.35	
Hansen J-Statistic [p-value]	0.252 [0.615]	
First-Stage Variance Decomposition:		
Demographics Component	26%	
Education Component (Orthogonal)	74%	

Notes: Dependent variable is regional inflation expectations. All specifications include region and year-quarter fixed effects plus controls for lagged regional inflation, unemployment rate, and consumer sentiment measures. Standard errors are Driscoll–Kraay. Instruments use leave-one-out Bartik construction with demographic×education shares (160 total groups). Demographics component captures variation from age, marital status, and family structure within education groups. Education component represents variation orthogonal to demographics after Gram-Schmidt orthogonalization. $^{\dagger}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table 11 presents results from this stacked IV exercise. When used separately, the demographics component yields a coefficient of 1.175 while the pure education component yields 0.875. The education component—which isolates variation purely from education differences—accounts for 74% of the first-stage variance despite representing only one dimension of demographic heterogeneity. This suggests education-based differences in expectations are particularly powerful predictors of regional inflation expectations.

However, the key test comes from estimating with both components. The Hansen J-statistic (0.252 with p-value 0.615) fails to reject the null hypothesis that both instruments identify the same parameter. This passing of the overidentification test is economically important: it suggests that while the magnitude of effects differs across demographic dimensions, both education and other demographic characteristics generate valid identifying variation for the same underlying expectations-inflation relationship.

The Kleibergen-Paap F-statistic of 6.35 indicates the instruments remain relatively weak when orthogonalized. This suggests some caution when interpreting. However, the pattern is clear: education matters quantitatively more for identification (providing nearly three-quarters of identifying variation), but both channels operate. This points to a world where there is an information channel proxied by educated groups *and* different demographic groups experience genuinely different inflation pass-through rates, possibly due to consumption basket differences.

Importantly, these results support the empirical strategy while highlighting an important subtlety. The failure to reject the overidentification test suggests the Bartik instrument identifies a well-defined parameter despite the underlying heterogeneity. However, the substantial difference in coefficients between education and demographic components (0.875 vs. 1.175) reveals that the aggregate estimate averages over meaningfully different group-level effects. This further reinforces the Rotemberg weight findings by helping to understand which groups drive identification.

Taken together, the Rotemberg decomposition and stacked IV estimates show that the aggregate pass-through estimate mainly reflects expectational feedback from younger, educated groups during periods of inflation volatility. These are the groups whose expectations matter most for inflation–those likely to be most informed and economically active. The heterogeneous group-specific effects, with β_g ranging from near zero to above 3 for different groups, suggests that policies aimed at stabilizing inflation expectations may have differential impacts across the population. These findings have important implications for the transmission of monetary policy and our understanding of how inflation expectations are formed.

4.4 Event studies: Volcker and COVID

The baseline 2SLS estimates identify the average causal response of regional inflation to expectations across all shocks that moved beliefs over the sample period. This represents a weighted average of potentially heterogeneous effects from different shock types–monetary policy/aggregate demand, mark up/aggregate supply, etc.–each of which might generate different pass-through rates. To explore this compositional issue, this section examines via an event-study style analysis the mechanism during two key periods for US inflation dynamics: the Volcker disinflation (1979-1984) and the COVID inflation surge (2020-2022).

What follows is descriptive, exploring how the reduced-form patterns evolve during well-understood historical episodes to help interpret the mixture of shocks our main estimates capture.

Event-study approach. The goal is to construct episode-specific measures of regional exposure and trace their relationship with inflation over the event. For each episode, we first calculate the average change in each group's expectations over the episode window: $\overline{\Delta\pi^e_g} = \bar{\pi}^e_{g,\text{post}} - \bar{\pi}^e_{g,\text{pre}}$, where pre and post periods are specific to each episode.[16] The instrument is the region's exposure to these episode-specific expectation changes:

$$\bar{z}_r^{\text{episode}} = \sum_g \mu_{r,g} \times \overline{\Delta\pi^e_g}$$

where $\mu_{r,g}$ are the baseline group shares (1978.1 CPS shares).

The event-study regression is:

$$\pi_{r,t} = \sum_k \beta_k \times \mathbf{1}(t=k) \times \bar{z}_r^{\text{episode}} + \gamma' x_{r,t} + \delta_r + \mu_t + \epsilon_{r,t} \tag{15}$$

where $\mathbf{1}(t=k)$ is an indicator function for months relative to the event and the pass-through right before the event is normalized $\beta_{-1} = 0$. The coefficient β_k captures how the change in inflation between high and low exposure regions evolves month-by-month through the event. The coefficients are estimated with 2sls with region and time fixed effects.

This approach aims to answer a more *descriptive* question than the main estimation: given the pattern of group expectation changes during an event, how do regions with different group exposures to these changes experience different inflation paths? The answer helps us understand

[16] For Volcker, the pre-period is January-June 1979 and the post-period is November 1979-December 1984. For COVID, the pre-period is September 2018-February 2020 and the post-period is July 2020-December 2022. These windows capture the characteristic expectation shifts associated with each episode.

the heterogeneous effects from different shocks that underlie this paper's estimates of the impact expectations have on inflation.

For each episode, the figures below report four panels that decompose the results:

(a) *Coefficient path* $\hat{\beta}_k$: How the relationship between group exposure and inflation evolves monthly relative to the event.

(b) *Inflation effect*: $\widehat{\Delta\pi}_k = \hat{\beta}_k \times \bar{z}_r^*$, where $\bar{z}_r^* = 1/4 \sum_r \bar{z}_r^{\text{episode}}$. The predicted inflation differential for the average regional exposure.

(c) *First stage path* $\hat{\beta}_k^{(\pi^e)}$: How demographic exposure predict regional expectations over the event.

(d) *Expectation effect*: $\widehat{\Delta\pi^e}_k = \hat{\beta}_k^{(\pi^e)} \times \bar{z}_r^*$. The predicted expectation change for the average regional exposure.

These panels focus in on the pass-through from expectations to inflation, as identified in the full sample, during specific episodes. The caution in interpreting these results is because, particularly for the Volcker episode, these are stretches of time rather than a specific, easily observable event.

The Volcker disinflation (1979-1984).

The Volcker episode captures a complex period where monetary policy dominates but operates alongside other shocks: see Figure 7. The period includes the second oil crisis (1979-1980), Volcker's appointment in August 1979, the implementation of monetary tightening, and fiscal changes. This multiplicity of forces creates more complex dynamics than a clean monetary experiment would produce. Nevertheless, the episode is still useful for understanding and interpreting the main estimation results: one might expect younger households are more sensitive or attentive to monetary policy (e.g. home purchases or mortgage refinancing) and so inflation expectations decline and pass-through to inflation.

Panel (c) reveals these competing forces. The first-stage coefficient initially shows regions with greater exposure to the shock – that is, those regions with more educated, young, and married groups – do not immediately show differential expectation declines but, rather, higher expectations. This likely reflects offsetting effects. While Volcker's appointment signaled eventual disinflation, the concurrent oil price shock pushed expectations upward, and the most aggressive monetary tightening was still months away.

Along these lines, panel (a) shows that, during the first three months, regions more exposed to these demographic groups had weakly (and stat. insignificant) positive pass-throughs. With more exposed regions featuring higher inflation expectations, but the pass-through effects weak, inflation during these first months of the event are stable.

Crucially, by month 4 onward, the monetary channel dominates and the expected pattern emerges strongly: high-exposure regions show both systematically lower expectations (panel d) and lower inflation (panel b). The pass-through coefficient stabilizes at positive values, confirming that even during disinflation, the structural relationship holds–it is regions with relatively higher expectations (or smaller expectation declines) that experience relatively higher inflation. This pattern indicates that the shift-share instrument primarily identifies monetary policy shocks during the Volcker period.

The COVID inflation surge (2020-2022). The COVID episode presents a cleaner shock-sudden supply chain disruptions and unprecedented fiscal transfers created immediate, visible inflation pressures. In this case, younger households with kids may be sensitive or attentive to fiscal transfers,

Figure 7: Event study analysis: Volcker disinflation (1979-1984)

