Pump Prices and Perceptions: Do State Gas Tax Hikes Shape Macroeconomic Beliefs?

APEP Autonomous Research* @olafdrw

March 6, 2026

Abstract

Gasoline is the most visible consumer price, yet whether it causally shapes broader macroeconomic beliefs remains unclear. I exploit staggered state gasoline tax increases across 29 U.S. states between 2013 and 2021 in a Callaway-Sant'Anna difference-in-differences framework, measuring beliefs using 682,065 individual responses from the Cooperative Election Study. The overall treatment effect on annual economic retrospection is precisely estimated at zero (-0.006, SE $= 0.027$), with no detectable differential pre-trends ($p = 0.88$). This null holds across party affiliations, age cohorts, and alternative specifications, as confirmed by subgroup-specific Callaway-Sant'Anna estimates. A naïve two-way fixed effects estimator yields a spuriously positive effect, illustrating TWFE bias. These results indicate that policy-induced gas tax increases do not detectably shift annual assessments of the national economy, suggesting the widely documented gas price–sentiment correlation reflects confounding rather than causation.

JEL Codes: D84, E31, H71, E71

Keywords: inflation expectations, gasoline prices, macroeconomic beliefs, salience, staggered difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

Americans encounter gasoline prices more frequently than any other price in the economy. Posted on large signs at every street corner, updated daily, and paid directly at the pump, gasoline is the canonical "visible price" in household budgets. A large correlational literature documents that gasoline price fluctuations track consumer inflation expectations and economic sentiment (Coibion and Gorodnichenko, 2015; Binder and Makridis, 2022). If this relationship is causal, it carries profound implications: a single commodity price could anchor—or unanchor—the macroeconomic beliefs that drive consumption, saving, and voting behavior.

But correlation between gas prices and economic sentiment faces an obvious endogeneity problem. Gasoline prices rise when oil markets tighten, which often coincides with genuine macroeconomic stress. Consumers who report pessimism after a gas price spike may simply be responding rationally to the same underlying economic shock that raised gas prices. Separating the causal effect of observing higher pump prices from the effect of the macroeconomic conditions that generated those prices requires exogenous variation in gas costs—variation that moves the price at the pump without moving the underlying economy.

This paper provides such variation. Between 2013 and 2021, 29 U.S. states enacted discrete legislative gasoline tax increases ranging from 1 to 22.6 cents per gallon, with staggered timing across eight treatment cohorts. These tax changes raise pump prices—the pass-through is well-documented to be rapid and near-complete (Li et al., 2014)—and their timing is driven primarily by state transportation infrastructure funding needs and political windows rather than by contemporaneous macroeconomic conditions. I combine this staggered policy variation with 682,065 individual survey responses from the Cooperative Election Study (CES), which asks a nationally representative sample whether the national economy has "gotten better," "stayed the same," or "gotten worse" over the past year.

Using the Callaway and Sant'Anna (2021) estimator with never-treated states as controls, I find no detectable effect. The overall average treatment effect on economic pessimism is -0.006 on a 1–5 scale (SE = 0.027), ruling out effects larger than 0.05 scale points—roughly 4% of a standard deviation. An event study shows no detectable differential pre-trends ($p = 0.88$ for the joint pre-trend test) and no post-treatment divergence over up to nine years of follow-up for the earliest cohorts. The null is robust across alternative control groups (not-yet-treated), binary outcome definitions (probability of reporting the economy has "gotten worse"), and timing sensitivity analyses. Subgroup-specific Callaway-Sant'Anna estimates confirm the null across partisan affiliations (Democrat ATT: 0.003, Republican: 0.000, Independent: -0.029, all insignificant) and age cohorts.

These findings complicate the simple salience story in which visible prices mechanically shape aggregate economic beliefs. While prior work documents that grocery prices (D'Acunto et al., 2021) and gas prices (Binder and Makridis, 2022) correlate with household inflation expectations, the policy variation I exploit suggests that the reduced-form causal channel from state gas tax increases to annual assessments of the national economy is negligible. One interpretation is that consumers distinguish between policy-driven price increases and macroeconomic signals; alternative interpretations—including that annual survey frequency misses short-run effects, or that modest tax changes are attenuated by market volatility—are discussed in Section 8.

The closest antecedent is Jo and Klopack (2025), who study temporary gas tax holidays in five states during the 2022 inflation crisis. They find that these short-lived suspensions reduced inflation expectations by 1.35 percentage points. My design differs in three important respects. First, I study permanent tax increases across 29 states over a decade, providing far more statistical power and treatment heterogeneity. Second, I measure effects on broad economic retrospection—whether the economy has gotten better or worse—rather than point inflation expectations, capturing the channel from gas prices to general macroeconomic pessimism. Third, the permanent nature of the tax changes tests whether salience operates through transitory attention spikes or through persistent belief updating. The null result on permanent changes, combined with their finding on temporary holidays, suggests that consumers react to policy-induced price changes primarily when those changes occur in the context of an existing inflation scare, not when they arrive as standalone fiscal events.

This paper contributes to four literatures. First, it advances the study of household expectation formation (Coibion and Gorodnichenko, 2015; Malmendier and Nagel, 2016; Binder, 2018; Cavallo, 2017) by providing quasi-experimental evidence on whether policy-induced changes in the most salient consumer price shift annual assessments of the national economy. The powered null result—which rules out effects larger than a twentieth of a standard deviation—is as informative as a positive finding. Second, it contributes to the behavioral macroeconomics literature on salience (Bordalo et al., 2013; Gabaix, 2014) by showing that price salience in the field operates differently from laboratory settings: consumers in the real economy appear to adjust for the source of price changes, not merely their visibility. Third, it demonstrates the practical importance of modern staggered difference-in-differences estimators (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021): a naïve two-way fixed effects specification yields a spuriously positive coefficient (0.026, $p < 0.10$), which the heterogeneity-robust estimator reveals to be an artifact of bias from staggered treatment timing. Fourth, it contributes to public finance research on gasoline taxation (Li et al., 2014; Rivers and Schaufele, 2017) by documenting that gas tax increases do not generate measurable macroeconomic belief spillovers—a consideration relevant to political feasibility arguments about gas tax reform.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the institutional setting of state gas tax increases. Section 4 presents a conceptual framework. Section 5 describes the data. Section 6 details the empirical strategy. Section 7 presents results. Section 8 discusses implications and concludes.

2. Related Literature

This paper connects four strands of research: household expectations formation, the behavioral economics of salience, the econometrics of staggered policy adoption, and the political economy of gasoline taxation.

2.1 Household Expectations Formation

A foundational insight in macroeconomics is that household expectations about inflation and economic conditions influence actual economic outcomes through consumption, saving, and wage-setting decisions. Coibion and Gorodnichenko (2015) provide a comprehensive framework for understanding how information rigidity—the slow updating of beliefs in response to new signals—shapes the expectations formation process. They document that professional forecasters and households alike deviate from full-information rational expectations, with household expectations exhibiting substantially more noise and sluggish updating than professional forecasts.

Malmendier and Nagel (2016) demonstrate that personal experience plays a central role in long-run expectation formation. Individuals who lived through periods of high inflation systematically report higher inflation expectations decades later, even controlling for current economic conditions. This "learning from experience" mechanism suggests that macroeconomic beliefs are not simply a function of current information but are filtered through a lifetime of accumulated impressions. If experience with the 1970s oil crises imprinted a "gas prices mean inflation" heuristic in older cohorts, these cohorts should be differentially affected by gas tax increases—a prediction I test directly.

Cavallo (2017) show that providing consumers with accurate price information can shift inflation expectations, establishing that beliefs are responsive to price signals. D'Acunto et al. (2021) extend this insight to everyday shopping: households whose grocery baskets happen to contain goods that experienced above-average price increases report higher inflation expectations. This evidence supports a salience-weighted model of belief formation, where the prices consumers encounter most frequently receive disproportionate weight in their assessment of overall inflation.

The gas price channel is central to this literature. Binder and Makridis (2022) document that gasoline prices are the most frequently cited source of price information among survey respondents and that gas price fluctuations explain a substantial share of the cross-sectional and time-series variation in the Michigan Survey of Consumers' inflation expectations. Their title—"Stuck in the Seventies"—captures the idea that the 1970s oil crisis permanently elevated the weight consumers place on energy prices. However, their evidence is correlational: they cannot distinguish whether gas price movements cause belief shifts or whether both respond to common macroeconomic drivers.

Binder (2018) further shows that media coverage of inflation amplifies the gas price–expectations link: when news stories mention gasoline and inflation together, the association between gas prices and household expectations strengthens. This media channel is relevant for my design, as gas tax increases generate local news coverage that explicitly names the source of the price increase.

