e61 Research Note | 1 July 2026

Who bears the burden of higher petrol prices?

Evidence on the elasticity of demand for petrol and its distributional implications.

Jack Buckley and Aditya Maitra

Summary

How does petrol consumption respond to price changes, and who bears the burden when prices rise? Using aggregated, consented and de-identified bank transaction data, we provide new evidence on both these questions. We find:

1. **Consumers are much more sensitive to petrol price changes than previously thought.** We estimate that a 10% increase in the price of petrol leads to a 3.8% decrease in consumption. This estimate of the elasticity of demand is three times as large as previous Australian evidence suggests, but consistent with more recent international research.
2. **Lower-income consumers are more responsive to price changes, but price shocks remain regressive.** Demand in the bottom income quintile is more than twice as elastic as the top (−0.62 vs −0.24), driven by stronger substitution effects. Yet price shocks remain regressive because petrol spending is a much larger share of income at the bottom. A 20% price rise equates to a 1.71% real income hit for the bottom income quintile compared to 0.47% for the top quintile.
3. **Consumers are less responsive to price changes in areas with greater motor vehicle reliance.** The elasticity of demand in urban areas (−0.47) is larger than in rural areas (−0.28). Within these urban areas, we find suggestive evidence that regions with lower motor vehicle reliance and greater public transport use have more elastic demand.

These findings have implications for policy, both in terms of how to best support households through global oil market volatility, and in the trade-offs for using price-based instruments to target the negative externalities of petrol use.

Implications for emissions reduction policies

- Our larger elasticity implies that price-based instruments targeting externalities associated with petrol use (fuel excise, carbon pricing) are more effective than previously thought, since a given price rise drives more behavioural change.
- However, our distributional findings show these instruments place a greater burden on low-income earners, so there may be a case for governments to pair them with compensatory transfers targeted at these households.

Household support around global oil price shocks

- Price rises hit low-income earners hardest, strengthening the case for crisis support. We estimate that the March 2026 price spike was equivalent to a 3.28% real income shock at the bottom of the distribution, 3.5 times larger than the 0.93% shock at the top.
- However, an excise cut is a poorly targeted response. Much of the outlay flows to high-income consumers who buy more petrol. It also carries a larger efficiency cost as consumers respond more to prices than previously thought.
- Two alternatives avoid both problems: direct income support through the transfer system or a lump-sum payment such as a vehicle registration fee cut. Both target exposed lower-income households more and leave the price signal intact.

The responsiveness of petrol consumption to price is an important input for policy. In times of crisis, it tells policy makers how readily households can adjust their behaviour to absorb a shock. It also helps clarify the trade-offs involved in policies that aim to bring down fuel prices, such as fuel excise cuts. In normal times, it helps policy makers understand the effect of changes in petrol tax rates and evaluate policies targeting the negative externalities associated with car use.

Despite the importance of this parameter, the most recent evidence on the price elasticity of demand for petrol in Australia is from almost 20 years ago. Using aggregate data from 1966 to 2006, Breunig and Gisz (2009) estimate the short-run price elasticity of fuel demand to be -0.13. This implies that a 10% rise in petrol prices leads to only a 1.3% fall in consumption.

This estimate may no longer reflect how Australians respond to petrol price changes today. Consumers now have more options to reduce their fuel consumption, through the rise of EVs, working from home, and the expansion of public transport networks. Each of these changes could make petrol demand more responsive to price.

The methods available to estimate the responsiveness of petrol demand have also improved. Recent overseas studies using more granular data have found substantially larger elasticities of demand, between −0.27 and −0.37.[1] These studies argue that highly aggregated data tends to bias elasticity estimates towards zero. This is because aggregation strips out shorter-run price variation that helps identify the demand elasticity, and makes demand shocks harder to control for.[2] This means older studies likely understate how much consumers actually adjust their behaviour when prices move.

This research note uses recent data from aggregated, consented and de-identified bank transaction records to provide an updated estimate of Australian households' elasticity of demand for petrol. The granularity of the data lets us strip out several sources of bias, including the aggregation bias likely present in earlier estimates.

The recent nature of the data, which covers the time period from 2020 to 2025, allows us to provide a more accurate picture of how Australian households currently respond to changes in price. We are also able to link petrol purchases to information on income and geographic characteristics, allowing us to study how the demand response varies across households, and what this implies for the distributional consequences of price increases. Appendix A provides further details on the data we use.