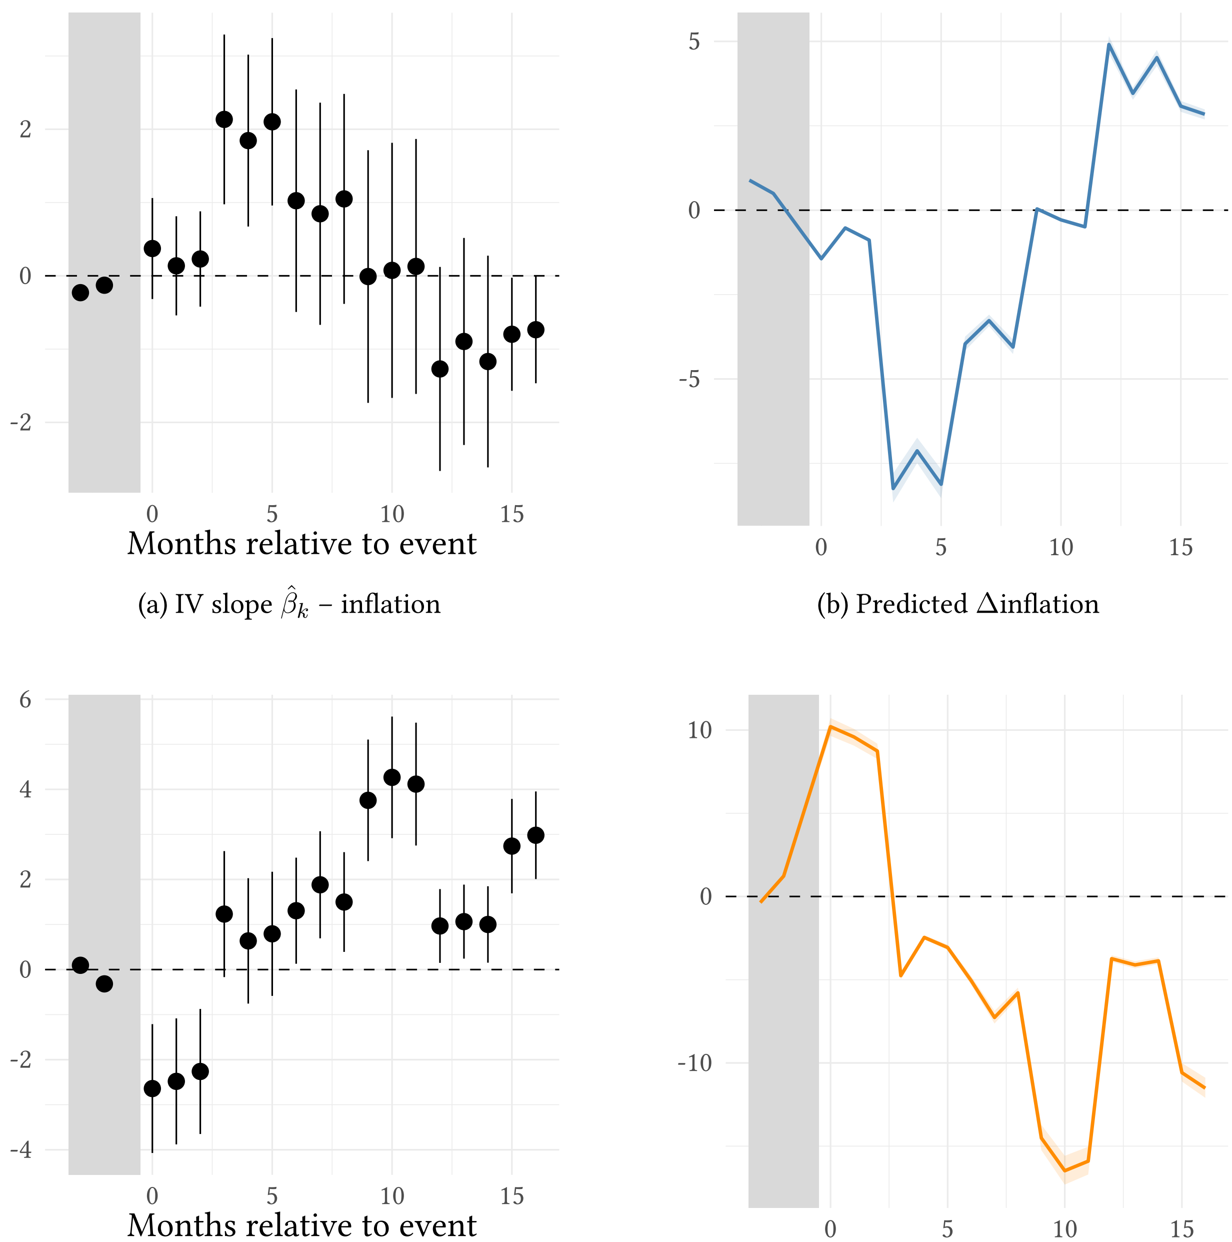

(a) IV slope $\hat{\beta}_k$ – inflation

(b) Predicted Δinflation

(c) First-stage $\hat{\beta}_k^{(\pi^e)}$

(d) Predicted $\Delta\pi^e$

impacting inflation expectations and passing through to inflation. Similarly, if the supply-chain disruptions had differential impact across groups, and hence across differentially exposed regions, then one can imagine an offsetting effect (higher goods prices induce substitution into travel by younger consumers). The results, though, suggest a strong positive pass-through.

Figure 8: Event study analysis: COVID inflation surge (2020-2022)

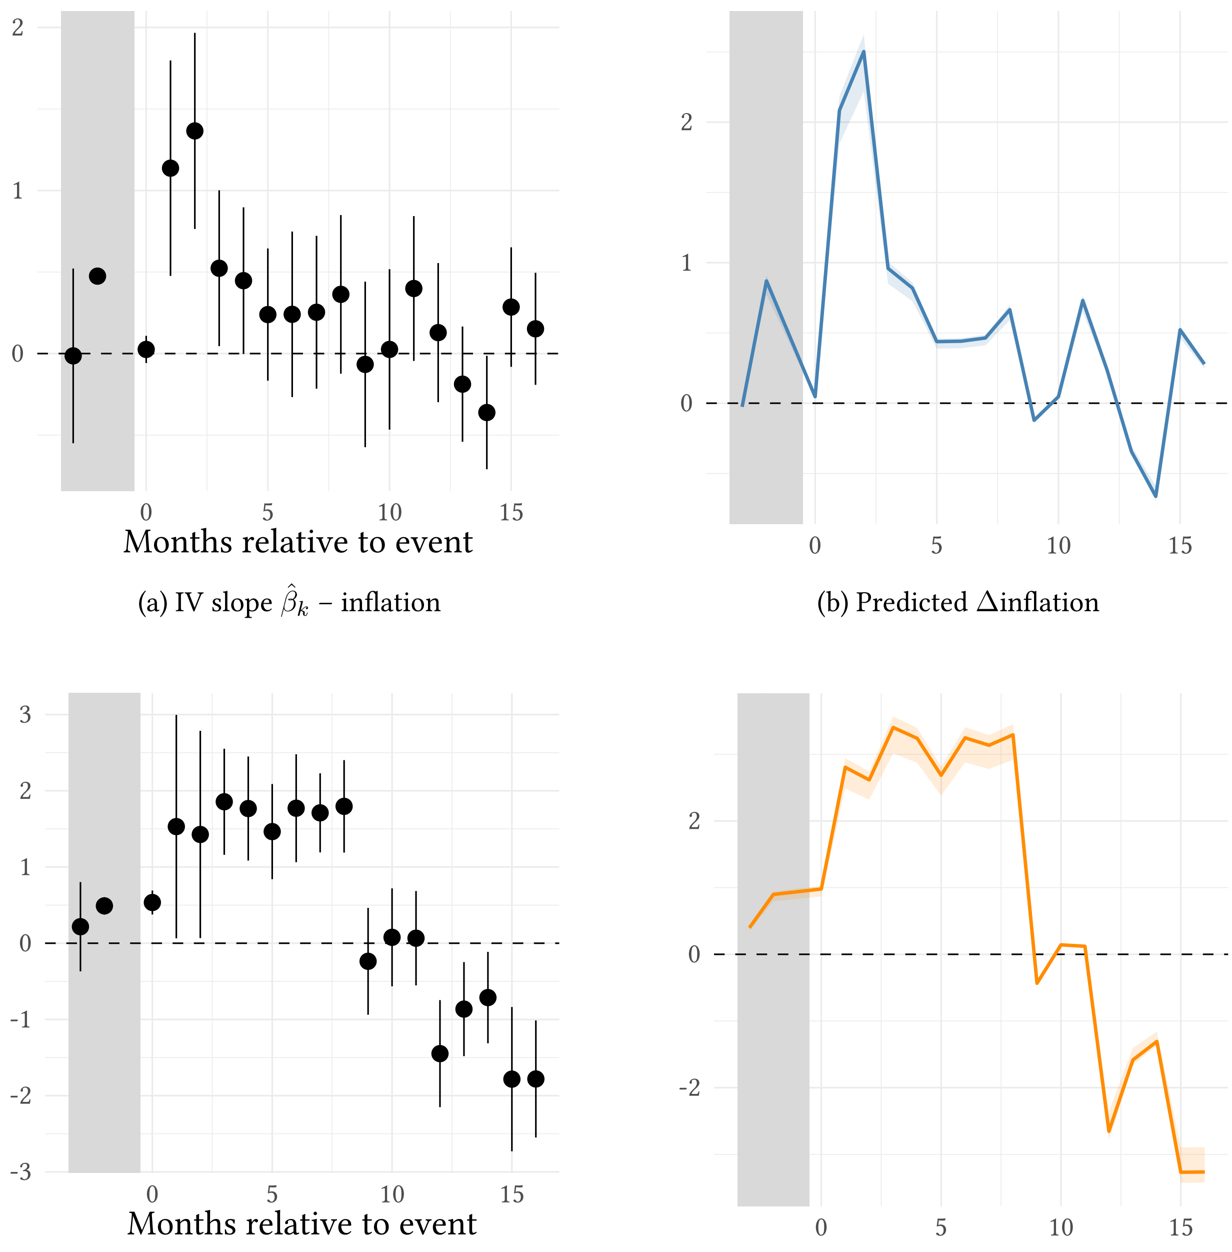

(a) IV slope $\hat{\beta}_k$ – inflation

(b) Predicted Δinflation

(c) First-stage $\hat{\beta}_k^{(\pi^e)}$

(d) Predicted $\Delta\pi^e$

Figure 8 details the results. Panel (c) shows immediate and sustained differences in regional expectations based on demographic exposure. Contrary to the Volcker episode, there is no lag or sign reversals. High-exposure regions (those with more young, educated households) immediately show higher expectations, likely reflecting both greater attentiveness and exposure to inflationary factors. Panel (d) confirms these differences occur rapidly and persist.

Panel (a) reveals the pass-through remains positive throughout—regions experiencing larger expectation increases see higher inflation—but with different dynamics than Volcker. Effects are front-loaded, peaking within 2-3 months then moderating as supply constraints and fiscal transfers ease, while monetary policy tightens.

Mechanics.

These descriptive events offer three insights into what the Bartik instrument identifies. First, the positive pass-through persists across different shock types, consistent with the structural interpretation. Whether expectations fall due to monetary contraction or rise due to supply disruptions or fiscal stimulus, higher expectations consistently feedback to higher inflation. The magnitude of the response, though, can vary according to the specific underlying shock.

Second, dynamics vary with shock characteristics. Monetary policy operates with lags and current expectations may not fully capture anticipatory effects, generating the oscillatory patterns in the Volcker episode. The supply shocks and fiscal transfers push inflation in the same direction, generating a clean inflation response during the pandemic. These differences confirm the identified effect captures weighted averages across different shock effects, with the full-sample estimates likely overweighting volatile periods.

Third, the episodes support the idea that monetary policy operates through the identified channel. The Volcker episode ultimately shows the expected pattern of the key households adjusting expectations to monetary policy and these adjustments driving inflation differentials across regions. The initial noise and lags make this confirmation more credible, as they reflect realistic policy transmission rather than an implausibly clean experiment.

These patterns do not prove the instrument is valid—that rests on the exclusion restriction. Rather, they illustrate that the identified reduced-form relationships operate sensibly during well-understood historical episodes. This strengthens confidence that the main estimates capture economically meaningful variation in how expectations transmit to inflation across different shock types, even if the precise mixture of shocks varies over the sample.

5 Robustness

This section probes robustness along four key dimensions: (i) coverage and coarser groupings; (ii) placebos that preserve the shift–share structure but remove relevant variation; (iii) alternative measurement choices; (iv) tests for channels that could violate exclusion (markups, cost correlates); and, (v.) a specification allowing for common correlated effects.

5.1 Alternative group definitions and coverage

A natural concern is whether the benchmark pass-through depends on the fine-grained nature of the 160-group partition or is driven by sparse cells. To address this, this subsection reports on reconstructing the entire instrument under a sequence of alternative partitions: medium (80 groups), coarse (32), very coarse (16), partitions that omit specific attributes (no education, no sex), and single-attribute partitions (age-only, education-only).

The analysis for each partition $\mathcal{P}$ proceeds in two steps: (i) re-aggregate Michigan survey responses into national group expectations $\pi_{g,t}^{e,\mathcal{P}}$ using the same leave-one-out procedure by region, and (ii) recompute regional exposure shares $\mu_{r,g}^{\mathcal{P}}$ using the CPS 1978 data, following the baseline specification. The resulting shift-share instrument is:

$$z_{r,t}^{\mathcal{P}} = \sum_{g \in \mathcal{G}_{\mathcal{P}}} \mu_{r,g}^{\mathcal{P}} \pi_{-r,g,t}^{e,\mathcal{P}} \tag{16}$$

where both the shifts and shares vary with the partition choice. Estimation maintains the identical region-month samples, controls, and fixed effects across all specifications.