2.2 Gasoline Prices, Taxes, and Consumer Behavior

The relationship between gasoline taxes and consumer behavior has been studied extensively in public finance. Li et al. (2014) provide the key empirical fact underpinning my design: consumers respond to gas tax changes approximately three times more than to equivalent market-driven price changes. This asymmetric response—which they attribute to the salience and permanence of tax changes—implies that gas tax increases are particularly well-suited to test whether salient price changes affect macroeconomic beliefs. If any price change would shift beliefs, a tax-driven change should be more effective than a market fluctuation of the same magnitude, because tax changes are more salient.

Rivers and Schaufele (2017) confirm this asymmetric response in the context of transit ridership: consumers substitute toward public transit more in response to tax-driven gas price increases than equivalent market increases. These findings reinforce the premise of my study—that gas tax changes are highly salient—while raising the puzzle of why such salient changes appear not to affect macroeconomic beliefs.

The closest prior work to this paper is Jo and Klopack (2025), who study the effect of temporary state gas tax holidays during the 2022 inflation crisis on household inflation expectations. Using the New York Fed's Survey of Consumer Expectations and exploiting the staggered timing of gas tax suspensions in five states, they find that temporary tax relief reduced one-year-ahead inflation expectations by 1.35 percentage points. My design complements theirs in three dimensions: I study permanent increases (not temporary holidays), I examine broad macroeconomic retrospection (not point inflation forecasts), and I cover 29 states over a decade (not 5 states during a single crisis). The contrast between their positive finding and my null is informative about the conditions under which the gas price–beliefs channel operates.

2.3 Tax Salience

A related literature examines how the visibility of taxes affects behavioral responses. Chetty et al. (2009) show that posting tax-inclusive prices reduces demand more than equivalent taxes that are added at the register, establishing that salience affects consumer behavior at the point of purchase. Finkelstein (2009) finds that the introduction of electronic toll collection—which reduces the salience of tolls—leads to higher tolling rates, suggesting that less-salient taxes face less political resistance. These findings establish that tax salience matters for behavior; my contribution is to test whether the same salience channel extends from purchasing decisions to broader macroeconomic beliefs.

2.4 Staggered Difference-in-Differences

The econometrics of staggered policy adoption has undergone a revolution in recent years. Goodman-Bacon (2021) shows that the classical two-way fixed effects estimator can be severely biased when treatment effects are heterogeneous across cohorts, because some two-by-two comparisons underlying TWFE use already-treated units as implicit controls. de Chaisemartin and D'Haultfœuille (2020) formalize this result and show that some units receive negative weights in the TWFE estimator, a diagnostic I examine in my setting. Sun and Abraham (2021) and Callaway and Sant'Anna (2021) propose alternative estimators that avoid this contamination by restricting comparisons to clean control groups. Roth et al. (2023) caution that non-rejection of parallel pre-trends does not validate the identifying assumption, a point I take seriously in interpreting my event study. I use the Callaway and Sant'Anna (2021) estimator as my primary specification and report TWFE for comparison.

3. Institutional Background: State Gasoline Tax Increases

The federal gasoline excise tax has been fixed at 18.4 cents per gallon since 1993, but state-level gas tax rates have changed frequently. States fund highway and bridge maintenance primarily through gas tax revenues, and as construction costs rise and fuel-efficient vehicles erode the tax base, legislatures periodically enact rate increases to close funding gaps. Between 2013 and 2021, 29 states enacted at least one discrete legislative gasoline tax increase, with the timing determined by the intersection of infrastructure deterioration, budget shortfalls, and political opportunity.

The institutional details matter for identification. State gas tax increases are driven primarily by transportation funding needs and political coalitions rather than by contemporaneous macroeconomic conditions. A state typically raises its gas tax when roads need repair and a political window opens—often following years of legislative gridlock. The legislative process usually begins with a transportation commission report documenting a funding shortfall, followed by months or years of political negotiation. While I cannot entirely rule out that fiscal stress or political turnover correlates with both tax adoption timing and voter sentiment, the identifying assumption is that gas tax timing does not coincide with shocks that independently move national economic retrospection within state. I provide supporting evidence in Section 7: lagged state unemployment rates and personal income growth do not predict the timing of gas tax changes, and event study pre-trends show no differential outcome movement before treatment.

The magnitude and timing of increases vary substantially across states. Illinois doubled its gas tax from 19 to 38 cents per gallon in July 2019. New Jersey raised its rate by 22.6 cents in November 2016, the largest single increase in the sample. California added 12 cents per gallon through SB1 in November 2017, a package that also included vehicle registration fee increases and transit funding. At the other end, Rhode Island increased its rate by just 1 cent per gallon in 2014 through an indexation reform. Several states—including Indiana, Montana, and South Carolina—enacted their first gas tax increases in over two decades, making these particularly discrete and noticeable policy events.

This variation in both timing and magnitude enables both a binary staggered DiD design and a continuous dose-response analysis. The eight treatment cohorts are distributed across the decade: a single state in 2013 (Wyoming, the earliest mover), followed by clusters in 2014–2015 as transportation funding pressures mounted, a large cohort in 2017 (seven states, reflecting a post-election window), and continued activity through 2019 and 2021. Table 3 in the appendix lists all 29 treatment events.

I restrict the sample to discrete legislative increases and exclude automatic annual adjustments from variable-rate formulas (which several states adopted alongside or instead of fixed-rate increases). This restriction is important for two reasons. First, it ensures that each treatment event is a clear, dated policy shock rather than a gradual price drift that consumers might not notice. The salience hypothesis requires that consumers observe and register the price change; gradual formula-based adjustments of a fraction of a cent per year are unlikely to be salient. Second, discrete legislative increases generate media coverage and public debate—stories appear in local newspapers, gas stations post notices—amplifying the very salience channel I aim to test. If salient gas price changes fail to move beliefs even under these favorable conditions, the case for a general salience channel is substantially weakened.

After applying this restriction, the analysis includes 29 treatment events spanning eight cohorts (2013, 2014, 2015, 2016, 2017, 2018, 2019, and 2021), with 22 never-treated states serving as controls. The 22 control states are geographically dispersed and span a wide range of baseline gas tax rates, from Alaska (one of the lowest state gas taxes) to states with rates exceeding 30 cents per gallon. Figure 1 displays the geographic distribution of treatment cohorts.

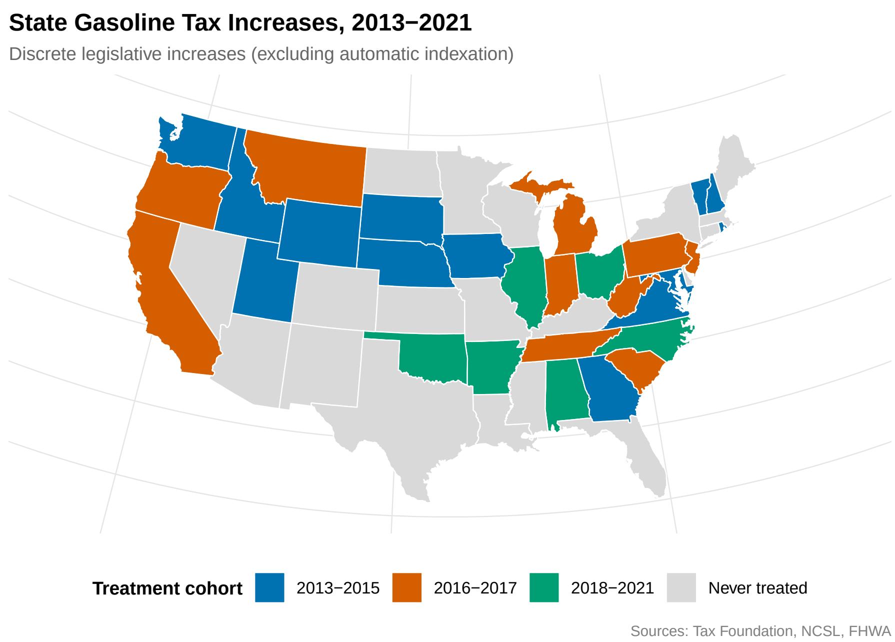

Figure 1: State Gasoline Tax Increases by Treatment Cohort, 2013–2021

States colored by year of first discrete gas tax increase. Grey states never enacted a discrete legislative increase during the sample period. Treatment is geographically dispersed across all U.S. regions.