We identify the elasticity of demand by exploiting movements in global oil prices alongside account-level variation in petrol purchases. This allows us to isolate price movements driven by supply rather than demand (Box 1 explains why this is important). We do this in two steps. First, we include account and time fixed effects to absorb persistent differences in driving behaviour, seasonality and long run trends. Second, we instrument the retail price of petrol with the Terminal Gate Price (TGP) — the wholesale spot price. Our assumption is that after accounting for seasonality and time trends, TGP movements are uncorrelated with unobserved shocks to Australian households' demand for petrol. The fact that the period we analyse has large, global-conflict driven variation in oil prices strengthens the identification, but also is important to keep in mind when interpreting the external validity of our results. Appendix B sets out the framework in full.

Box 1: Estimating the elasticity of petrol demand

The core challenge with estimating demand elasticities is simultaneity bias. Price and quantity are jointly determined by demand and supply. When demand rises — say because more people start driving to work — both price and quantity increase together. The raw correlation between price and quantity therefore does not identify the true response of demand to price, because it also reflects shifts in demand itself.

Figure 1 illustrates this challenge. In panel A, demand shifts between periods (D1, D2, and D3), which means the observed equilibria trace the slope of the supply curve, not demand. Panel B shows what we want. When supply shifts and demand is held fixed, the equilibria trace the demand curve – giving the elasticity of demand, which is what we want to estimate.

Figure 1: Inference from observed prices and quantities

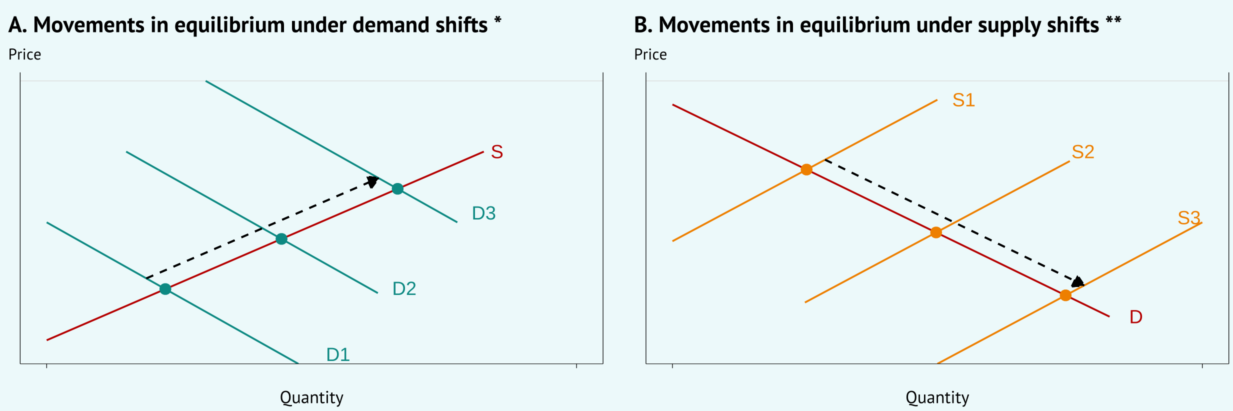

* Panel A simulates market equilibrium under shifts to the demand curve. When supply is fixed and demand is shifting, observed equilibriums trace out the supply curve, inducing a positive correlation between price and quantity.
** Panel B simulates market equilibrium under shifts to the supply curve. When demand is fixed and supply is shifting, observed equilibriums now trace out the demand curve.
Source: e61 Institute

1 For instance, Levin et al. (2017) obtain a range of estimates between −0.27 and −0.35, Knittel and Tanaka (2021) find an elasticity of −0.37, Kilian and Zhou (2024) estimate an elasticity of −0.31, and Coglianese et al. (2017) estimate the elasticity to be −0.37.
2 In Appendix E, we overview the existing literature on petrol demand elasticity estimation, including this discussion of aggregation bias, in more detail.