Figure 9b reports the 2SLS pass-through estimates with 95% confidence intervals for each partition, while Figure 9a displays the corresponding first-stage F statistics. Two key patterns emerge:

1. **Effect stability across groupings.** Reconstructing both components of the instrument yields point estimates close to the benchmark, with overlapping confidence intervals across specifications. Neither dropping education from the partition nor adopting coarser groupings significantly changes the estimates. Single-attribute groupings (age-only, education-only) maintain positive estimates with strong first stages, though the age-only specification yields a somewhat smaller point estimate.

2. **First-stage strength and coverage.** The Kleibergen-Paap F statistics remain well above conventional thresholds across all partitions. As group definitions become coarser, the F statistic typically increases—a consequence of (a) larger sample sizes per group-month when constructing national shifts and (b) reduced measurement error from aggregating cells—without affecting the point estimates.

Table 12 quantifies coverage by partition, reporting the number of groups, panel size, median respondents per group-month used to construct $\pi_{g,t}^{e,\mathcal{P}}$, and the share of months with at least 2 and 5 respondents. While coverage improves monotonically with coarser partitions, the benchmark fine partition already achieves substantial coverage. The stability documented in Figure 9 confirms that the main results are not driven by sparse cells.

The pass-through effect is robust to partition choice—it neither depends on including education in the demographic grouping nor on the specific level of granularity. Reconstructing the entire Bartik instrument under alternative, economically meaningful demographic groupings yields similar estimates with strong first stages. Combined with the stacked-IV evidence in Section 4.3, these results support the existence of an effect that remains consistent regardless of how demographic composition is characterized.

5.2 Probing identification

Exogeneity of initial shares. The first set of diagnostics asks whether 1978 CPS demographic shares are systematically aligned with macroeconomic factors that could independently drive regional price movements. Table 14 reports (i) a cross-section test that relates the 1978 state-level share vector to 1977Q1 $\rightarrow$ 1978Q1 changes in three state covariates – per-capita personal income (PCPI), the unemployment rate (UR), and the FHFA state house price index (STHPI)– and (ii) panel tests that measure the partial R^2 of the orthonormal basis of 1978 shares for state inflation conditioning on time fixed effects.[17] In the cross-section, the joint fit of changes in $\{PCPI, STHPI, UR\}$ to the share vector is modest, and permutation tests fail to reject the null of no alignment at conventional levels. In the panel, the share-basis explains only a small fraction of residual variation after absorbing time

[17]The adding-up constraint (shares sum to one) makes the share matrix rank-deficient. We construct an orthonormal basis via QR decomposition to preserve all cross-state variation while avoiding collinearity.

Figure 9: Robustness of estimates to alternative demographic groupings. Pass-through estimates remain stable across partitions ranging from 160 fine groups to 2 coarse groups, while first-stage strength generally improves with coarser groupings.

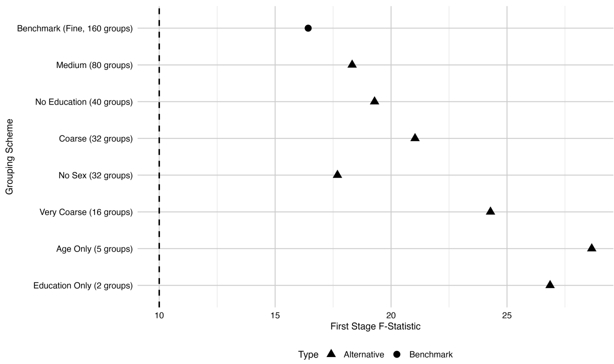

(a) First stage F-statistics across grouping schemes. Gray dashed line shows weak instrument threshold (F = 10). All specifications exceed the threshold.

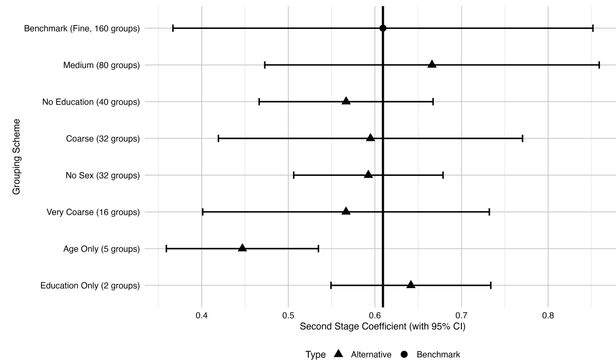

(b) Second-stage coefficients with 95% confidence intervals. Red dashed line shows benchmark estimate (0.61). All estimates fall within 0.45-0.67 range.

Table 12: Coverage and first-stage strength by grouping scheme

Grouping scheme	# groups	N (panel)	Med. g-m	% months $\geq$2	% months $\geq$5	KP F	$\hat{\beta}$ (SE)
Benchmark (Fine, 160 groups)	160	49253	3.8	51	23	16.4	0.6093 (0.1237)
Medium (80 groups)	80	22519	8.0	83	55	18.3	0.6659 (0.0985)
No Education (40 groups)	40	17805	7.0	79	55	19.3	0.5667 (0.0512)
Coarse (32 groups)	32	15590	12.0	87	66	21.0	0.5949 (0.0896)
No Sex (32 groups)	32	20272	7.0	74	50	17.7	0.5924 (0.044)
Very Coarse (16 groups)	16	8220	24.0	94	79	24.3	0.5666 (0.0844)
Age Only (5 groups)	5	2660	95.0	100	100	28.7	0.447 (0.0448)
Education Only (2 groups)	2	1064	244.5	100	100	26.9	0.6416 (0.0471)

Note:
"Med. g-m" is the median number of respondents per group-month used to form national group shocks. Percent columns report the share of survey months with at least 2 or 5 respondents in that grouping. The corresponding $\hat{\beta}$ and KP F are shown in Figure 9b and Figure 9a, respectively.

fixed effects, and permutation tests likewise fail to reject. Overall, these results are consistent with the exogeneity of the 1978 demographic shares for the pass-through estimates reported previously.

Markups and Group Composition. A more subtle concern arises from consumption heterogeneity in the model. If demographic groups systematically differ in their willingness to pay for identical goods, regions with different demographic compositions might exhibit different markup levels (see Section 3.1.4). Such markup heterogeneity could create a spurious correlation between group shares and inflation that operates outside the expectations channel.

To investigate this possibility, state-level markup proxies are constructed using labor's share of output from Nakamura and Steinsson (2014) and examine their relationship with the demographic shares of key groups identified in Table 8. Figure 10 presents scatter plots of state-level markups against group shares for groups present in all states in the 1978 CPS. The relationships are predominantly flat or slightly negative, suggesting that states with higher concentrations of key groups, if anything, have *lower* markups.

Table 13 formalizes this analysis, regressing state markups on the combined share of top-10 groups.

The coefficient is statistically insignificant (p = 0.27) with an R^2 of 0.004, indicating that demographic composition explains virtually none of the cross-state markup variation. Panel regressions with group fixed effects (not shown) yield similarly null results.

These findings alleviate concerns about endogenous markup variation. Even if markups and expectation pass-through were correlated with demographic distributions, the exclusion restriction would remain valid as long as group shares do not directly affect inflation rates—which these tests support.

Figure 10: State-level markups versus demographic group shares, 1978. Each panel plots a key group's population share against the state markup proxy. Groups shown are those present in all states in the CPS 1978.1 sample.

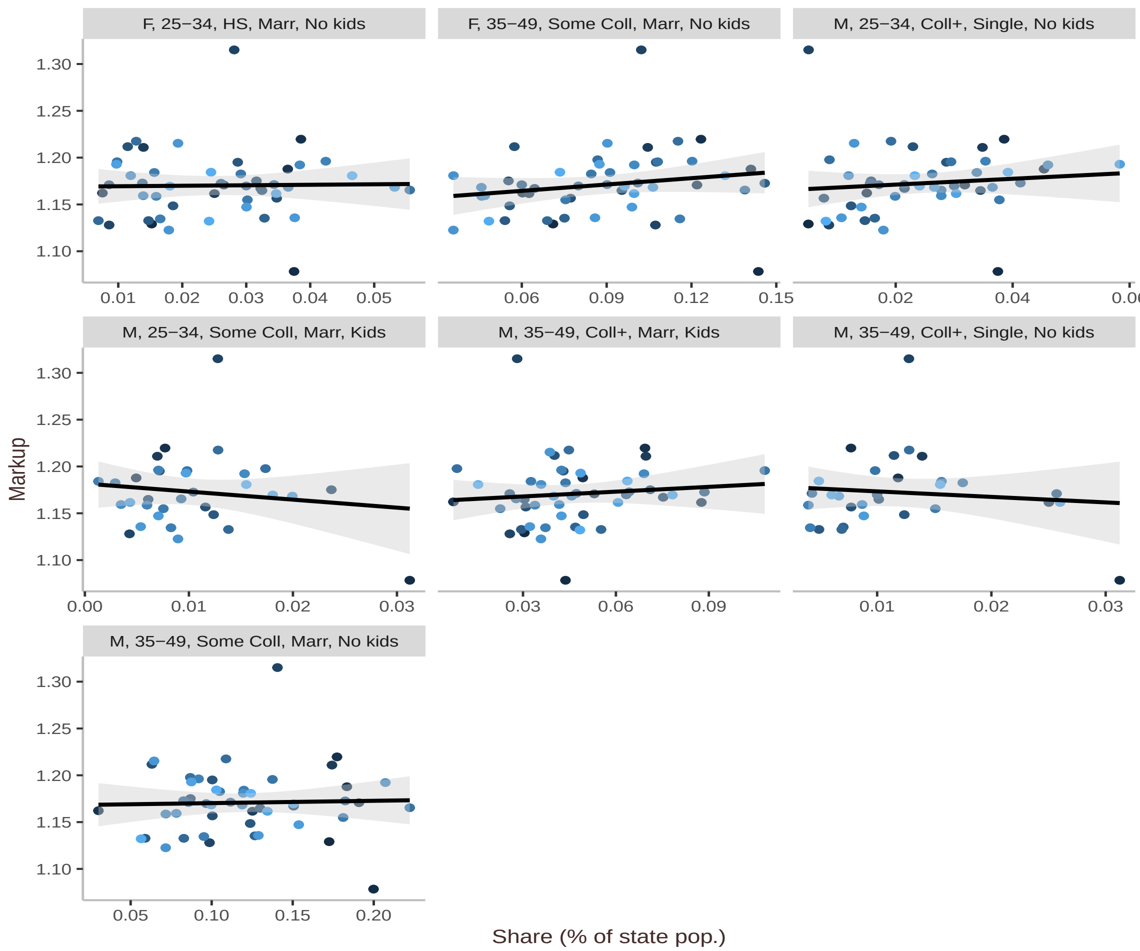

5.3 Sectoral Inflation

The next analysis examines how expectations pass through to different components of the price index. This decomposition identifies which sectors drive the aggregate effect and tests predictions from the literature. In particular, D'Acunto, Malmendier, Ospina, and Weber (2021) suggest that pass-through should be strongest for non-durable goods, where consumers observe prices most frequently.