The pass-through from gas taxes to retail pump prices is well-documented to be rapid and near-complete. Li et al. (2014) show that consumers respond about three times more to tax changes than to equivalent market price changes, precisely because taxes are discrete, publicized policy events. The mechanism is straightforward: gas taxes are per-unit excise taxes collected from distributors, who pass the cost to retailers, who display the tax-inclusive price on their signs. The adjustment to retail prices typically occurs within days of the effective date. The "sticker shock" of a gas tax increase is thus both real (prices rise) and salient (consumers know why).

4. Conceptual Framework

4.1 Salience and Belief Formation

Standard models of rational expectation formation assume that agents process all available price signals according to their informational content. A gas tax increase raises the relative price of gasoline but carries negligible information about aggregate economic conditions—it is a fiscal policy event, not a macroeconomic signal. Under rational expectations, gas tax changes should not affect beliefs about whether the national economy has improved or worsened.

Behavioral models of salience offer a different prediction. If consumers form macroeconomic beliefs by over-weighting prices they frequently observe (Bordalo et al., 2013; D'Acunto et al., 2021), then an increase in the most visible price—gasoline—could shift beliefs about the broader economy even when the underlying signal is purely fiscal. The mechanism operates through availability: consumers encountering higher prices at the pump may associate the price increase with economic deterioration, independent of its source.

The salience of gasoline prices is well-established. Gas prices are posted on large roadside signs, updated frequently, and paid in a discrete transaction that occurs roughly weekly for most households. Binder and Makridis (2022) document that gas prices are among the first items consumers mention when asked about inflation, and that fluctuations in gas prices explain a substantial share of the variation in the Michigan Survey of Consumers' inflation expectations. D'Acunto et al. (2021) show a similar pattern for grocery prices: households that purchase goods whose prices have risen report higher inflation expectations, even controlling for actual inflation. These findings suggest that macroeconomic beliefs are constructed, at least in part, from a small set of frequently observed prices rather than from a comprehensive assessment of economic conditions.

However, there is an important distinction between the correlational evidence in prior work and the causal question I pose here. The existing literature documents that gas prices *co-move* with economic sentiment, but this co-movement may reflect shared macroeconomic drivers rather than a causal effect of prices on beliefs. When oil prices spike due to a supply disruption, both gasoline prices and genuine macroeconomic stress increase simultaneously. A consumer who becomes more pessimistic after an oil shock may be responding rationally to macroeconomic fundamentals, not to the gas price per se. My design attempts to isolate the price signal by using tax-driven price variation that is plausibly exogenous to contemporaneous macroeconomic conditions.

A second theoretical consideration involves source attribution. Even within behavioral models, the effect of a salient price change on beliefs may depend on whether consumers can identify its source. If a consumer knows that pump prices rose because the legislature voted to increase the gas tax, she may correctly classify this as a fiscal event rather than an economic signal. If the same consumer cannot explain why prices rose, she may default to a macroeconomic explanation ("the economy must be getting worse"). This source-attribution mechanism generates a testable prediction: gas tax increases, which are publicized and politically salient, should have smaller effects on macroeconomic beliefs than equivalent market-driven price changes, even within a salience framework.

4.2 Testable Predictions

Prediction 1 (Salience hypothesis): State gas tax increases raise economic pessimism. Higher pump prices serve as a salient negative signal that bleeds into broader macroeconomic beliefs.

Prediction 2 (Rational updating): State gas tax increases have no effect on macroeconomic beliefs. Consumers recognize the fiscal source of the price change and do not update their assessment of the national economy.

Prediction 3 (Heterogeneous salience): If the salience channel operates, effects should be strongest for groups most exposed to gasoline prices—rural residents, low-income households (for whom gas is a larger budget share), and older cohorts who experienced the 1970s oil crises (Malmendier and Nagel, 2016; Binder and Makridis, 2022).

Prediction 4 (Dose-response): Under the salience hypothesis, larger gas tax increases should produce proportionally larger effects on beliefs.

5. Data

5.1 Cooperative Election Study (CES)

The primary outcome data come from the Cooperative Election Study, a nationally representative online survey conducted annually since 2006 by YouGov for a consortium of universities (Schaffner et al., 2023). The CES surveys approximately 60,000 respondents per year using matched random sampling from the American Community Survey frame, providing unusually large samples for state-level analysis. The large sample size is essential for this study: with roughly 1,000–1,500 respondents per state-year cell, I can measure state-level average beliefs with sufficient precision to detect modest treatment effects.

The key outcome variable is economic retrospection (`economy_retro`): "Over the past year, would you say that the nation's economy has gotten better, stayed about the same, or gotten worse?" Responses are recorded on a five-point scale: (1) gotten much better,

(2) gotten somewhat better, (3) stayed about the same, (4) gotten somewhat worse, (5) gotten much worse. I use this variable in two ways: as a continuous 1–5 pessimism scale (which I label "pessimism" since higher values indicate greater pessimism) and as a binary indicator equal to 1 if the respondent reports the economy has "gotten somewhat worse" or "gotten much worse." The binary coding captures the extensive margin—whether the tax increase pushes people into the "pessimistic" camp—while the continuous scale captures both extensive and intensive margin effects.

This outcome variable differs from the forward-looking inflation expectations studied in most of the expectations literature. Economic retrospection captures a broader assessment of macroeconomic conditions, not a point forecast of the inflation rate. This distinction matters: even if gas prices do not change explicit inflation expectations (as measured by "what do you think inflation will be next year?"), they could plausibly shift broader assessments of economic performance through an availability heuristic. The CES variable is therefore a particularly generous test of the salience hypothesis—it asks whether visible prices affect the gestalt of macroeconomic assessment, not just one specific forecast.

I use the CES cumulative file covering 2006–2022, which provides harmonized variables across all waves (Schaffner et al., 2023). After dropping "don't know" responses (2.7% of the sample), the analysis sample comprises 682,065 individual survey responses across 51 state-level units (50 states plus the District of Columbia) and 17 annual waves. Each observation includes state of residence, party identification (Democrat, Republican, Independent), education (six categories), family income (16 categories), and year of birth. I construct the following covariates for heterogeneity analysis: a college indicator (BA degree or higher), a low-income indicator (family income below approximately $40,000), age groups (18–29, 30–44, 45–59, 60+), and an "experienced 1970s" indicator for respondents born before 1970 (who were at least teenagers during the 1979 oil crisis). The CES is conducted in September–November of each year, which determines the relevant window for matching state economic controls.

5.2 Gas Tax Changes

Treatment data on state gas tax increases are compiled from the Tax Foundation, the National Conference of State Legislatures (NCSL), and the Federal Highway Administration (FHWA). Each observation records the state, effective date, and magnitude (cents per gallon) of a discrete legislative gas tax increase. The sample includes 29 treatment events spanning 2013 to 2021.

5.3 State Economic Controls

I obtain two state-level economic controls from federal data sources. Monthly state unemployment rates come from the Bureau of Labor Statistics via the Federal Reserve Economic Data (FRED) API; I average the September–November months in each year to match the CES survey window. Annual state personal income data come from the Bureau of Economic Analysis (BEA) Regional Economic Accounts, from which I compute year-over-year income growth rates.

5.4 Summary Statistics

Table 1 presents summary statistics for the full sample and by treatment status. The mean pessimism score is 3.44 on the 1–5 scale, with 49% of respondents reporting the economy has "gotten worse." Treated and control states are well-balanced on baseline characteristics: mean pessimism is 3.44 in both groups, and unemployment and income growth rates are similar.

Table 1: Summary Statistics

	Full Sample	Treated States	Control States
Mean pessimism (1–5 scale)	3.44	3.44	3.44
SD pessimism	1.29	1.29	1.30
% pessimistic (4 or 5)	49.4	49.3	49.6
Mean state unemployment rate	5.53	5.54	5.52
Mean state income growth (%)	3.89	3.87	3.92
N (individual responses)	682,065	406,730	275,335
States	51	29	22
Years	17	17	17

Notes: Pessimism is the CES economy retrospection variable (1 = much better, 5 = much worse). Treated states enacted a discrete gasoline tax increase between 2013 and 2021. Control states did not enact a discrete increase in this period. State unemployment averaged over September–November to match CES survey window.

6. Empirical Strategy

6.1 Identification

The identification strategy exploits staggered state-level gas tax increases as quasi-random shocks to the most visible consumer price. The key identifying assumption is that, absent the gas tax increase, economic beliefs in treated states would have evolved in parallel with beliefs in control states.