The clearest way to see how this strategy works in practice is through the events of 2022. Following Russia's invasion of Ukraine, global oil prices surged, then fell, then surged again, pulling wholesale costs and retail prices up and down with them. Figure 2 shows the path of prices and quantities throughout the year. Panel A shows the co-movement of wholesale costs and retail petrol prices. Panel B shows the quantity response to these price movements. Our analysis takes advantage of this kind of variation across the entire time period in our sample to estimate the elasticity of demand.[3]

Figure 2: Movements of costs, prices and quantities over 2022

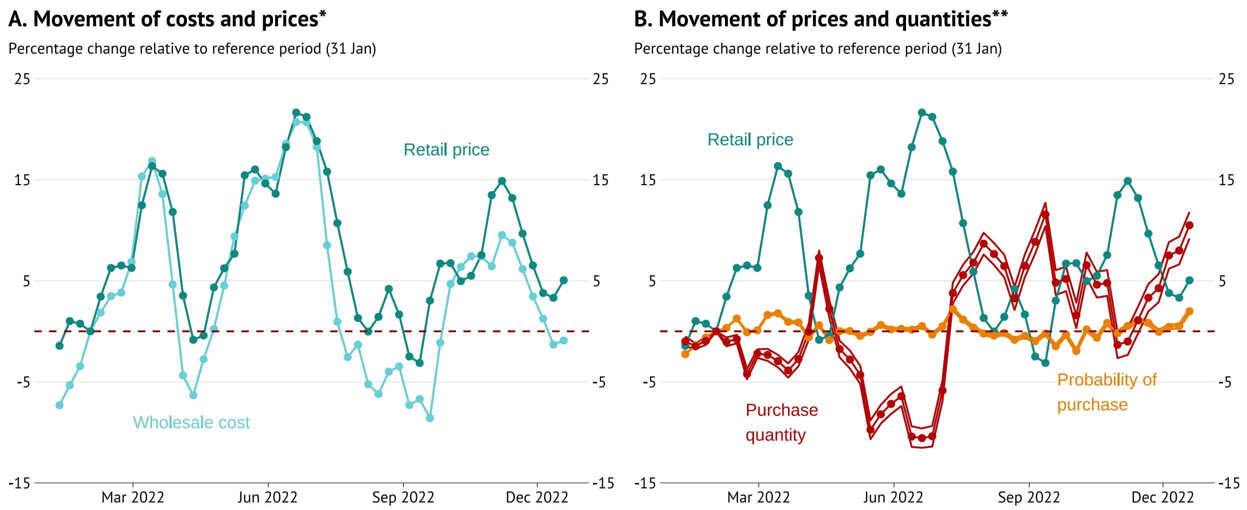

* Panel A plots the time-series movement in average wholesale costs (TGP) and average retail petrol prices across 2022. For each, we regress the log of the outcome variable (either cost or price) on weekly dummies, and plot these coefficients.
** Panel B plots the time-series movement in average retail petrol prices and measures of quantity across 2022. For the probability of purchase time-series we regress dummies for if the account makes a petrol purchase in the week on weekly dummies with account level fixed effects. For the purchase quantity series we regress the log of the conditional purchase quantity in litres on weekly dummies with account level fixed effects. For both, we plot the resulting coefficients on the weekly dummies as well as the 95% confidence intervals.
*** In both panels, the reference period is the week beginning on the 31st of January 2022. All coefficients should be interpreted as the change (or percentage change) in the outcome variable relative to this reference period.
Source: e61 Institute

How responsive is petrol demand to price movements?

We estimate the price elasticity of demand for petrol to be −0.38. This implies that a 10% increase in petrol prices leads to a 3.8% fall in demand.

Figure 3: Demand elasticity estimates

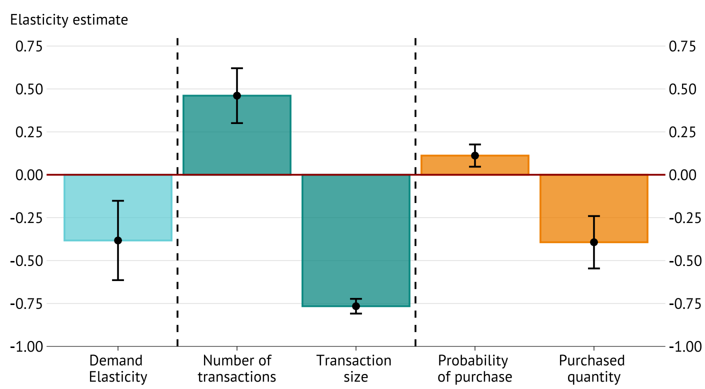

* Figure plots estimates of demand elasticities as well decompositions across behavioural margins. Point estimates and the 95% confidence interval are shown. Column 1 shows the baseline estimate for the elasticity of demand for petrol. Columns 2-4 show estimates of elasticites for the specified outcome variable. For instance, number of transactions gives the percentage change in the number of petrol transactions in a month caused by a percent increase in price.
Source: e61 Institute