Table 15 presents 2SLS estimates for major CPI components. The results reveal heterogeneity across sectors:

Table 13: State markups and demographic composition

	Coefficient	Std. Error	t-statistic	p-value
Top-10 group share	0.0939	0.0842	1.12	0.2656

Note:
Dependent variable: 1978 state-level markup (output/labor compensation). Independent variable: demographic group share (proportion of state population in top-10 weighted groups). Fixed effects regression estimated across groups, controlling for group-specific effects.
* Within R^2 = 0.0038

Table 14: Exogeneity diagnostics using CPS 1978 shares.

Specification	R^2	Perm. p
Cross-section		
All shares	0.0864	0.0775
Top-10 shares	0.0966	0.0845
Panel		
Year FE	0.0430	0.2900
Year×Quarter FE	0.0484	0.5010

Notes: Cross-section: regress 1978 CPS state share vector on 1977Q1→1978Q1 changes in PCPI (BEA), UR (BLS), and STHPI (FHFA). "Top-10 shares" keeps the ten most variable shares across states. Panel: partial R^2 of the 1978 share-basis for state inflation with Year or Year×Quarter fixed effects. Permutation p: one-sided right-tail tests—cross-section permutes the state pairing between changes and shares; panel permutes time-FE-residualized inflation across states; 2,000 (cross) / 1,000 (panel) draws. Small p rejects the null of no association (cross) or zero explanatory power (panel); otherwise the test fails to reject. State counts reflect available Q1 data.

Commodities show strong pass-through. The effect is particularly pronounced for non-durables, with a coefficient of 1.59 (SE = 0.61). The overall commodities coefficient of 1.17 (SE = 0.45) is nearly double the baseline aggregate estimate. In contrast, durables show no meaningful pass-through (-0.06, SE = 0.22), consistent with their infrequent purchase and longer planning horizons.

Services exhibit weaker effects. The services coefficient of 0.15 (SE = 0.10) is positive but imprecisely estimated. Excluding housing services strengthens the estimate to 0.35 (SE = 0.17), suggesting that the stickiness of rental contracts may soften pass-through in that sector. Services excluding medical care also show modest effects.

These patterns are intuitive: sectors with more frequent transactions and flexible pricing (non-durables) exhibit stronger expectation pass-through than those with sticky prices or infrequent purchases (durables, housing services).

5.4 Alternative Specifications

Table 16 addresses potential data and specification concerns through four robustness checks, each maintaining the Bartik instrument construction with Michigan survey shares.

Table 15: 2SLS by component inflation

	Commodities			Services		
	commodities	non-durables	durables	services	services-house	services-med.
Dependent Var.:	commod. infl.	non-dur. infl.	dur. infl.	serv. infl.	serv.-rent infl.	infSlm
pe	1.168** (0.4451)	1.589** (0.6083)	-0.0646 (0.2223)	0.1469 (0.0956)	0.3478* (0.1747)	0.1601 (0.1016)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fixed-Effects:	——————-	——————-	——————-	——————	——————-	——————
REGION	Yes	Yes	Yes	Yes	Yes	Yes
TIME	Yes	Yes	Yes	Yes	Yes	Yes
________________	__________________	__________________	__________________	________________	__________________	________________
S.E. type	Dris.-Kra. (L=4)	Dris.-Kra. (L=4)	Dris.-Kra. (L=4)	Dri.-Kra. (L=4)	Dris.-Kra. (L=4)	Dri.-Kra. (L=4)
Observations	1,403	1,403	1,403	1,403	1,395	1,403
R2	0.93427	0.90670	0.98696	0.95343	0.92898	0.94972
Within R2	0.13476	0.06935	0.64615	0.77100	0.51282	0.77032

Note:
Reports panel regression results where the outcome variable is component inflation. Commodities include all non-durable and durable goods. Services-housing and services-med. remove housing and medical services, respectively.

Sample restrictions (columns 1-2). The first specification excludes outlier expectations exceeding ±25% to address measurement error concerns. The resulting coefficient of 0.67 (SE = 0.25) slightly exceeds the baseline, remaining significant at the 1% level. While reassuring, the full sample remains preferable as extreme expectations may reflect genuine behavioral biases. The second specification restricts the analysis to first-time survey respondents (approximately 60% of the sample) to address potential panel conditioning effects. The estimate of 0.58 (SE = 0.47) remains positive but less precise, likely due to the reduced sample size.

Alternative inflation measure (column 3). This specification reconstructs regional inflation using state-level CPIs from Hazell, Herreno, Nakamura, and Steinsson (2022), weighted by consumer expenditure shares. This approach addresses the limited time coverage of BLS regional CPIs and provides an independent inflation measure. The coefficient of 0.67 (SE = 0.31) closely matches the benchmark, despite a shorter sample period.

Lagged instrument (column 4). To address potential contemporaneous endogeneity concerns, this specification instruments using twelve-month lagged Michigan survey shares rather than contemporaneous shares. This approach maintains the CPS 1978 property that shares are predetermined while introducing time-variation in those shares. The estimate of 0.64 (SE = 0.29) is comparable to the higher end of baseline IV estimates.[18]

Across all specifications, the pass-through coefficient remains positive, economically meaningful, and generally statistically significant. The comparability of estimates across diverse robustness checks–different samples, inflation measures, and instrument timing–reinforces confidence in the main findings.

[18]The Appendix also reports 2sls estimates where contemporaneous inflation is regressed on lagged inflation expectations.

Table 16: Alternative estimates

	small	first-only	state-CPI	lag michigan shares
Dependent Var.:	RegInf	RegInf	RegInf	RegInf
pe	0.6669** (0.2546)	0.5838 (0.4747)	0.6558* (0.3012)	0.6409* (0.2938)
Controls	Yes	Yes	Yes	Yes
Fixed-Effects:	------	-----	------	------
REGION	Yes	Yes	Yes	Yes
TIME	Yes	Yes	Yes	Yes
______________	________________	______________	_______________	_______________
S.E. type	Drisc.-Kra. (L=4)	Dri.-Kra. (L=4)	Dris.-Kra. (L=4)	Dris.-Kra. (L=4)
Observations	1,387	1,387	1,399	1,387
R2	0.93325	0.91304	0.92265	0.93074
Within R2	0.35334	0.15758	0.28862	0.32908

Note:
Reports panel regression results for a variety of alternative specifications. "small" removes large survey responses. "first" includes only first-time survey respondents. "state-cpi" measures regional inflation by aggregating state-level CPI's. "lag michigan shares" instruments with 12 month lagged survey shares.

5.5 Common Correlated Effects

The baseline specification could overlook two important features in firms' price-setting behavior. First, national firms often set uniform prices across regions (DellaVigna and Gentzkow (2019)), potentially creating common pricing effects. Second, aggregate shocks could generate correlated time effects across regions that standard fixed effects fail to capture. This section addresses both concerns using a common correlated effects (CCE) framework.

Following Pesaran (2006) and Bai (2009), this subsection presents estimates from a model with a multi-factor error structure that allows for both observable and unobservable common effects:

$$\pi_{r,t} = \beta\pi^e_{r,t} + \gamma' x_{r,t} + \phi\pi^e_t + \upsilon_{r,t} \tag{17}$$
$$\upsilon_{r,t} = \lambda'_r F_t + \epsilon_{r,t} \tag{18}$$
$$\pi^e_{r,t} = \zeta z_{r,t} + \vartheta\pi^e_t + \Lambda'_r F_t + u_{r,t} \tag{19}$$

This specification decomposes the pass-through effect into two components: regional expectations ($\pi^e_{r,t}$) and national expectations (π^e_t). The term $\lambda'_r F_t$ captures unobserved common factors with heterogeneous regional loadings—these "interactive fixed effects" allow different regions to respond differently to the same aggregate shocks. The national expectation measure π^e_t uses the median inflation expectation from the Michigan survey (FRED: MICH).

The estimation follows Harding and Lamarche (2011), who extend the CCE approach to handle endogenous regressors through instrumental variables. The procedure augments the 2SLS regression with cross-sectional averages of all explanatory variables, which proxy for the unobserved factors F_t and address both the standard endogeneity concern and potential correlation between the interactive fixed effects and regional expectations.

Table 17: Decomposition of expectation pass-through using common correlated effects. IV estimation with interactive fixed effects and jackknife bias correction separates regional (π^e_R) and aggregate (π^e) expectation channels. Results show national expectations dominate, with total pass-through ($\beta + \phi$) of 1.32.

	no controls		controls	
π^e_R coeff.	0.3370	**	0.4715	**
π^e coeff.	0.9803	*	0.8500	*

Note:
Estimates pass-through of regional inflation expectations and median national expectations using the IV common effects estimator with interactive fixed effects. The first column excludes exogenous controls. The second column includes controls. Applies the Jacknife bias correction.
[1] sign. levels: ** = 0.05, * = 0.10

Table 17 presents the second-stage estimates with jackknife bias corrections. Two key conclusions are apparent from the results:

1. **Regional effects remain significant.** The regional pass-through coefficient decreases from approximately 0.6 in the baseline to 0.47 with controls (0.34 without controls), reflecting the portion of regional variation that is genuinely local rather than driven by common factors.

2. **Aggregate expectations dominate.** The coefficient on aggregate expectations ranges from 0.85 to 0.98, roughly double the regional coefficient. The total pass-through effect—combining regional and national channels—equals $\hat{\beta} + \hat{\phi} = 1.32$ without controls and 1.32 with controls, indicating that a one percentage point increase in both regional and national expectations translates to a 1.3 percentage point increase in regional inflation.

These estimates require careful interpretation. The CCE estimator achieves consistency as the number of regions grows large, but with only four census regions, finite-sample concerns remain relevant. Additionally, while the jackknife correction addresses some small-N bias, its optimality in a model with interactive fixed effects and instrumental variables is not established.

Apart from these caveats, the results suggest that accounting for common effects matter for the estimated aggregate pass-through effect. The dominance of aggregate expectations highlights how aggregate sentiment drives regional price dynamics, consistent with national pricing strategies and common macroeconomic factors. If the relative magnitudes from Table 17 apply to the baseline estimates in Table 7, the implied aggregate pass-through would be approximately 1.8—consistent with the upper range of estimates reported in Section 4.2. This finding underscores the importance of distinguishing between local and aggregate expectation channels when measuring inflation pass-through.