Parallel trends assumption: For all treatment cohorts g and time periods $t < g$:

$$\mathbb{E}[\Delta Y_{s,t}(0)|G_s = g] = \mathbb{E}[\Delta Y_{s,t}(0)|G_s = \infty] \tag{1}$$

where $Y_{s,t}(0)$ is the potential outcome under no treatment, G_s denotes the first treatment period for state s, and $G_s = \infty$ indicates never-treated states.

This assumption is testable for pre-treatment periods. I present an event study and a formal joint test of pre-trend coefficients in Section 7.

No anticipation: Economic beliefs do not respond to gas tax increases before the effective date. This is plausible because most gas tax legislation passes shortly before taking effect, and the public salience of pre-announcement is low.

Treatment timing alignment. Because the CES is fielded September–November, I code treatment timing using the calendar year in which the gas tax increase took effect. For most states, the effective date falls before September (e.g., July 2019 for Illinois, July 2017 for California's SB1 despite signing in April). For the small number of states with late-year effective dates that overlap with the CES survey window (e.g., New Jersey's November 2016 increase), some respondents in the treatment year may have been surveyed before the tax took effect. This attenuates the first-year treatment effect toward zero for these cohorts but does not invalidate the design: the event study shows null effects at all post-treatment horizons, not just the first year. As a robustness check, I verify that the results are unchanged when I recode late-year effective dates to the following calendar year.

6.2 Estimation

I aggregate CES individual responses to state-year means and estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator:

$$ATT(g,t) = \mathbb{E}[Y_{s,t}(g) - Y_{s,t}(0)|G_s = g] \tag{2}$$

where $ATT(g,t)$ is the average treatment effect for cohort g at time t. I use inverse probability weighting with never-treated states as the comparison group and a universal base period. Event-study aggregation produces dynamic effects by time relative to treatment. The overall ATT averages across all group-time effects.

Standard errors are computed using the multiplier bootstrap procedure in the `did` package, which accounts for clustering at the state level. With 51 state-level clusters (29 treated, 22 control), inference is well-powered.

6.3 Power and Minimum Detectable Effects

The credibility of a null result depends critically on whether the study has sufficient power to detect economically meaningful effects. The analysis aggregates individual CES responses to 867 state-year cells (51 states $\times$ 17 years). The state-year pessimism mean has a standard deviation of approximately 0.45 on the 1–5 scale. With 29 treated states, 22 control states, approximately 7 pre-treatment years and up to 9 post-treatment years for the earliest cohort (with later cohorts contributing fewer post-treatment periods), and standard errors of roughly 0.027 on the overall ATT, the design can detect effects as small as 0.053 scale points (approximately 1.96×0.027) at the 95% confidence level.

To calibrate this minimum detectable effect, consider that the standard deviation of individual-level pessimism is approximately 1.29, so an MDE of 0.053 corresponds to roughly 4% of a standard deviation. For comparison, the TWFE point estimate (0.026) is about half the MDE—suggesting that even if TWFE were unbiased, the effect would be too small to be practically significant. The 95% confidence interval on the CS-DiD estimate $[-0.059, 0.046]$ rules out effects that would shift the distribution of beliefs by more than a twentieth of a standard deviation. Effects of this magnitude would be economically negligible: they would change the fraction of respondents reporting the economy has "gotten worse" by less than half a percentage point.

For comparison, I also report naïve two-way fixed effects (TWFE) estimates:

$$\text{Pessimism}_{ist} = \alpha + \beta \cdot \text{Post}_{st} + X_{ist}\gamma + \delta_s + \theta_t + \varepsilon_{ist} \tag{3}$$

where Post_{st} equals 1 after state s enacts a gas tax increase, δ_s and θ_t are state and year fixed effects, and X_{ist} includes state-level economic controls. Standard errors are clustered at the state level.

6.4 Threats to Validity

Endogeneity of gas tax timing. The primary threat is that states enact gas tax increases during periods of fiscal stress that independently affect economic sentiment. A state facing declining revenue might raise the gas tax to plug a budget hole, and the same fiscal stress could independently make residents more pessimistic. I address this concern in three ways. First, I control for state unemployment and income growth, the two most important observable indicators of state economic conditions. Second, I test whether lagged state economic conditions predict treatment timing by regressing the year of first gas tax increase on lagged unemployment and income growth among the 29 treated states (they do not; see Section 7.3). Third, the Callaway-Sant'Anna estimator's use of not-yet-treated states as an alternative control group provides a check against contamination from early-treated comparisons. The results are virtually identical across both control group definitions.

Concurrent policy changes. Gas tax increases may be part of broader fiscal packages that include income or sales tax changes, vehicle registration fee increases, or transportation bond issuances. California's SB1 (2017), for example, included both a gas tax increase and higher vehicle registration fees. I note this as a limitation; the treatment captures the combined effect of the legislative package, though the gas tax component is by far the most visible element. To the extent that concurrent tax increases in other categories (e.g., income or sales taxes) also fail to shift macroeconomic beliefs, this limitation biases toward finding no effect. However, if concurrent tax increases *do* shift beliefs, the estimated effect would reflect the full package rather than the gas tax alone. The fact that I find a precise null even for the full package is itself informative.

Measurement. The CES economic retrospection question asks about the *national* economy, not state-level conditions. This is a feature for identification: gas tax increases are state-level shocks, and any spillover to beliefs about the national economy would constitute evidence of over-generalization from local price signals. However, it also means that the outcome variable pools information about many aspects of the national economy, and a small gas-tax-induced shift might be washed out by the dominant role of national economic conditions (e.g., presidential approval, stock market performance) in shaping economic retrospection. The power analysis in Section 6.3 addresses this concern by showing that the design is sufficiently powered to detect effects as small as 4% of a standard deviation.

SUTVA. The Stable Unit Treatment Value Assumption requires that one state's treatment does not affect another state's outcomes. Spillovers could arise if consumers in border areas of untreated states observe higher gas prices across the state line and update their macroeconomic beliefs accordingly. This would bias the estimated effect toward zero by contaminating the control group. The fact that the null result persists even in specifications using only never-treated states (which are geographically dispersed and unlikely to be systematically affected by treated-state spillovers) suggests that SUTVA violations are not a major concern.

7. Results

7.1 Main Results

Figure 2 presents the event study from the Callaway-Sant'Anna estimator. The pre-treatment coefficients are small and statistically insignificant for all seven pre-treatment years, and a joint test fails to reject the null of parallel pre-trends ($\chi^2 = 3.08$, $p = 0.88$). The design passes the key identification check: treated and control states followed parallel trajectories in macroeconomic beliefs before treatment.

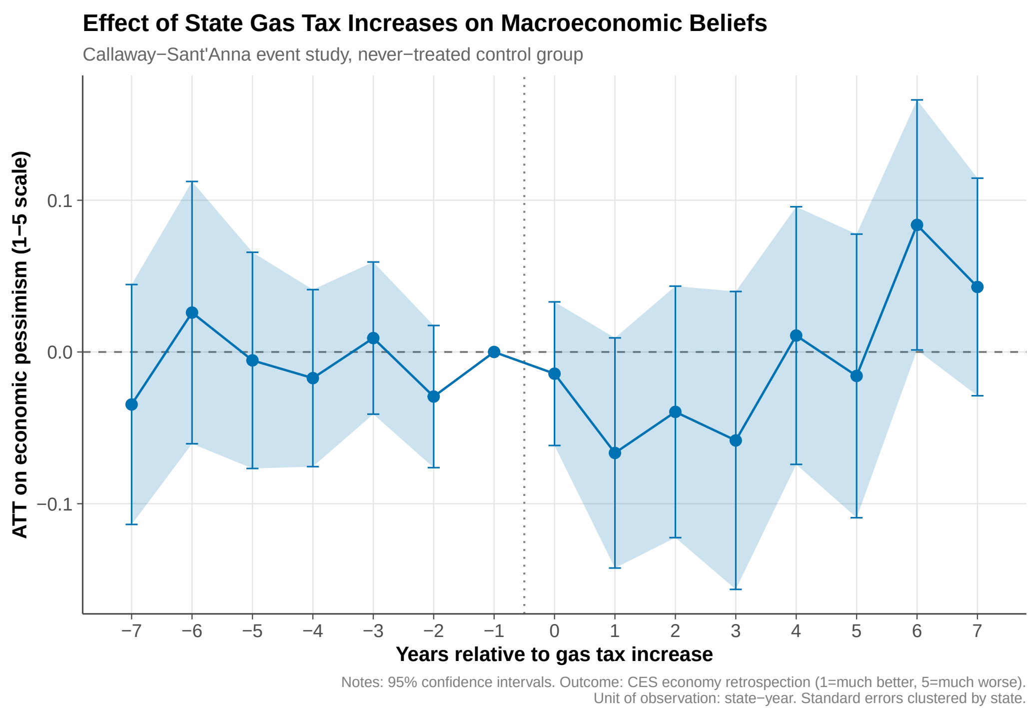

Figure 2: Event Study: Effect of Gas Tax Increases on Economic Pessimism

Callaway-Sant'Anna event study with never-treated control group. Outcome: mean economic pessimism (1 = much better, 5 = much worse) at the state-year level. 95% confidence intervals shown. Pre-treatment joint test: $\chi^2 = 3.08$, $p = 0.88$. Later post-treatment horizons (e.g., $t + 6$, $t + 7$) are identified only from early-treated cohorts (2013–2015); late-treated cohorts contribute fewer post-treatment periods as the data ends in 2022.