This overall response masks (partially) offsetting movements along two margins: the extensive margin (how often consumers refill) and the intensive margin (how much they refill each time). We find that when prices rise, consumers respond by making

3 It is important to note that other sources of variation might produce different estimates. Our strategy relies on largely short-run changes in local prices driven by global supply shocks. Longer-run responses to more gradual or permanent shifts may differ.

smaller, more frequent purchases. A price increase of 10% is associated with a 7.7% decrease in the average size of each transaction, but a 4.6% increase in the number of monthly transactions (Figure 3).[4] When prices fall, the pattern reverses, with consumers making larger, less frequent transactions.[5]

This transaction size 'smoothing' behaviour is consistent with a number of potential explanations. One is that because petrol is storable, there is an opportunity cost of filling up. When prices spike, consumers might refill a smaller amount in the hope that prices will fall before their next purchase. Another explanation is that consumers may target expenditure amounts rather than litre amounts when filling up their tank.

This elasticity is much larger than earlier Australian estimates, but consistent with international research

Our elasticity estimate is substantially larger than the −0.13 short-run elasticity reported by Breunig and Gisz (2009), the most recent published Australian estimate. Appendix D reports a range of robustness checks that test whether the gap reflects specific choices in our approach, none of which substantially changes our results.

The gap instead likely reflects two factors. The first is structural. Breunig and Gisz (2009) use data from 1966 to 2006. Consumers today have more options to reduce their petrol consumption than they did two decades ago (let alone 60 years ago), with the rise of EVs, expanded public transport networks, and more flexible work arrangements. The second is methodological. Breunig and Gisz (2009) only had access to highly aggregated time-series data. Estimates using such data are susceptible to attenuation bias due to the challenge of controlling for shifts in demand (Levin et al., 2017). Our estimate sits closer to a range of recent overseas studies that find elasticities between -0.27 and -0.37 using more granular data (Coglianese et al., 2017; Kilian & Zhou, 2024; Knittel & Tanaka, 2021; Levin et al., 2017).

The difference between elasticity estimates is important for policy design

Applied to the peak of the price spike following the 2026 Iran shock, an elasticity of −0.13 implies only a 4.7% reduction in petrol demand. Our estimate implies a 13.7% reduction – a meaningfully different picture of the demand destruction caused by price increases (Figure 4 panel B).[6]