6 Long-Horizon Expectations

Having established that one-year inflation expectations pass through to regional inflation, we now turn to the question of what role long-run expectations play. The Michigan survey asks consumers about inflation over the "next 5-10 years," providing a direct measure of long-horizon inflation expectations. Central banks often assign primacy to long-run expectation anchoring, viewing it as a key determinant for price stability. If long-horizon expectations drive inflation independently of short-run beliefs, this would underscore their importance.

This subsection presents estimates of the pass-through from both short and long-horizon expectations using separate Bartik instruments for each horizon. The long-horizon series begins with continuous coverage in 1990.

Figure 11 summarizes the findings. Panel (a) plots the first-stage and makes apparent that the instrument is a good predictor of long-horizon expectations, while panel (b) shows the reduced form–the relationship between the long-horizon instrument and inflation is flat, hovering near zero across the entire sample. Table 20 presents first-stage results when only including, and instrumenting for, long-run expectations. The Bartik instrument successfully predicts five-to-ten-year beliefs across all specifications, with F-statistics ranging from 62 to 91 without time fixed effects. Including time effects reduces but does not eliminate predictive power (F-statistics of 10-14). However, the first-stage R^2 values reveal that, even in the strongest specifications, long-run expectations exhibit far less variation than their short-run counterparts.[19] With two-way fixed effects, the first-stage R^2 reaches only 0.38 compared to approximately 0.78 for one-year expectations in comparable specifications.

Short and long expectations. Table 19 presents findings from estimating:

$$\pi_t^r = \beta_1 \pi_{r,t}^e + \beta_5 \pi_{r,t}^{e,5-10} + \gamma' x_{r,t} + \delta_r + \mu_t + \varepsilon_{r,t} \tag{20}$$

The results show that short-run expectations maintain their pass-through: with two-way fixed effects (CPS shares), $\hat{\beta}_1 = 0.609$ (SE = 0.244), virtually identical to the baseline estimate. Long-run expectations, by contrast, have no measurable effect: $\hat{\beta}_5 = 0.00$ (SE = 0.052). This pattern holds across all specifications. Without any fixed effects—where identification is strongest and both instruments have maximum power—the long-run coefficient reaches only 0.044 (SE = 0.022) compared to 0.310 (SE = 0.044) for short-run expectations. Even this small effect vanishes once time trends are controlled. Michigan shares yield identical conclusions.

Long expectations. Table 21 examines whether long-run expectations affect inflation when estimated alone, without short-run beliefs. The answer remains no. With two-way fixed effects, the coefficient is -0.041 (SE = 0.047) for CPS shares and -0.002 (SE = 0.041) for Michigan shares—statistically and economically indistinguishable from zero. This null result is robust. Across eight different specifications (varying fixed effects and share constructions), not one produces an economically meaningful positive coefficient. The absence of any effect, even without controlling for potentially collinear short-run expectations, reinforces that long-horizon beliefs simply do not drive current regional inflation.

[19] Figure 11, panel (a) scatters regional inflation expectations against the long-horizon instrument, with the trend line capturing the unconditional correlation between the instrument and long expectations. The strong positive slope indicates the instrument is a good predictor of long-horizon expectations. Table 20, on the other hand, reports conditional first-stage regressions that include region and time fixed effects and additional controls. Once common variation is absorbed by time fixed effects, the remaining variation is limited, and the estimated coefficients are attenuated and even negative.

Interpretation. In post-1980s US data, long-run inflation expectations do not affect current regional inflation through the variation identified by the shift-share instrument. While one-year expectations show robust pass-through of approximately 0.6, five-to-ten-year expectations contribute nothing measurable—whether estimated jointly with short-run beliefs, in isolation, or in reduced form. This stark difference has important implications.

First, the results challenge the common premise that long-run expectation anchoring directly stabilizes current inflation. At least through the regional variation channel studied here, five-to-ten-year expectations appear disconnected from near-term price dynamics. Firms and consumers making current pricing and purchasing decisions respond to short-run beliefs. Second, the findings help explain why inflation can surge despite well-anchored long-run expectations, as occurred in 2021-2023. If current inflation responds primarily to short-run beliefs, then stable long-run anchoring provides little immediate restraint on price dynamics.

However, it is important to note several caveats. The sample period with long-horizon expectations coincides with the Great Moderation, when long-run expectations exhibited historically low volatility. The first-stage R^2 values confirm that long-run beliefs vary less than short-run expectations, making effects harder to detect. Additionally, this analysis examines only current inflation. Long-horizon expectations may influence investment decisions, wage contracts, or inflation persistence in ways not captured here. Nevertheless, in modern US data, regional inflation responds to what consumers expect over the next year, not what they expect over the next decade. This finding suggests that stabilizing short-run expectations may be more important for price stability than previously recognized.

7 Conclusion

Understanding how subjective inflation expectations feed into realized inflation remains central to both macroeconomic theory and policy. Many household and firm decisions depend on beliefs about future prices and real interest rates, yet the quantitative importance of these beliefs continues to be debated. Standard approaches estimate the New Keynesian Phillips curve under rational expectations, instrumenting for expected inflation. While informative, these methods require explicit assumptions about expectation formation and face challenges of weak instruments and limited cross-sectional variation.

This paper takes a different route. Using the Michigan Survey of Consumers, it exploits rich microdata on demographic heterogeneity to identify the causal impact of subjective inflation expectations on (regional) inflation. The empirical design is a quasi-experimental shift–share: regional exposure to aggregate shocks in group-level expectations is weighted by as-good-as-random demographic shares. This differential-exposure framework provides a credible instrument without requiring assumptions about how expectations are formed. Demographic heterogeneity may arise from differences in consumption baskets, attentiveness, or information acquisition—the analysis is agnostic on this point—and identification relies only on cross-sectional variation in exposure.

The results indicate a significant and economically meaningful effect of short-horizon expectations on realized inflation. A one–percentage point increase in expected inflation raises regional inflation by roughly 60 basis points. The effect is substantially larger than the OLS correlation, underscoring the endogeneity of expectations. At the same time, long-horizon expectations (5–10 years ahead) have little or no predictive power once short-run expectations are controlled for—a pattern consistent with models such as Werning (2022) where short-term expectations mainly impact price-setting decisions, as well as with the view that long-run expectations are anchored under credible inflation

Table 18: 2SLS with short and long expectations: first stage

Panel A: CPS shares				
	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	1.08 (0.05)	1.10 (0.07)	0.23 (0.04)	0.31 (0.05)
π^e_5	0.57 (0.07)	0.58 (0.08)	-0.30 (0.09)	-0.29 (0.09)
u^r	0.53 (0.06)	0.46 (0.04)	0.74 (0.09)	0.53 (0.04)
R^2	0.688	0.716	0.766	0.781
KP-F stat	351.4	386.5	6.2	11.1

Panel B: Michigan shares				
	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	0.74 (0.03)	0.74 (0.03)	0.22 (0.01)	0.23 (0.01)
π^e_5	0.35 (0.04)	0.37 (0.04)	-0.21 (0.05)	-0.18 (0.05)
u^r	0.52 (0.06)	0.46 (0.04)	0.74 (0.09)	0.53 (0.04)
R^2	0.696	0.715	0.770	0.784
KP-F stat	373.4	382.3	14.7	16.4

Note: Driscoll–Kraay standard errors in parentheses. KP-F is the Kleibergen–Paap rank-Wald weak-instrument statistic for the first-stage regression of long-run inflation expectations p^e on the Bartik shift–share instrument. Instruments: Bartik shift–share for the indicated horizons (1-year and 5-10 year ahead expectations) together with the 12-month lag of the unemployment rate u^r_{t-12}. All specifications include two lags of regional inflation, the full set of survey-covariate controls, and the fixed effects indicated at the top of each column.

Table 19: 2SLS estimates: both 1-year and 5-year expectations

Panel A: CPS shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	0.310 (0.044)***	0.303 (0.043)***	0.750 (0.399)*	0.609 (0.244)**
π_5^e	0.044 (0.022)**	0.037 (0.021)*	-0.004 (0.059)	-0.000 (0.052)
u^r	-0.007 (0.031)	-0.000 (0.041)	0.027 (0.047)	-0.134 (0.083)
R^2	0.923	0.926	0.911	0.932
Within R^2	0.923	0.925	0.189	0.340
N	1,531	1,531	1,531	1,531

Panel B: Michigan shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	0.268 (0.040)***	0.271 (0.041)***	0.387 (0.156)**	0.378 (0.141)***
π_5^e	0.033 (0.023)	0.030 (0.022)	0.005 (0.036)	-0.005 (0.040)
u^r	0.004 (0.033)	0.009 (0.043)	0.009 (0.028)	-0.137 (0.064)**
R^2	0.927	0.928	0.948	0.951
Within R^2	0.927	0.928	0.528	0.526
N	1,531	1,531	1,531	1,531

Note: Driscoll–Kraay standard errors in parentheses. Instruments: Bartik shift–share for the one-year and five-10 year horizons shown together with the 12-month lag of unemployment u^r_{t-12}. All regressions include two lags of regional inflation, survey-covariate controls, and the fixed effects indicated above each column.

Table 20: 2SLS with long expectations: first stage

Panel A: CPS shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π_5^e	0.59 (0.08)	0.61 (0.09)	-0.29 (0.08)	-0.28 (0.08)
u^r	0.52 (0.06)	0.46 (0.04)	0.74 (0.09)	0.53 (0.04)
R^2	0.153	0.162	0.366	0.379
KP-F stat	88.1	91.2	11.8	10.7

Panel B: Michigan shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π_5^e	0.36 (0.04)	0.38 (0.05)	-0.20 (0.05)	-0.18 (0.05)
u^r	0.51 (0.06)	0.45 (0.04)	0.74 (0.09)	0.53 (0.04)
R^2	0.126	0.140	0.368	0.379
KP-F stat	62.4	69.0	14.1	10.5

Note: Driscoll–Kraay standard errors in parentheses. KP-F is the Kleibergen–Paap rank-Wald weak-instrument statistic for the first-stage regression of long-run inflation expectations p^e on the Bartik shift–share instrument. Instruments: Bartik shift–share for the long horizon together with the 12-month lag of the unemployment rate u^r_{t-12}. All specifications include two lags of regional inflation, the full set of survey-covariate controls, and the fixed effects indicated at the top of each column.