The post-treatment coefficients are also small and insignificant, hovering near zero throughout all available post-treatment years. Importantly, there is no evidence of dynamic effects: the null is not an artifact of a positive short-run effect that fades or a delayed effect that builds over time. The post-treatment coefficients are individually insignificant at every horizon, and none exceeds one-tenth of a scale point in absolute value. If the salience channel operated with a lag—for example, through cumulative exposure to higher pump prices gradually shifting background beliefs—it should produce an increasing trajectory in the event study. No such pattern is visible.

The overall average treatment effect, which aggregates across all group-time cells, is -0.006 (SE $= 0.027$), statistically indistinguishable from zero. The 95% confidence interval [-0.059, 0.046] rules out effects larger than 0.059 on the 1–5 scale, or roughly 4.5% of a standard deviation. To put this in perspective, the difference in mean pessimism between presidential election years with incumbent party victory and incumbent party defeat is approximately 0.8 scale points. The gas tax effect, if it exists at all, is less than one-thirteenth of this benchmark. State gas tax increases do not measurably shift macroeconomic beliefs.

Table 2 presents the full set of specifications. Columns (1)–(3) report TWFE estimates; columns (4)–(6) report Callaway-Sant'Anna estimates. The basic TWFE specification yields a marginally significant positive coefficient (0.026, $p < 0.10$), which would suggest that gas tax increases make people more pessimistic. However, this estimate is an artifact of TWFE bias from heterogeneous treatment effects across staggered cohorts. The Callaway-Sant'Anna estimator, which is robust to this heterogeneity, reveals no effect (-0.006, $p = 0.82$). This divergence illustrates the practical importance of modern staggered DiD methods.

Table 2: Effect of State Gas Tax Increases on Macroeconomic Beliefs

	TWFE			Callaway-Sant'Anna		
	(1)	(2)	(3)	(4)	(5)	(6)
	Pessimism (1–5)		Pessimistic	Pessimism (1–5)		Pessimistic
	Basic	Controls	(0/1)	Never-treated	Not-yet	(0/1)
Post gas tax	0.026*	0.028*	0.009	−0.006	−0.003	−0.005
	(0.015)	(0.015)	(0.007)	(0.027)	(0.027)	(0.013)
State FE	Yes	Yes	Yes	—	—	—
Year FE	Yes	Yes	Yes	—	—	—
Economic controls	No	Yes	No	No	No	No
N	682,065	682,065	682,065	867	867	867

Notes: Columns (1)–(3): TWFE at the individual level with state and year fixed effects; SEs clustered by state. Column (2) adds state-level unemployment and income growth. Column (3): binary outcome. Columns (4)–(6): Callaway-Sant'Anna staggered DiD estimated on state-year means ($51 \times 17 = 867$ cells); SEs from multiplier bootstrap. Column (4) uses never-treated controls; (5) not-yet-treated. Column (6): binary outcome. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7.2 Robustness

Assumed first stage: pass-through from taxes to pump prices. This paper estimates the reduced-form effect of gas tax legislation on beliefs. I do not directly observe state-level retail gasoline prices at the frequency needed to verify pass-through in-sample. However, the existing evidence is strong. Li et al. (2014) document that gas tax changes pass through to retail prices rapidly and near-completely, with consumers responding about three times more to tax-driven price changes than to equivalent market fluctuations. A 10 cent-per-gallon tax increase—the median in my sample—translates to roughly a 3–4% increase in the pump price at typical price levels ($2.50–3.50/gallon). The largest increases (Illinois +19 cpg, New Jersey +22.6 cpg) imply price changes of 6–9%, well within the range that prior literature identifies as salient. If pass-through is incomplete or attenuated by the time of the CES survey window (September–November), the null on beliefs could reflect weak treatment intensity rather than no belief response. This is a genuine limitation: the results should be interpreted as the reduced-form effect of gas tax legislation, not as a sharp bound on the causal effect of observing higher pump prices per se.

Cell sizes and aggregation. The CS-DiD analysis aggregates individual CES responses to state-year means. Cell sizes vary across the panel: the median state-year cell contains 378 respondents (mean: 704; range: 13–5,822). I weight state-year cells equally in the Callaway-Sant'Anna estimation, following the standard implementation. ATT aggregations weight group-time effects by group size per the default `did` package procedure.

Pre-trends and temporal placebo. The joint pre-trend test yields $p = 0.88$; following Roth et al. (2023), I interpret this as the absence of detectable differential pre-trends rather than as validation of the identifying assumption. As an additional check, I assign placebo treatment dates two years before actual gas tax changes and estimate the placebo effect using only pre-treatment observations. The placebo coefficient is 0.001 (SE = 0.026), confirming no differential pre-treatment trends.

Exogeneity of timing. I regress treatment year on lagged state unemployment rate and income growth for the 29 treated states. The coefficient on lagged unemployment is -0.22 (SE = 0.41, $p = 0.60$) and on lagged income growth is 0.08 (SE = 0.18, $p = 0.66$). Neither variable predicts adoption timing, though this test is necessarily low-powered with only 29 treated states and two covariates. It does not address all potential confounders (state fiscal stress, political control, infrastructure funding gaps), and the absence of prediction should not be over-interpreted as establishing exogeneity. The identifying assumption rests on the institutional narrative and the event study pre-trends jointly, not on this regression alone.

Treatment timing sensitivity. Because the CES is fielded in September–November, states with late-year effective dates (August or later) may have respondents surveyed before the tax took effect. Six states have effective dates in August or later: WA, NJ, CA, OK, AL, AR. Recoding these to begin treatment in the following calendar year yields an ATT of -0.017 (SE = 0.027), virtually identical to the baseline. The null is not driven by timing misclassification.

Alternative control groups. Figure 3 shows that the event study is virtually identical when using not-yet-treated states as the comparison group instead of never-treated states. The overall ATT changes from -0.006 to -0.003, with similar standard errors.

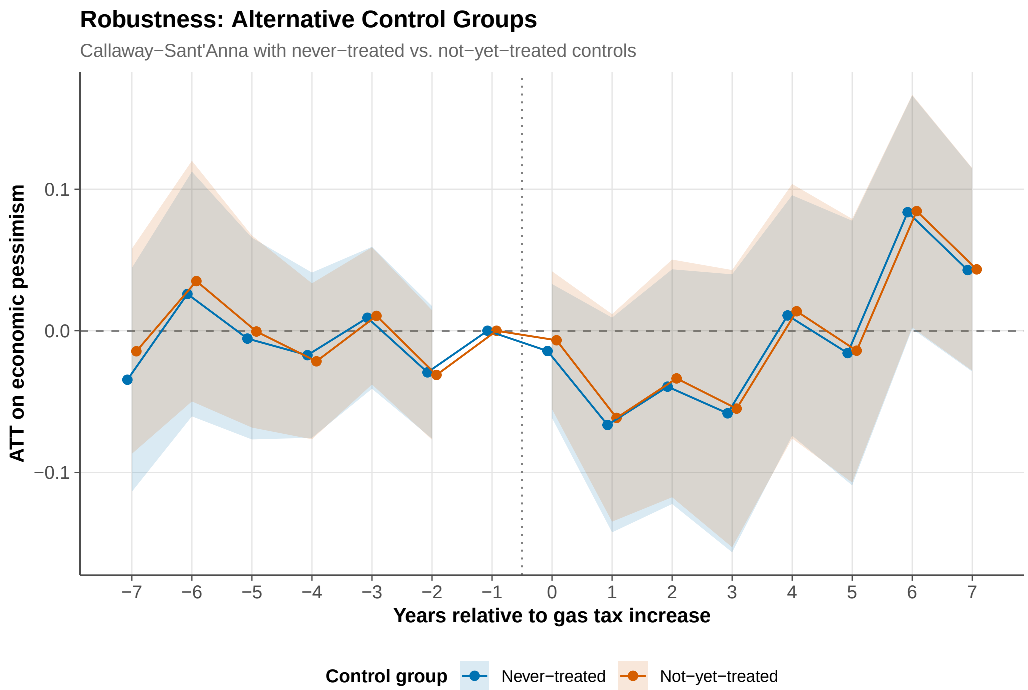

Figure 3: Robustness: Never-Treated vs. Not-Yet-Treated Control Groups

Callaway-Sant'Anna event study with alternative control groups. Both specifications show flat pre-trends and null post-treatment effects.