Figure 4: Movements of observed quantities and simulated quantities in 2026

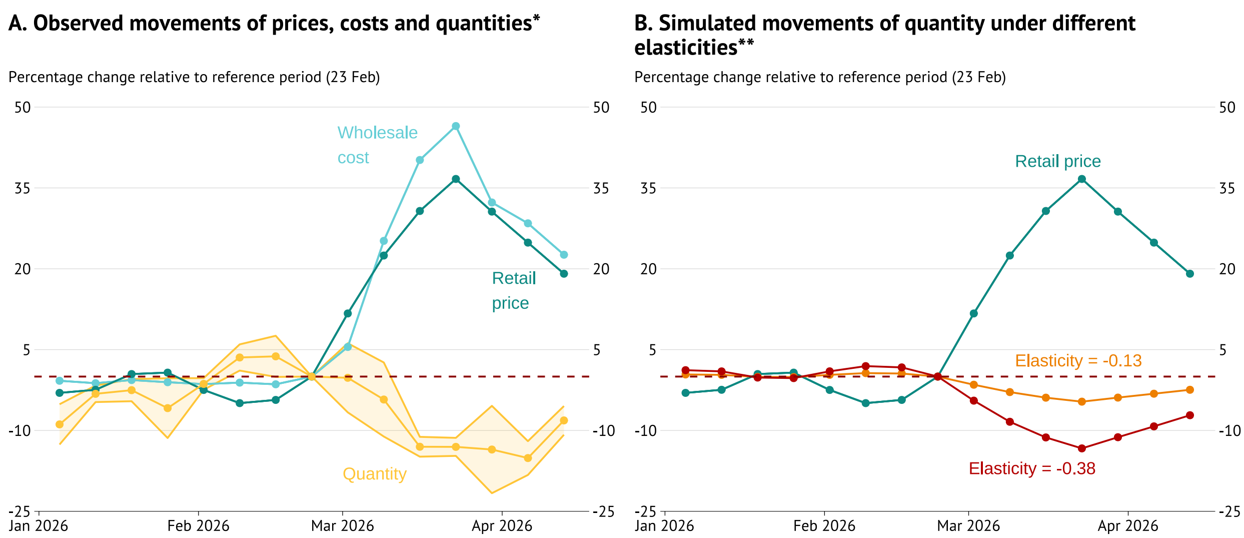

* Panel A plots the time-series movement in average wholesale costs (TGP) and average retail petrol prices across 2022. For each we regress the log of the outcome variable on weekly dummies and plot these coefficients. For the quantity series, we regress quantity in levels on weekly dummies with individual fixed effects. We then plot these weekly dummy coefficients, divided by the average purchase quantity, as well as the 95% confidence interval for these imputed coefficients. The reference period is the week begining on the 23rd of February 2026. All coefficients should be interpreted as the percentage change in the outcome variable relative to this reference period.
** Panel B plots the same time-series movement in retail price, as well as simulated movements in quantites under assumed elasticities of -0.13 and -0.39 respectively. Quantity movements are simulated by taking the percent change in price relative to the reference period and multiplying by the chosen elasticity value
Source: e61 Institute

4 Note that these extensive margin (0.46) and intensive margin (−0.77) elasticities do not add up to the aggregate (−0.38). We explain why in Appendix D.1.
5 This transaction size 'smoothing' implies that over any time-horizon (monthly in Figure 3), transaction probability correlates positively with price. So we see a small positive price effect on transaction probability, and a larger negative effect on quantity conditional on purchasing.
6 The price spike peaked in the week beginning 23 March 2026, when average prices were around 36% higher than four weeks earlier (Figure 4).

Notably, our estimate does not use data from 2026. Yet it appears close in magnitude to preliminary data on the observed drop in petrol quantities following the shock (Figure 4 panel A), adding further support to the credibility of our findings.

Our results also have policy relevance beyond this immediate context and global oil shocks in general. As Australia works towards decarbonising its transport use, an important question is how effectively price-based instruments, such as carbon pricing or road user charges, can drive reductions in petrol consumption. Our estimates suggest that these instruments are likely to be much more effective than previous evidence suggests.

What are the distributional implications of higher petrol prices?

The aggregate elasticity we report above captures the average demand response to price movements, which has implications for crisis support and tax design. But this average obscures heterogeneity that is likely also important. Policymakers may wish to know how the response varies by factors such as income and geography to better understand the distributional consequences of large price spikes.

Lower-income consumers are more exposed to price increases

Higher petrol prices are widely thought to be regressive. Petrol is a normal good, meaning that as incomes rise, individuals purchase and consume greater amounts (Figure 5 panel A; Maitra 2026).[7] However, as with other necessities, consumption does not scale proportionately with income.[8] Therefore the proportion of income spent on petrol (the budget share) is highest at the bottom of the income distribution (Figure 5 panel B).[9]