Table 21: 2SLS estimates: 5-year expectations

Panel A: CPS shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π_5^e	0.008 (0.022)	0.003 (0.022)	-0.028 (0.044)	-0.041 (0.047)
u^r	0.061 (0.041)	0.076 (0.054)	-0.002 (0.020)	-0.142 (0.044)***
R^2	0.922	0.921	0.956	0.957
Within R^2	0.922	0.920	0.601	0.580
N	1,531	1,531	1,531	1,531

Panel B: Michigan shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π_5^e	-0.001 (0.025)	-0.004 (0.025)	0.008 (0.037)	-0.002 (0.041)
u^r	0.064 (0.042)	0.078 (0.055)	-0.008 (0.018)	-0.141 (0.044)***
R^2	0.922	0.920	0.958	0.960
Within R^2	0.922	0.920	0.617	0.611
N	1,531	1,531	1,531	1,531

Note: Driscoll–Kraay standard errors in parentheses. Instruments: Bartik shift–share for the five-ten year horizon shown together with the 12-month lag of unemployment u^r_{t-12}. All regressions include two lags of regional inflation, survey-covariate controls, and the fixed effects indicated above each column.

targeting. Limited variation in long-run expectations during the Great Moderation also contributes to the attenuation.

The identifying variation is concentrated in periods of elevated volatility—1978–82, 2007–09, and 2021–22—when group-level heterogeneity in expectations was greatest. This state dependence suggests that the estimated pass-through should be interpreted as an average effect of expectation shocks during turbulent periods, rather than normal times. Similar to Binder, Kamdar, and Ryngaert (2024), who use political heterogeneity to identify the expectation–inflation link, the results here imply that expectations matter most when disagreement is high and uncertainty elevated.

Overall, the findings support the view that short-term inflation expectations are an important causal driver of inflation dynamics, while long-run expectations play a stabilizing role. The approach provides a flexible framework for studying expectations in other contexts—macroeconomic, financial, or behavioral—where heterogeneous beliefs interact with aggregate shocks to shape economic outcomes.

A What shocks are identified by the instrument?

The 2SLS estimator captures the average causal response of inflation to expectations across the realized mixture of shocks that move expectations in our sample. This Appendix formalizes this interpretation. All variables below are the orthogonal complements with respect to controls and fixed effects (omitted for notational simplicity). The 2SLS estimator can be written, as in the main text, as

$$\widehat{\beta} = \frac{\mathrm{Cov}_t(\pi_{r,t}, z_{r,t})}{\mathrm{Cov}_t(\pi^e_{r,t}, z_{r,t})}.$$

To interpret this ratio, suppose group-level aggregate expectation changes are driven by a set of K (independent) underlying shocks u_{kt} with group coefficients λ_{gk}:

$$\pi^e_{g,t} = \sum_{k=1}^{K} \lambda_{gk} u_{kt} + \eta_{g,t}, \qquad \mathrm{Cov}(u_{kt}, u_{lt}) = 0 \ (k \neq l).$$

The shift–share instrument aggregates these into the regional instrument

$$z_{r,t} \equiv \sum_{g} \mu_{r,g} \pi^e_{g,t} = \sum_{k=1}^{K} b_{r,k} u_{kt} + \zeta_{r,t}, \quad b_{r,k} \equiv \sum_{g} \mu_{r,g} \lambda_{gk}.$$

Now assuming that regional expectations respond to the same shocks according to,

$$\pi^e_{r,t} = \sum_{k=1}^{K} \theta_{r,k} u_{kt} + \varepsilon^e_{r,t},$$

and denoting the shock-specific expectations passthroughs as $\{\beta_k\}_{k=1}^{K}$ so that,

$$\pi_{r,t} \ = \sum_{k=1}^{K} \beta_k \, \theta_{r,k} u_{kt} + \varepsilon_{r,t}.$$

Then, using orthogonality and exogeneity of the shift-share instrument $z_{r,t}$,

$$\mathrm{Cov}_t(\pi_{r,t}, z_{r,t}) = \sum_{k=1}^{K} \beta_k \underbrace{\left(\sum_{r} \theta_{r,k} b_{r,k} \mathrm{Var}_t(u_{kt}) \right)}_{\equiv \Xi_k}, \qquad \mathrm{Cov}_t(\pi^e_{r,t}, z_{r,t}) = \sum_{k=1}^{K} \Xi_k.$$

Hence

$$\widehat{\beta} = \sum_{k=1}^{K} \Omega_k \, \beta_k, \qquad \Omega_k \equiv \frac{\Xi_k}{\sum_{l=1}^{K} \Xi_l}, \quad \Omega_k \geq 0, \quad \sum_{k} \Omega_k = 1.$$

It follows that the 2SLS estimate is a weighted average of the shock-specific passthroughs $\{\beta_k\}$ with weights Ω_k that are larger when (i) the shock k impacts more strongly into expectations ($\theta_{r,k}$), (ii) the instrument is more exposed to that particular shock through the shift–share ($b_{r,k}$), and (iii) the shock is more volatile ($\mathrm{Var}(u_{kt})$). Equivalently, as explained in the main text, $\widehat{\beta}$ can be written as a Rotemberg-weighted average of just identified effects. When policy/demand-like episodes account for most of the impact, the estimate shifts toward the corresponding β_k; energy/relative-price episodes receive smaller weights and thus have limited influence.

B Robustness to survey weights

This appendix shows that the main results are robust to the use of survey weights in constructing the Bartik instruments. In the baseline specification, aggregate expectations from the Michigan Survey are computed without incorporating the survey weights designed to make responses nationally representative and account for variations in survey methodology over the sample. This section presents results using an intermediate demographic classification with survey weights applied throughout the instrument construction.

The Bartik instrument construction is modified in two key ways. First, define an intermediate demographic grouping that excludes marital and parental status while preserving the core demographic dimensions. This aligns the grouping with the survey design and sampling weights.

Second, national aggregate expectations are computed as:

$$\bar{\pi}^e_{gt} = \frac{\sum_{i \in g,t} w_i \pi^e_i}{\sum_{i \in g,t} w_i} \tag{21}$$

where w_i is the raw sampling weight. In the case of the CPS shares, this reweighted group aggregate expectation is interacted with the CPS 1978.1 shares. Michigan shares, on the other hand, are computed now as:

$$s^g_{rt} = \frac{\sum_{i \in g,r,t} w_i}{\sum_{i \in r,t} w_i} \tag{22}$$

This approach addresses potential concerns that the baseline specification may not be representative of the national expectations of each group.

Tables 22 and 23 present the first-stage and second-stage results, respectively, for one-year inflation expectations using the intermediate granularity with survey weights.

The first-stage results in Table 22 show that the weighted intermediate specification produces strong instruments. The second-stage estimates in Table 23 demonstrate that the main conclusions remain robust to this alternative instrument construction. Though, the estimated effects of inflation expectations on realized inflation are slightly smaller in magnitude compared to the baseline specification. The robustness exercise confirms that the main results do not depend critically on using unweighted national expectations.

C Placebo Tests for Shift-Share Instrument Validity

To validate the shift-share identification strategy, this section conducts a placebo test that randomizes the components of the Bartik instrument. The double randomization placebo tests whether the results reflect genuine economic relationships or just spurious correlations.

The placebo test breaks the two key links in the identification strategy:

1. Michigan survey participants are randomly reassigned to different demographic groups, severing the connection between group specific conditions and survey responses within demographic categories. That is, a young college-educated married woman with no kids who provided an inflation expectation might be randomly reassigned the label of, say, 'older high school-educated single men,' but her recorded expectation remains unchanged.

2. Randomly shuffle the 1977 CPS employment shares across demographic groups within each region, breaking the link between regional demographic composition and regional economic exposure.

Table 22: First-stage: intermediate granularity with survey weights (1-year expectations)

Panel A: CPS shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	0.62 (0.02)	0.61 (0.02)	0.32 (0.04)	0.32 (0.03)
u^r	0.52 (0.05)	0.46 (0.04)	0.78 (0.06)	0.63 (0.02)
R^2	0.703	0.716	0.792	0.802
KP-F stat	582.5	578.0	65.8	65.4

Panel B: Michigan shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	0.93 (0.03)	0.92 (0.03)	0.85 (0.03)	0.83 (0.02)
u^r	0.52 (0.06)	0.46 (0.04)	0.78 (0.06)	0.63 (0.02)
R^2	0.846	0.851	0.857	0.861
KP-F stat	1889.0	1835.4	432.2	413.1

Note: Driscoll–Kraay standard errors in parentheses. Both Michigan and CPS demographic groups use intermediate granularity (40 groups: age × education × gender, excluding marital status and children) with survey weights for Michigan expectations and CPS shares computed without weights. Instruments: Bartik shift–share for the one-year horizon together with the 12-month lag of the unemployment rate u^r_{t-12}. All specifications include two lags of regional inflation, the full set of survey-covariate controls, and the fixed effects indicated at the top of each column.

Table 23: 2SLS estimates: intermediate granularity with survey weights (one-year expectations)

Panel A: CPS shares				
	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	0.242 (0.040)***	0.247 (0.041)***	0.284 (0.092)***	0.313 (0.092)***
u^r	0.005 (0.032)	0.011 (0.041)	-0.007 (0.023)	-0.109 (0.042)***
R^2	0.936	0.936	0.958	0.958
Within R^2	0.936	0.936	0.591	0.570
N	1,639	1,639	1,639	1,639

Panel B: Michigan shares				
	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e	0.205 (0.029)***	0.213 (0.030)***	0.133 (0.048)***	0.136 (0.048)***
u^r	0.012 (0.033)	0.018 (0.042)	-0.016 (0.020)	-0.124 (0.035)***
R^2	0.937	0.937	0.962	0.964
Within R^2	0.937	0.936	0.632	0.624
N	1,639	1,639	1,639	1,639

Note: Driscoll–Kraay standard errors in parentheses. Both Michigan and CPS demographic groups use intermediate granularity (40 groups: age × education × gender, excluding marital status and children) with survey weights for Michigan expectations and CPS shares computed without weights. Instruments: Bartik shift–share for the one-year horizon shown together with the 12-month lag of unemployment u^r_{t-12}. All regressions include two lags of regional inflation, survey-covariate controls, and the fixed effects indicated above each column.

With the randomization, a new shift-share instrument is constructed that interacts the reassigned group expectations with the randomized CPS shares. The experiment does this 500 times, and reports the distribution of the first-stage F-statistics and second-stage coefficients.

Under the null hypothesis of valid identification, the placebo tests should exhibit:

- **Weak First Stage**: Randomized instruments should have little predictive power for expectations, yielding low F-statistics
- **Null Second Stage**: Placebo coefficients should cluster around zero, indicating no systematic relationship between artificial instruments and actual regional inflation

Figure 12 displays the distribution of first-stage F-statistics from 500 placebo iterations alongside the benchmark estimate. The placebo F-statistics cluster tightly in the range 0-2, while the benchmark F-statistic of approximately 15 lies far above this distribution. This demonstrates that instrument relevance depends critically on the specific economic relationships embedded in the demographic classifications and regional compositions.