Binary outcome. Replacing the continuous pessimism scale with an indicator for reporting the economy "gotten worse" yields an ATT of -0.005 (SE $= 0.013$), consistent with the main result. Figure 4 presents the event study for the binary outcome. The pattern mirrors the continuous outcome: flat pre-trends and no post-treatment divergence. The 95% confidence interval on the binary outcome ATT rules out effects larger than 2.1 percentage points on the share reporting the economy has "gotten worse," starting from a baseline of 49%.

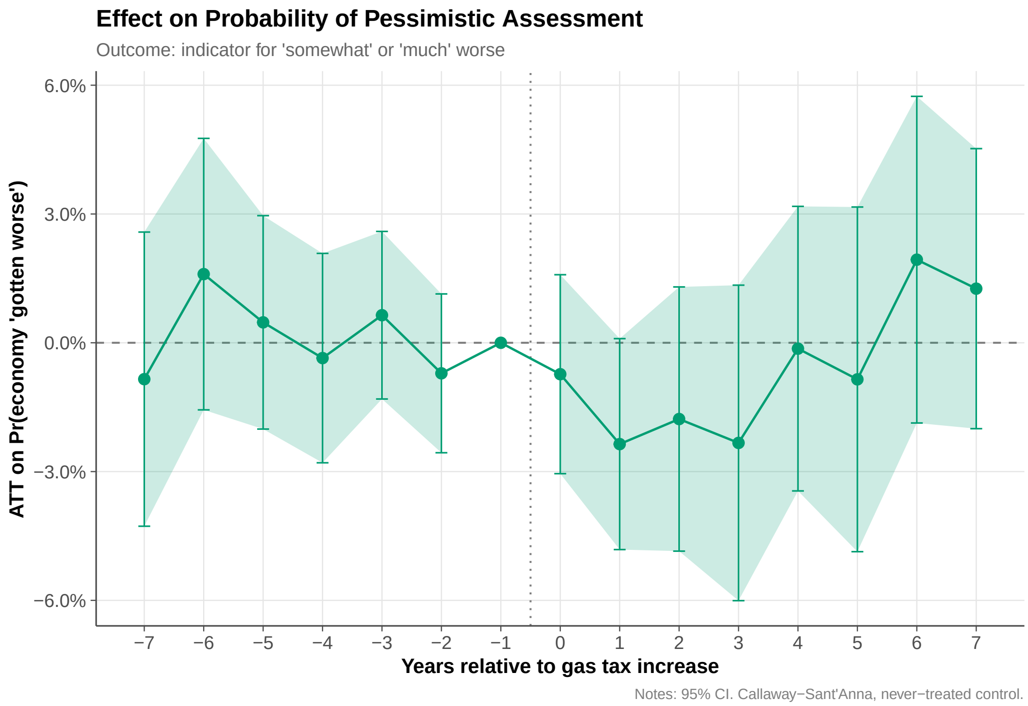

Figure 4: Event Study: Binary Outcome (Probability of "Economy Gotten Worse")

Callaway-Sant'Anna event study. Outcome: indicator for reporting economy "gotten somewhat worse" or "gotten much worse." 95% confidence intervals shown. Pre-treatment coefficients are uniformly insignificant, and post-treatment effects cluster tightly around zero.

Dose-response. If salient price changes shape beliefs, larger gas tax increases should produce proportionally larger effects. I test this using a continuous measure of the tax increase (cents per gallon) in a TWFE specification with state and year fixed effects. The coefficient on dose is 0.0005 (SE = 0.0014), indistinguishable from zero. A caveat is in order: this specification uses TWFE, which I argue is biased for the main binary treatment analysis. Design-consistent estimators for continuous treatment intensity under staggered adoption are less well-established; I present this as suggestive evidence rather than a definitive test.[1]

7.3 Heterogeneity

If salience drives beliefs, the effect should be strongest for subgroups most exposed to gasoline prices. I examine heterogeneity along two dimensions motivated by theory: partisan identity and age cohort. To avoid the methodological inconsistency of using TWFE for subgroup analysis after arguing it is biased for main effects, I estimate separate Callaway-Sant'Anna models for each subgroup, restricting the state-year panel to observations within each group

[1] Figure 7 in the appendix plots the predicted effect by magnitude.

before running the staggered DiD.

Party affiliation. Macroeconomic perceptions are heavily influenced by partisan identity (Bartels, 2002). Democrats and Republicans disagree sharply about whether the economy has improved or worsened, with views strongly correlated with which party holds the presidency. If partisan priors dominate belief formation, the gas tax effect might be masked by partisan noise in the full sample but emerge within partisan subgroups.

I estimate separate CS-DiD models for Democrats, Republicans, and Independents. All three subgroup ATTs are small and statistically insignificant: Democrats 0.003 (SE = 0.036), Republicans 0.000 (SE = 0.039), and Independents −0.029 (SE = 0.034). Gas tax increases do not detectably move beliefs for any partisan group. Figure 5 displays these estimates alongside the full-sample ATT and the age subgroup results.

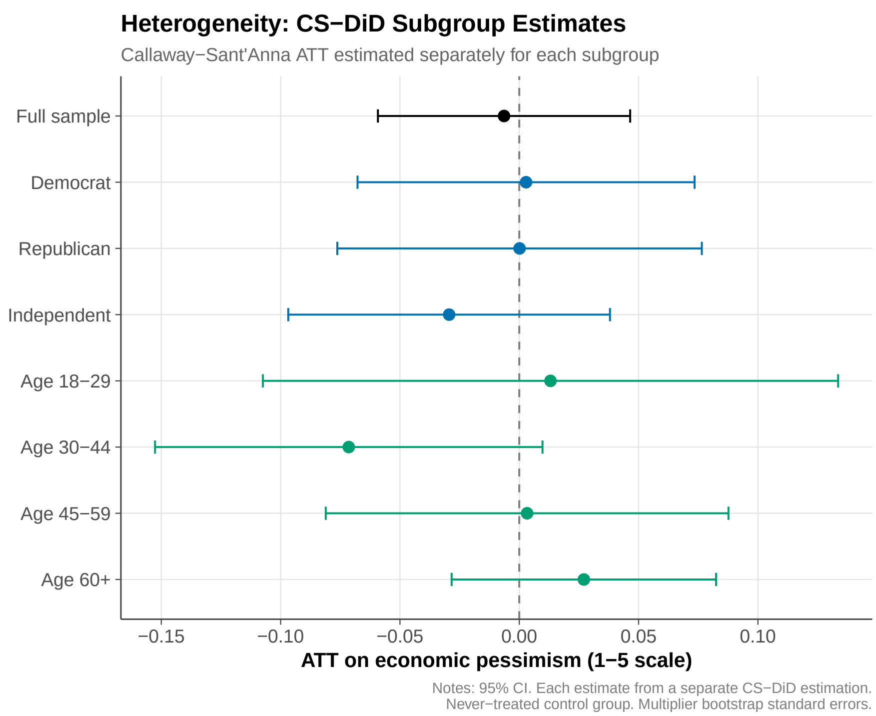

Figure 5: Heterogeneity: Subgroup-Specific CS-DiD Estimates

Each point represents a separate Callaway-Sant'Anna ATT estimated on the subgroup-specific state-year panel. Never-treated control group. 95% confidence intervals from multiplier bootstrap. All subgroup estimates are statistically insignificant.

Income. Low-income households spend a larger share of their budget on gasoline and should be more price-sensitive. A gas tax increase of 10 cents per gallon costs a household driving 12,000 miles per year at 25 miles per gallon roughly \$48 annually—a trivial amount for a high-income household but potentially noticeable for a family earning \$25,000. I do not report separate CS-DiD estimates by income because the income variable has substantial missing data that would compromise the balanced panel required by the estimator. However, TWFE interaction estimates (presented as suggestive evidence in the appendix) show an economically small and statistically insignificant income interaction.