Figure 5: Petrol consumption and budget shares across the income distribution

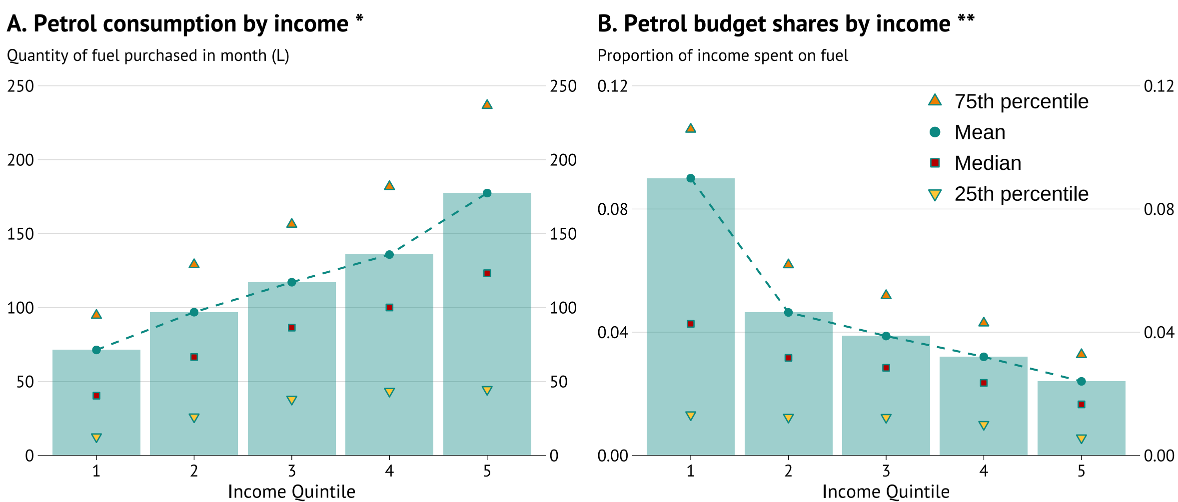

* Panel A shows average monthly consumption of petrol in litres across the income distribution.
** Panel B shows the average proportion of income spent on petrol across the income distribution.
Source: e61 Institute

The pattern in Figure 5 panel B implies that lower-income consumers are more exposed to increases in petrol prices. If all households were to respond to a given price increase in the same way, it will take a larger bite out of low-income households' budgets than high-income ones, because petrol spending makes up a larger share of their income.

In practice, whether price rises are regressive, and their degree of regressivity, will depend not only on this budget exposure, but also on the behavioural response of households. If lower-income households cut their consumption more sharply when

7 All the results on income presented in the body of the note condition on an account purchasing petrol at least once (i.e. we condition on owning a motor vehicle that requires petrol). In Appendix D.2, we discuss how our distributional results are affected by differences in this extensive margin (buying petrol at all vs not) across the income distribution.
8 A similar pattern is often observed with consumption and spending on items such as groceries and electricity.
9 In this section we display our results by income quintiles. In our sample, the lowest quintile has real income of less than (roughly) $30k per year. The other quintiles respectively have (approximately) incomes of $30k-$50k, $50k-$70k, $70k-$100k, and $100k +. All numbers are expressed in 2026 dollars. In Appendix D.2 we present the results in terms of income 'buckets' for greater interpretability.

prices rise, the burden on them will be reduced.[10] Thus, the overall welfare effect of higher petrol prices depends on the combination of budget exposure and the demand response across the income distribution.

Petrol demand is most sensitive at the bottom of the income distribution

Which way the demand response varies with income is not obvious in theory. Higher-income households may find it easier to substitute away from petrol consumption if they can work from home, live closer to public transport, or engage in more discretionary driving. Lower-income consumers may face more of a financial incentive to reduce fuel consumption when prices spike because petrol accounts for a higher budget share at the bottom of the income distribution.

To answer this question empirically, we re-estimate the price elasticity of demand separately for each income quintile.[11] We find that the elasticity of demand for petrol is decreasing across the income distribution. Households in the bottom income quintile are more than twice as responsive to price movements as those in the top income quintile (elasticity of −0.62 compared to −0.24; Figure 6, panel A). The implication of this is that a given increase in price induces more demand destruction (a larger fall in the quantity of petrol consumed) at the bottom of the income distribution than at the top.