Figure 13 shows the distribution of second-stage coefficients from the placebo tests. The placebo coefficients are tightly centered around zero, while the benchmark coefficient of approximately 0.60 falls outside this distribution. This indicates that the estimated effect reflects genuine economic relationships rather than mechanical correlations.

The placebo test results provide strong evidence for the validity of the shift-share identification strategy:

1. **Instrument Relevance**: The dramatic difference between benchmark and placebo F-statistics (15 vs. 0-2) confirms that the instrument's predictive power stems from genuine economic relationships.
2. **Exclusion Restriction**: The clustering of placebo coefficients around zero supports the interpretation that the instrument generates quasi-random variation in expectations across regions.
3. **Economic Significance**: The benchmark coefficient's position outside the placebo distribution indicates a meaningful economic relationship that disappears when the theoretically motivated links in the identification strategy are broken.

D Component Inflation Analysis with 5-10 Year Expectations

To examine the channels through which longer-term inflation expectations affect regional inflation, one can estimate the effect of 5-10 year expectations (pe5) on various components of regional inflation using the same 2SLS specification as the main analysis.

The component inflation analysis uses FRED regional CPI data for six inflation categories: commodities, non-durables, durables, services, services excluding rent, and services excluding medical care. The instrument is 5-10 year expectations with the Bartik shift-share instrument constructed using longer-term expectations.

Table 24 reports the second-stage coefficients. The results show that 5-10 year expectations have no significant effect on any component of regional inflation. All coefficients are small in magnitude (ranging from -0.064 to 0.028) and statistically insignificant, with standard errors that rule out economically meaningful effects.

Table 24: 2SLS by component inflation using pe5

Dependent Var.:	commod. infl.	non-dur. infl.	dur. infl.	serv. infl.	serv.-ren
pe5	-0.0284 (0.1142)	-0.0643 (0.1567)	0.0275 (0.0548)	0.0064 (0.0269)	0.0221 (0
Controls	Yes	Yes	Yes	Yes	Yes
Fixed-Effects:	——————-	——————-	——————	——————	——————
REGION	Yes	Yes	Yes	Yes	Yes
TIME	Yes	Yes	Yes	Yes	Yes
________________	__________________	__________________	_________________	_________________	________
S.E. type	Dris.-Kra. (L=4)	Dris.-Kra. (L=4)	Dri.-Kra. (L=4)	Dri.-Kra. (L=4)	Dris.-Kra
Observations	1,380	1,380	1,380	1,380	1,376
R2	0.95380	0.93524	0.98756	0.95441	0.93561
Within R2	0.38881	0.35222	0.65817	0.76324	0.56349

Note:
Reports panel regression results where the outcome variable is component inflation using pe5 (5-10 year expec
non-durable and durable goods. Services-housing and services-med. remove housing and medical services, res

These null results are consistent with the null results for longer-term expectations in regional inflation rates.

E Alternative timing of expectations

The benchmark model reflects the standard timing assumption in modern macroeconomic models that current outcomes are determined by contemporaneous beliefs about future payoff relevant aggregate outcomes. However, alternative approaches may break that simultaneity because of delays in decision-making, learning, or a process-oriented approach to decision-making and market-clearing as in **?**. To address concerns that equilibrium timing might suggest lagged expectations rather than contemporaneous beliefs, this Appendix re-estimates the baseline with $\pi_{r,t}$ on $\pi^e_{r,t-1}$, instrumenting $\pi^e_{r,t-1}$ with a lagged Bartik constructed from leave-one-out aggregate group expectations:

$$\pi_{r,t} = \alpha_r + \tau_t + \beta_\ell \pi^e_{r,t-1} + X_{r,t}\gamma + \varepsilon_{r,t}$$

Table 25 shows the first-stage results. First-stage strength is comparable to the baseline. Table 26 shows the second-stage results. The 2SLS pass-through remains of similar magnitude, with somewhat wider intervals with the predetermined CPS shares. These results provide evidence that the main result is not driven by the contemporaneous timing assumption.

Table 25: 2SLS: first stage (lagged expectations)

Panel A: CPS shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e_{t-1}	1.09 (0.05)	1.11 (0.07)	0.20 (0.05)	0.29 (0.07)
u^r	0.51 (0.06)	0.45 (0.04)	0.74 (0.09)	0.53 (0.04)
R^2	0.661	0.692	0.748	0.762
KP-F stat	479.7	533.7	8.5	12.5

Panel B: Michigan shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e_{t-1}	0.75 (0.04)	0.74 (0.04)	0.21 (0.02)	0.22 (0.02)
u^r	0.51 (0.06)	0.45 (0.04)	0.75 (0.09)	0.53 (0.04)
R^2	0.674	0.694	0.752	0.765
KP-F stat	530.3	545.5	19.9	21.0

Note: Driscoll–Kraay standard errors in parentheses. K̈P-Fïs the Kleibergen–Paap rank-Wald weak-instrument statistic for the first-stage regression of lagged one-year inflation expectations p^e_{t-1} on the lagged Bartik shift–share instrument. Instruments: Lagged (t-1) Bartik shift–share for the one-year horizon together with the 12-month lag of the unemployment rate u^r_{t-12}. All specifications include two lags of regional inflation, the full set of survey-covariate controls, and the fixed effects indicated at the top of each column. The endogenous variable is p^e_{t-1} instead of p^e_t to address simultaneity concerns.

Table 26: 2SLS estimates (lagged expectations)

Panel A: CPS shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e_{t-1}	0.523 (0.081)***	0.516 (0.080)***	0.621 (0.493)	0.458 (0.325)
u^r	0.089 (0.083)	0.118 (0.110)	0.034 (0.054)	-0.222 (0.092)**
R^2	0.787	0.786	0.898	0.918
Within R^2	0.787	0.785	0.067	0.204
N	1,527	1,527	1,527	1,527

Panel B: Michigan shares

	(1) No FE	(2) Region FE	(3) Time FE	(4) Two-way FE
π^e_{t-1}	0.473 (0.075)***	0.482 (0.076)***	0.512 (0.219)**	0.491 (0.190)***
u^r	0.099 (0.087)	0.125 (0.113)	0.027 (0.042)	-0.221 (0.090)**
R^2	0.790	0.787	0.909	0.915
Within R^2	0.790	0.786	0.169	0.179
N	1,527	1,527	1,527	1,527

Note: Driscoll–Kraay standard errors in parentheses. Instruments: Lagged (t-1) Bartik shift–share for the one-year horizon shown together with the 12-month lag of unemployment u^r_{t-12}. All regressions include two lags of regional inflation and the fixed effects indicated above each column. The endogenous variable is lagged one-year inflation expectations p^e_{t-1} instead of current expectations p^e_t to address simultaneity concerns.

References

Almas, I. (2012): "International income inequality: measuring PPP bias by estimating engel curves for food," *American Economic Review*, 102(2), 1093–1117.

Anderson, E., S. Rebelo, and A. Wong (2020): "Markups Across Space and Time," working paper.

Andrade, P., O. Coibion, E. Gautier, and Y. Gorodnichenko (2021): "No Firm is an Island? How Industry Conditions Shape Firms' Expectations," *Journal of Monetary Economics*, forthcoming.

Angelico, C., and F. Di Giacomo (2022): "Heterogeneity in inflation expectations and personal experience," working paper.

Bachmann, R., T. O. Berg, and E. R. Sims (2015): "Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence," *AMerican Economic Journal: Economic Policy*, 7(1), 1–35.

Bai, J. (2009): "Panel data models with interactive fixed effects," *Econometrica*, 77(4), 1229–1279.

Bartik, T. J. (1991): *Who Benefits from State and Local Economic Development Policies?* Upjohn Press.

Binder, C. (2015): "Whose Expectations Augment the Phillips Curve," *Economics Letters*, 136, 35–38.

Binder, C., R. Kamdar, and J. M. Ryngaert (2024): "Partisan Expectations and COVID-era Inflation," *Journal Of Monetary Economics*, 148, 1036–49.

Blanchard, O., and L. F. Katz (1992): "Regional Evolutions," *Brookings Papers on Economic Activity*, (1).

Boivin, J., M. Giannoni, and I. Mihov (2009): "Sticky prices and monetary policy: evidence from disaggregated us data," *American Economic Review*, 99(1), 350–384.

Branch, W. A. (2004): "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations," *Economic Journal*, 114, 592–621.

——— (2007): "Sticky Information and Model Uncertainty in Survey Data on Inflation Expectations," *Journal of Economic Dynamics and Control*, 31(1), 245–276.

Brock, W. A., and C. H. Hommes (1997): "A Rational Route to Randomness," *Econometrica*, 65, 1059–1095.

Bryan, M. F., and G. Venkatu (2001): "The Curiously Different Inflation Perspectives of Men and Women," *FRB Cleveland Economic Commentary*.

Burke, M. A., and A. K. Ozdagli (2022): "Household Inflation Expectations and Consumer Spending: Evidence from Panel Data," working paper.

Calvo, G. A. (1983): "Staggered Pricing in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12, 383–398.

Carroll, C. D. (2003): "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics*, 118(1), 269–298.

Coibion, O., and Y. Gorodnichenko (2012): "What Can Survey Forecasts Tell us About Information Rigidities," *Journal of Political Economy*, 120, 116–159.

——— (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105, 2644–2678.

Coibion, O., Y. Gorodnichenko, and T. Ropele (2020): "Inflation Expectations and Firm Decisions: New Causal Evidence," *Journal Of Monetary Economics*, 109, 122–140.

Cravino, J., T. Lan, and A. A. Levchenko (2020): "Price stickiness along the income distribution and the effects of monetary policy," *Journal of Monetary Economics*, 110(April), 19–32.

Crump, R., S. Eusepi, A. Tambalotti, and G. Topa (2021): "Subjecctive Intertemporal Substitution," *Journal of Monetary Economics, forthcoming*.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber (2021): "Exposure to Grocery Prices and Inflation Expectations," *Journal of Political Economy*, 129(5), 1615–1639.

D'Acunto, F., U. Malmendier, and M. Weber (2021): "Gender roles produce divergent economic expectations," *Proceedings of the National Academy of Sciences*, 118(21), 1–10.

Das, S., C. M. Kuhnen, and S. Nagel (2020): "Socioeconomic Status and Macroeconomic Expectations," *Review of Financial Studies*, 33(1), 395–432.

de Bruin, B., W. Van Der Klauuw, J. Downs, B. Fischhoff, G. Topa, and O. Armantier (2010): "2010," *Journal of Consumer Affairs*, 44(2), 381–402.

DellaVigna, S., and M. Gentzkow (2019): "Uniform Pricing in U.S. Retail Chains," *Quarterly Journal of Economics*, 134(4), 2011–2084.

Evans, G. W., and R. Gulamani (1984): "Tests for Rationality of Carlson-Parkin Inflation Expectations Data," *Oxford Bulletin of Economics and Statistics*, 46(1), 1–19.

Evans, G. W., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.