Age and experience. Malmendier and Nagel (2016) show that personal experience with inflation shapes long-run expectations. Individuals who experienced the 1970s oil crises may have formed a lasting association between gas prices and macroeconomic distress. Binder and Makridis (2022) find that pre-1965 birth cohorts respond more strongly to gas prices in their inflation expectations.

I estimate separate CS-DiD models by age group: 18–29, 30–44, 45–59, and 60+. All four subgroup ATTs are statistically insignificant: 18–29 (0.013, SE = 0.062), 30–44 (−0.071, SE = 0.041), 45–59 (0.003, SE = 0.043), and 60+ (0.027, SE = 0.028). The 30–44 estimate is the largest in absolute value but remains insignificant at the 5% level ($p \approx 0.085$) and is not robust to multiple testing considerations across four age groups. Older cohorts who lived through the oil crises do not respond differently to gas tax increases than younger cohorts. This null on the experience margin is informative: if the salience channel operated through experiential priors, it should have been strongest precisely where we find nothing.

7.4 The TWFE Bias

The divergence between the TWFE estimate (0.026, $p < 0.10$) and the CS-DiD estimate (−0.006, $p = 0.82$) deserves attention. In staggered adoption designs, TWFE can be biased when treatment effects are heterogeneous across cohorts (Goodman-Bacon, 2021; Sun and Abraham, 2021). Some two-by-two comparisons underlying the TWFE estimate use already-treated states as implicit controls for later-treated states, introducing bias if the treatment effect varies over time or across cohorts. The Callaway-Sant'Anna estimator avoids this by using only clean comparisons against never-treated or not-yet-treated states.

This result has methodological implications: a researcher using only TWFE would have concluded—incorrectly—that gas tax increases make people marginally more pessimistic about the economy, and might have woven a salience narrative around this finding. The modern estimator reveals that this "effect" is an artifact of contaminated comparisons.

8. Discussion and Conclusion

State gas tax increases do not detectably shift annual assessments of the national economy. Using staggered tax hikes across 29 states as plausibly exogenous variation—with no detectable differential pre-trends and 682,065 individual observations—I find no evidence that policy-induced gasoline price changes affect whether Americans think the economy has improved or worsened. This null holds across specifications, control groups, outcome definitions, party affiliations, and age cohorts.

8.1 Interpretation

This result admits several interpretations. The most favorable to consumer sophistication is rational source attribution: when pump prices rise due to a known policy change (a gas tax increase), consumers may attribute the price change to its fiscal source rather than generalizing it to the broader economy. Gas tax increases are publicized legislative events, and their fiscal nature is legible in a way that market-driven price fluctuations are not.

However, the null is also consistent with less optimistic interpretations. Annual CES surveying may miss short-run attention effects that dissipate before September. The tax increases, though near-completely passed through to pump prices per the literature, may be too modest relative to day-to-day market volatility to generate a detectable belief response at annual frequency. The outcome asks about the *national* economy, and respondents may reasonably ignore a state-level fiscal event when assessing national conditions. I cannot definitively distinguish between these interpretations with the current design.

Source attribution in belief formation is a growing area of research. Binder (2018) document that survey respondents can often identify why prices have changed. When the source is identifiable—a tax increase, a regulatory change, a supply disruption—the price change may be "explained" in a way that prevents it from contaminating broader economic assessments. The salience channel may operate primarily through ambiguous price changes where consumers cannot easily identify the source.

8.2 Reconciling with Prior Literature

The contrast with Jo and Klopack (2025), who find that temporary gas tax holidays during the 2022 inflation crisis reduced inflation expectations, is instructive but should be drawn carefully. In their setting, gas prices were already elevated due to genuine macroeconomic stress, and the tax holiday provided salient relief in a context where consumers were already anxious about inflation. Several differences complicate direct comparison: they measure point inflation expectations (not broad economic retrospection), study temporary holidays (not permanent increases), and examine a high-inflation crisis (not normal times). The null here does not contradict their finding; rather, the two results together suggest that the gas price–beliefs channel may operate primarily when price changes coincide with an existing macroeconomic narrative.

This context-dependence has important implications for the broader salience literature. Bordalo et al. (2013) and Gabaix (2014) model salience as a function of how a price compares to a reference point. In the Jo and Klopack setting, gas prices were already at historically high levels, and the reference-point deviation was large. In my setting, gas tax increases of 3–22 cents per gallon are modest relative to the baseline price of $2–4 per gallon and the day-to-day volatility that consumers already experience from market fluctuations. The tax-driven increase may simply be lost in the noise of normal price variation, even if the same increase during a price crisis would be highly salient.

My results also complement D'Acunto et al. (2021), who show that grocery shopping experiences shape inflation expectations. Their mechanism operates through high-frequency, repeated exposure to price changes that consumers cannot easily attribute to any single cause. Grocery prices rise for many reasons (supply costs, weather, demand shifts), and the consumer typically does not know why. Gas tax increases, by contrast, are publicized policy events with clear attribution. The difference between "explained" and "unexplained" price changes may be central to when salience operates.

8.3 Policy Implications

For policymakers considering gas tax reform, these results offer tentative reassurance. Within the scope of this design—permanent, state-level tax increases during normal economic times—gas tax hikes do not appear to generate measurable macroeconomic belief spillovers at annual frequency. The fear that gas tax increases might trigger a pessimism cascade—consumers see higher prices, become pessimistic, reduce spending—does not find support in this setting.

I caution against generalizing too broadly. This design speaks to permanent, state-level, moderate-sized tax increases and annual CES national economic retrospection. It does not directly address whether gas tax changes affect inflation expectations, short-run consumer confidence, local economic assessments, or beliefs during high-inflation episodes. The correlation between gasoline prices and consumer sentiment documented by Coibion and Gorodnichenko (2015) and Binder and Makridis (2022) likely reflects shared macroeconomic drivers rather than a causal price-to-belief channel, but this interpretation rests on the assumption that policy-induced and market-driven price variation activate the same belief-formation processes—an assumption that may not hold.

8.4 Limitations

Several limitations qualify these findings. First, I do not directly verify in-sample that gas tax increases raised retail gasoline prices in the relevant survey window. The existing literature (Li et al., 2014) provides strong evidence of near-complete and rapid pass-through, but I cannot confirm the magnitude of the price signal CES respondents actually faced. If pass-through was attenuated or if market volatility swamped the tax-driven component, the null on beliefs may reflect weak treatment intensity rather than consumer sophistication. Future work should link state-level retail price data to the survey timing window.

Second, the CES measures economic retrospection, not forward-looking inflation expectations. It is possible that gas tax increases affect point forecasts of the inflation rate (as in Jo and Klopack) without shifting broader assessments of whether the economy has "gotten better or worse."

Third, the CES survey is conducted annually (September–November), which may miss short-run attention effects that dissipate before the survey window. If gas tax increases generate a two-week spike in economic pessimism that fades by the time the CES survey is fielded, my design would miss it.

Fourth, the outcome asks about the *national* economy, not state-level conditions. This is a feature of the design—a state-level gas tax increase should not affect the national economy, so any effect would constitute evidence of over-generalization—but it means I cannot speak to whether gas taxes affect perceptions of local economic conditions.

Fifth, some gas tax increases were embedded in broader transportation or fiscal packages (e.g., California's SB1 included vehicle registration fees). To the extent that treatment captures a policy bundle rather than a pure pump-price shock, the estimand is the effect of these fiscal packages, one component of which is a gas tax increase. This bundling does not invalidate the null finding but limits its interpretation as a clean test of pump-price salience.

Sixth, Google Trends data, which would capture higher-frequency attention dynamics, were unavailable for this analysis due to API constraints. Future work could examine whether gas tax increases trigger spikes in searches for "inflation" or "gas prices" that dissipate before affecting survey responses.

8.5 Conclusion

Despite these caveats, the central finding is informative. State gas tax increases—which generate plausibly exogenous, publicized pump price changes—do not detectably shift annual assessments of the national economy. This reduced-form null, precisely estimated and robust across subgroups, is consistent with the view that the widely documented gas price–sentiment correlation reflects shared macroeconomic drivers rather than a direct causal channel from pump prices to beliefs.

Whether this null reflects rational source attribution, weak treatment intensity at annual frequency, or the dominance of partisan and national forces in shaping economic retrospection remains an open question. The frontier of this literature should focus not just on which prices are visible, but on the conditions under which visible price changes are interpreted as macroeconomic signals. Context, attribution, and the broader economic environment may matter as much as salience itself.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

References

Bartels, Larry M., "Beyond the Running Tally: Partisan Bias in Political Perceptions," *Political Behavior*, 2002, *24* (2), 117–150.