Figure 6: Differences in demand response by income

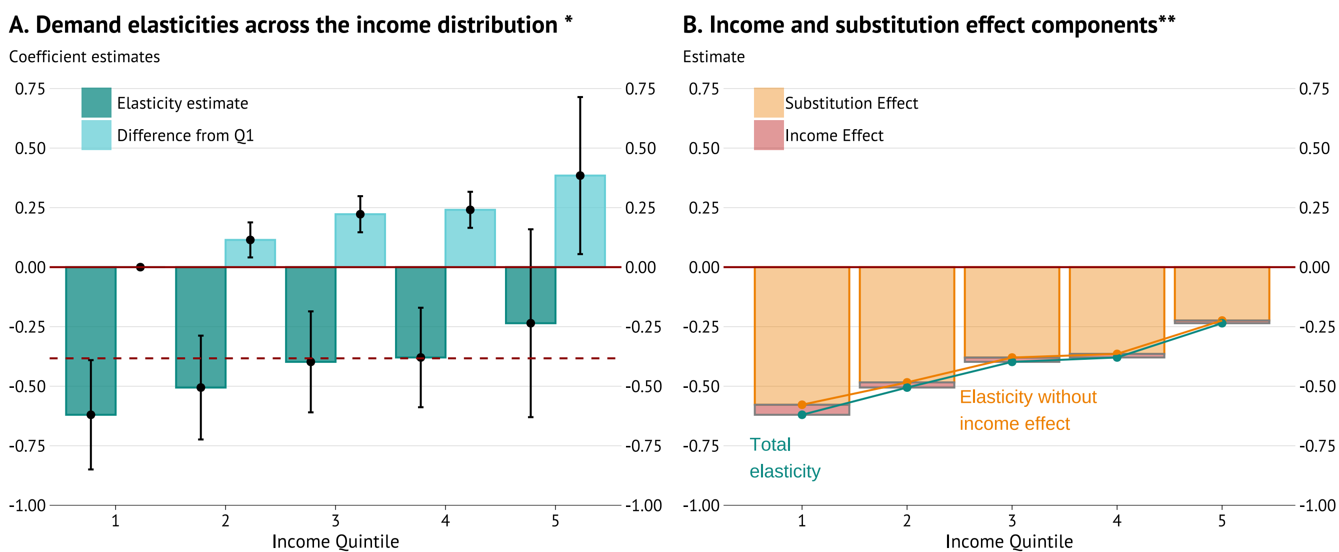

* Panel A shows estimates for the elasticity of demand across income quintiles. Point estimates and the 95% confidence interval are shown. The dashed line gives the pooled estimate. The dark green shows the estimate for each income quintile. Light blue shows the estimate of the difference between a quintile and quintile 1.
** Panel B shows the contribution of income effects and substitution effects to the total elasticity in panel A. This decomposition assumes an income elasticity of demand for petrol of 0.6. The total elasticity line traces out the same point estimates as from panel A. The no income effect line traces out the Hicksian demand elasticities (i.e. the elasticities with income effect components removed.
Source: e61 Institute

Care should be taken when interpreting the welfare implications of this result. The elasticity of demand for any good reflects both income and substitution effects. When the price of a good rises, people may consume less because they switch to other products (the substitution effect), or because they are now relatively poorer (the income effect). Both contribute to a downwards sloping demand curve, but have different welfare implications.

If lower-income households are more elastic primarily because price rises make them relatively poorer, then this implies large welfare costs. On the other hand, if they are more elastic because they can substitute away more readily, the welfare loss may be smaller for them. In other words, what the elasticity gradient in Figure 6 panel A above means for welfare depends on whether it is driven by differential income or substitution effects across the income distribution. One way to get at this is to ask how much of the gradient could plausibly be explained by income effects (i.e. by petrol's budget share being larger at the bottom of the income distribution).

We find that the gradient is far too steep to be driven solely by greater income effects at the bottom of the distribution. In panel B of Figure 6, we show what happens to the elasticity gradient when we 'account' for the plausible contribution of income

10 A caveat here is it will matter whether the consumption response is driven by income or substitution effects. We discuss this explicitly in the next section.
11 In Appendix C.3 we lay out the econometric approach for this.

effects.[12] Accounting for income effects flattens the gradient slightly, but not by much. This implies that the difference in elasticity between high- and low-income households is driven largely by lower-income households being more willing to substitute away from petrol use when prices rise.[13] This softens the regressivity of price increases, relative to a world where all households substitute equally.

Overall, price increases remain regressive

Even with the stronger behavioural response of low-income households, price increases remain highly regressive. The core reason is that the greater budget share exposure of lower-income households (Figure 5, panel B) dominates the offsetting effect of low-income households having higher elasticities of demand (Figure 6, panel A). Petrol takes up so much more of low-income budgets that their sharper cutback isn't enough to even things out.

We measure this overall regressivity in two ways. First, we estimate how price rises shift petrol spending as a share of income. We find that lower-income households see a larger increase in their petrol budget share following a price rise (Figure 7, panel A). On average, a 20% increase in price lifts petrol's budget share by 0.82 percentage points for the lowest income quintile, compared to 0.4 percentage points for the top quintile.[14]