Fernandez-Val, I., and M. Weidner (2018): "Fixed effect estimation of large T panel data models," *Annual Reviews*, 10, 109–138.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020): "Bartik Instruments: What, When, Why and How," *American Economic Review*.

Harding, M., and C. Lamarche (2011): "Least squares estimation of a panel data model with multifactor error structure and endogenous covariates," *Economics Letters*, 111(3), 197–199.

Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson (2022): "The Slope of the Phillips Curve: Evidence from U.S. States," *Quarterly Journal of Economics*, 137(3), 1299–1344.

Hobijn, B., and D. Lagakos (2005): "Inflation Inequality in the United States," *The Review of Income and Wealth*, 51(4), 581–606.

Jorda, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95(1), 161–182.

Kaplan, G. W., and S. Schulhofer-Wohl (2017): "Inflation at the Household Level," *Journal of Monetary Economics*, 91, 19–38.

Kim, G., and C. Binder (2023): "Learning-through-Survey in Inflation Expectations," *American Economic Journal: Macroeconomics*, 15(2).

Koo, B., S. Lee, M. H. Seo, and M. Takano (2024): "What Impulse Response Do Instrumental Variables Identify?," working paper.

Mavroeidis, S., G. Chevillon, and M. Massmann (2009): "Inference in Models with Adaptive Learning," *Journal of Monetary Economics.*

Mavroeidis, S., M. Plagborg-Moller, and J. H. Stock (2014): "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, 52(1), 124–188.

Nakamura, E., and J. Steinsson (2014): "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review*, 104(3), 753–792.

Pesaran, M. H. (2006): "Estimation and infererence in large heterogeneous panels with a multifactor error structure," *Econometrica*, 74(4), 967–1012.

Tanaka, M., N. Bloom, J. David, and M. Koga (2020): "Firm Performance and Macro Forecast Accuracy," *Journal of Monetary Economics*, October.

Werning, I. (2022): "Expectations and the Rate of Inflation," working paper.

Woodford, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ.

——— (2013): "Macroeconomic Analysis without the Rational Expectations Hypothesis," in *Annual Review of Economics.*

Figure 11: Long-run inflation expectations: reduced-form evidence

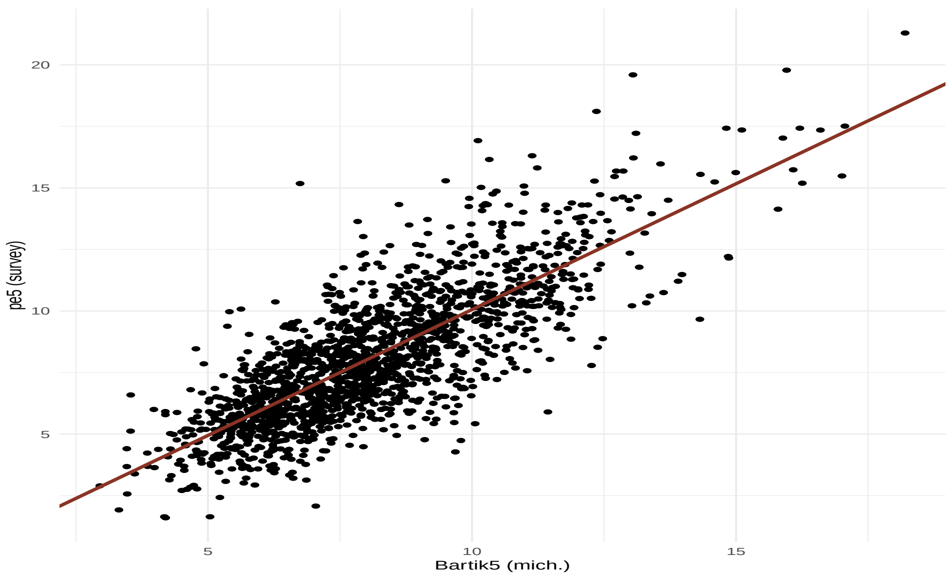

(a) First stage for five-to-ten-year expectations

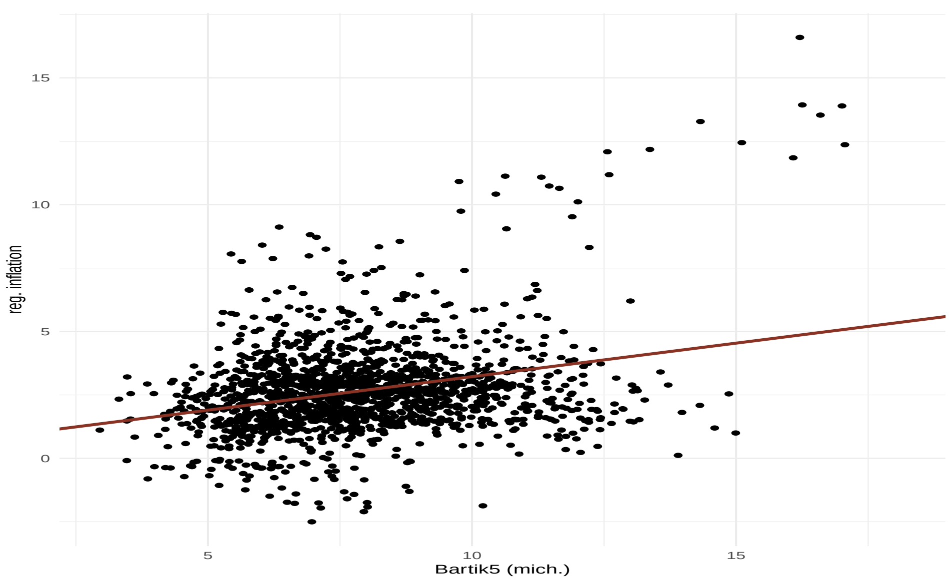

(b) Reduced form: inflation on long-horizon instrument

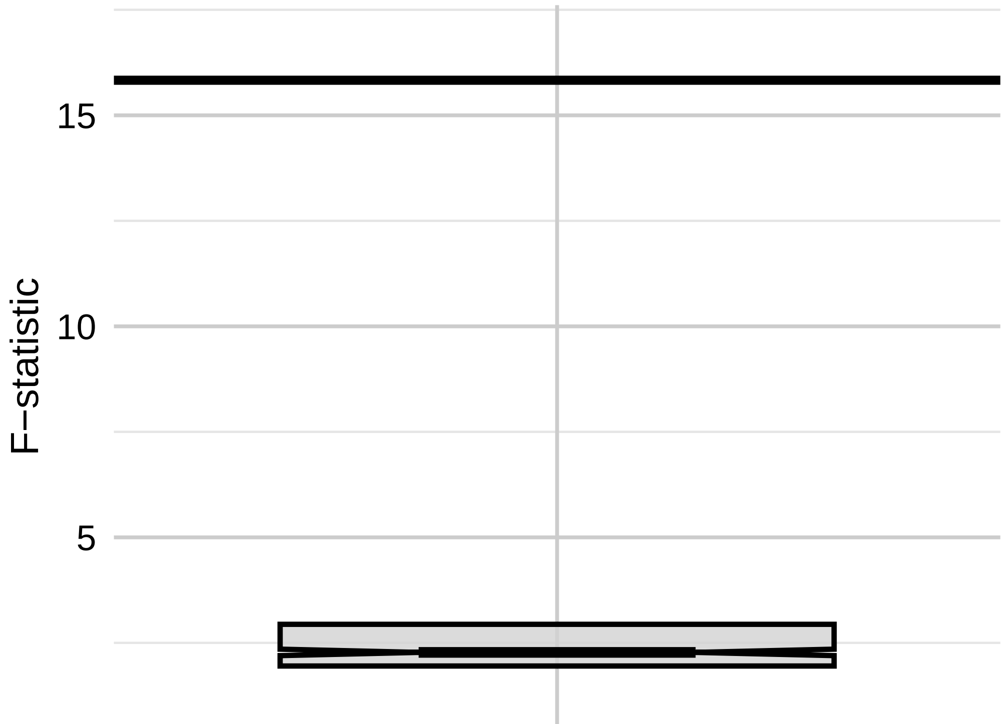

Figure 12: First-Stage F-Statistics: Placebo Distribution vs. Benchmark. The boxplot shows first-stage F-statistics from 500 placebo iterations with randomized industry assignments and CPS shares. The solid line shows the benchmark F-statistic.

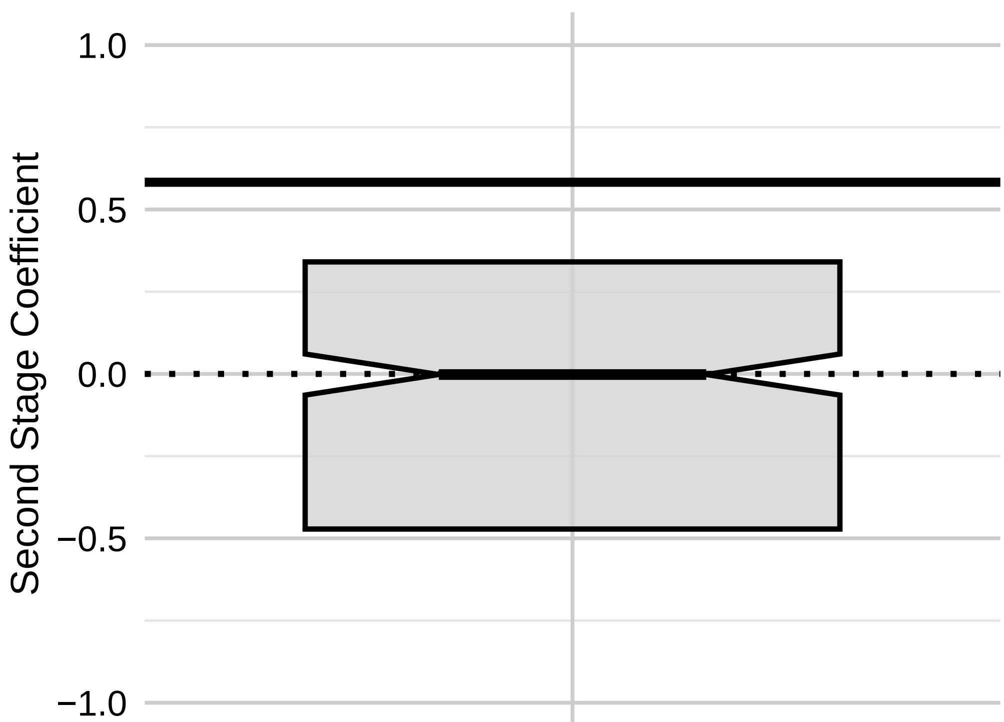

Figure 13: Second-Stage Coefficients: Placebo Distribution vs. Benchmark. The boxplot shows second-stage coefficients from 500 placebo iterations. The solid line shows the benchmark coefficient.