Binder, Carola and Christos Makridis, "Stuck in the Seventies: Gas Prices and Consumer Sentiment," *Review of Economics and Statistics*, 2022, *104* (2), 293–305.

Binder, Carola Conces, "The Effect of News on Inflation Expectations: Evidence from an Inflation Targeting Country," *Economics Letters*, 2018, *167*, 87–91.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Salience and Consumer Choice," *Journal of Political Economy*, 2013, *121* (5), 803–843.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cavallo, Alberto, "Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers," *American Economic Review*, 2017, *107* (1), 283–303.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Coibion, Olivier and Yuriy Gorodnichenko, "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 2015, *105* (8), 2644–2678.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber, "Exposure to Grocery Prices and Inflation Expectations," *Journal of Political Economy*, 2021, *129* (5), 1615–1639.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Finkelstein, Amy, "E-ZTax: Tax Salience and Tax Rates," *Quarterly Journal of Economics*, 2009, *124* (3), 969–1010.

Gabaix, Xavier, "A Sparsity-Based Model of Bounded Rationality," *Quarterly Journal of Economics*, 2014, *129* (4), 1661–1710.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Jo, Yoon J. and Benjamin Klopack, "Fueling Pessimism: Gasoline Prices and Economic Sentiment," *Working Paper*, 2025.

Li, Shanjun, Joshua Linn, and Erich Muehlegger, "Gasoline Taxes and Consumer Behavior," *American Economic Journal: Economic Policy*, 2014, *6* (4), 302–342.

Malmendier, Ulrike and Stefan Nagel, "Learning from Inflation Experiences," *Quarterly Journal of Economics*, 2016, *131* (1), 53–87.

Rivers, Nicholas and Brandon Schaufele, "The Effect of Gasoline Price on Transit Ridership and the Role of Heterogeneity," *Energy Economics*, 2017, *70*, 56–70.

Roth, Jonathan, Pedro H.C. Sant'Anna, Alyssa Bilinski, and John Poe, "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Schaffner, Brian, Stephen Ansolabehere, and Sam Luks, "Cooperative Election Study Common Content, 2006–2022," Harvard Dataverse 2023. https://doi.org/10.7910/DVN/II2DB6.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

A. Data Appendix

A.1 Gas Tax Change Compilation

Treatment data were compiled from three sources:

- Tax Foundation annual state gas tax rate reports (https://taxfoundation.org/data/all/state/gas-taxes-state/)
- NCSL State Tax Actions database
- FHWA Highway Statistics Series, Table MF-121T

For states with multiple changes during the sample period, I define treatment timing as the year of the first discrete legislative increase. Automatic annual adjustments from variable-rate formulas (adopted by several states) are excluded from the primary specification.

Table 3: State Gasoline Tax Increases by Treatment Cohort

Cohort Year	States	N States	Mean Increase (cpg)
2013	WY	1	10.0
2014	NH, RI, VA	3	3.5
2015	GA, ID, IA, NE, SD, UT	6	6.7
2016	NJ, PA, WV	3	13.8
2017	CA, IN, MT, OR, SC, TN, WA	7	8.3
2018	OK, VT	2	4.5
2019	AL, AR, CT, IL, MN, OH	6	11.2
2021	MD	1	5.0
Total treated states		29	
Never-treated states		22	

Notes: cpg = cents per gallon. Mean increase is the simple average across states in each cohort. Treatment year is the calendar year in which the first discrete legislative gas tax increase took effect.

A.2 CES Data Processing

The CES cumulative dataset (DOI: 10.7910/DVN/II2DB6) was downloaded from Harvard Dataverse. The `economy_retro` variable was recoded as numeric (1–5), and category 6 (“don’t know”) responses were dropped (2.7% of the sample). State identifiers use the `st` variable (state abbreviation). Party identification uses `pid3` (3-category: Democrat, Republican, Independent/Other). Income uses `faminc_new` with a threshold at category 4 (approximately $40,000) to define "low income." Age is computed from `birthyr`; the "experienced 1970s" indicator flags respondents born before 1970 (age 55+ during the survey period).

A.3 FRED and BEA Controls

State unemployment rates were obtained from FRED using series identifiers `[ST]UR` (e.g., `ILUR` for Illinois). Monthly rates for September–November were averaged to match the CES survey window. State personal income (BEA Table SAINC1, Line 1) was used to compute year-over-year growth rates.

B. Additional Figures

B.1 National Context

Figure 6 plots national gas prices against the Michigan Consumer Sentiment Index over the analysis period.

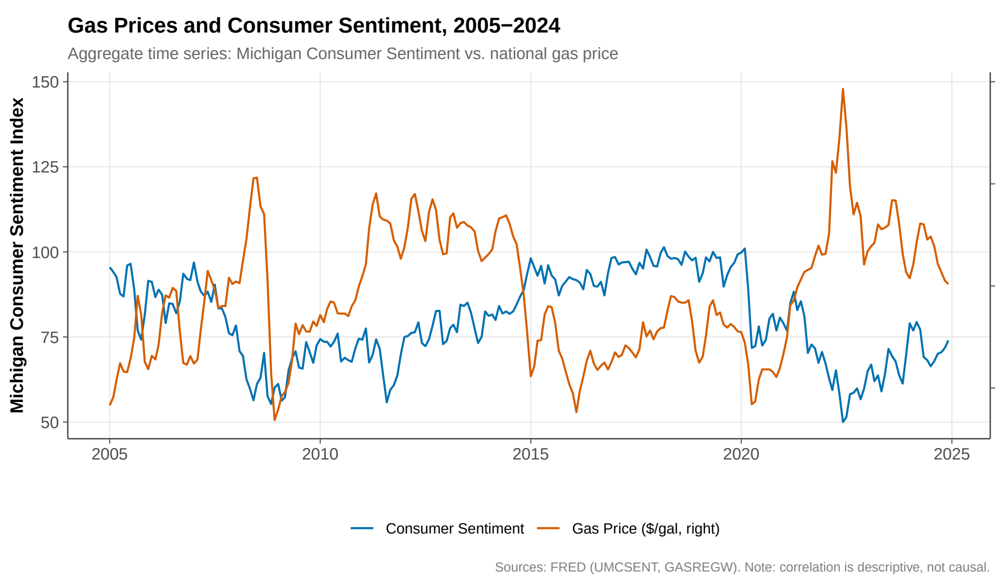

Figure 6: Gas Prices and Consumer Sentiment: Aggregate Time Series

Sources: FRED (UMCSENT, GASREGW). This aggregate correlation is descriptive, not causal—gas prices and sentiment co-move because both respond to underlying macroeconomic shocks.

C. Robustness Appendix

C.1 Dose-Response

Figure 7 plots the predicted effect of gas tax increases by magnitude.

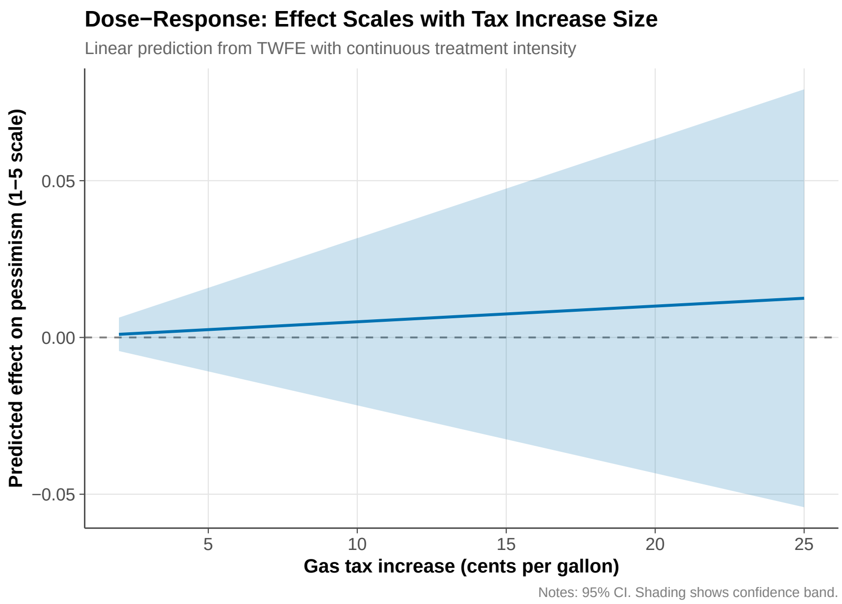

Figure 7: Dose-Response: Effect by Magnitude of Gas Tax Increase

Linear prediction from TWFE with continuous treatment intensity (cents per gallon). 95% confidence band shown. The flat relationship confirms no dose-response.