Figure 7: Incidence of price increases

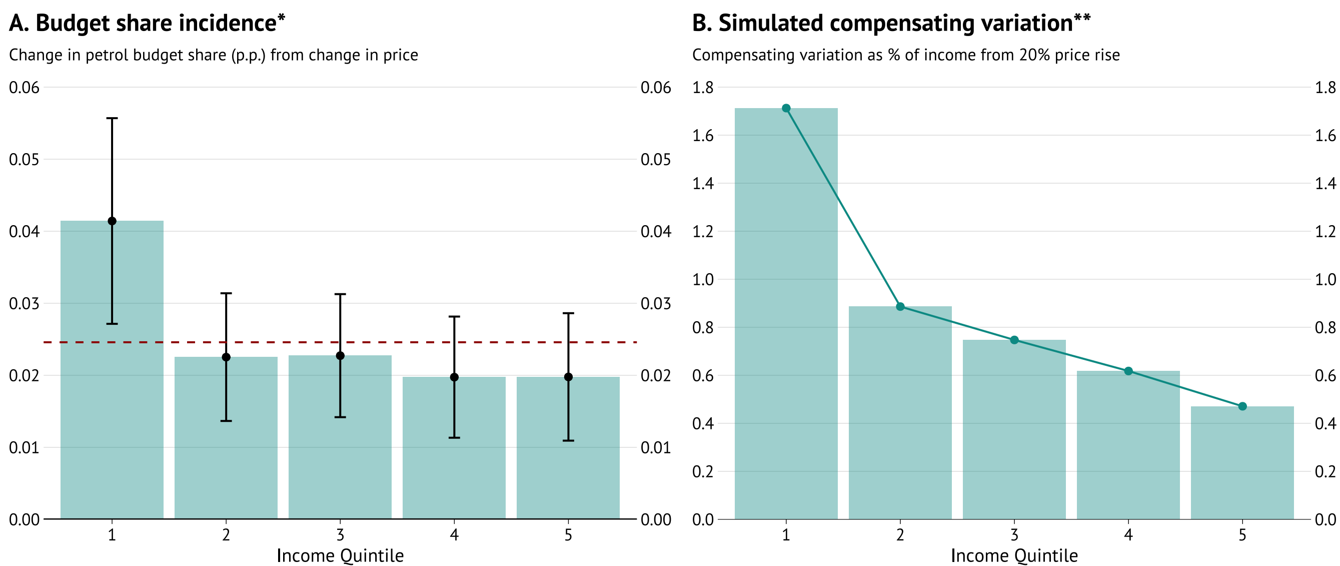

* Panel A shows the coefficient estimates for the effect of price changes on the proportion of income spent on petrol. The dashed line gives the pooled estimate of approximately 0.02. The interpetation of this estimate is that a 1% increase in price, on average leads to a consumer spending 0.02 more percentage points of their budget on fuel.
** Panel B shows the simulated compensating variation from a 20% price increase as a proportion of income. A consumer in the lowest income quintile would need to be paid 1.7% of their income to make them equally well off compared to before a 20% price increase.
Source: e61 Institute

Second, we use the decomposition from Figure 6 panel B to simulate the compensating variation from a 20% price spike across the income distribution. This compensating variation represents the income transfer required to keep a household as well off (their utility constant) following a price change.[15] Households in the bottom income quintile require roughly 3.6 times more compensation as a share of income than those in the top quintile (Figure 7, panel B).[16] Applied to the 2026 Iran

12 More precisely, we use the Slutsky decomposition to decompose our estimated price elasticity of demand into an income effect component and a substitution effect component. The income effect is a function of budget shares and the income elasticity of demand (how much consumption of petrol changes by when income changes). We use our panel data to estimate the income elasticity of demand, as well as mean budget shares for each income group. Combining this with our price elasticity estimates allows us to recover the income and substitution effect components for each income group. In Appendix C.3, we formally lay out the assumptions made to derive this result, and discuss its robustness.
13 This might be consistent with, for instance, lower-income households having a lower opportunity cost of time, and therefore being more willing to substitute towards more time-consuming commuting options.
14 We outline the econometric specification for this result in Appendix C.3.
15 In other words, the compensating variation is an income-equivalised measure of the welfare costs of price changes. In Appendix C.3, we provide more details on this exercise, including how to derive the compensating variation formula, and how to apply it to get an income-equivalised measure of the welfare cost of any given price change.
16 In dollar terms (rather than share of income terms), this compensating variation is increasing with income. In other words, higher-income consumers require a larger transfer (in dollars) to compensate for a price increase. We show this in Appendix C.5. This is because petrol consumption (and hence spending) increases with income. We argue that effects expressed as a share of income are the appropriate way to think about welfare in this context.

shock, this suggests that at its peak the rise in petrol prices was equivalent to a 3.28% real income shock at the bottom of the distribution, compared to a 0.93% shock at the top.

These income results have two main takeaways. First, price rises reduce consumption of petrol more at the bottom of the income distribution than at the top, because lower earners substitute away more. Second, price rises hurt consumers at the bottom of the income distribution more than those at the top, because petrol spending takes up a larger share of their income.

Regions with greater motor vehicle reliance are less responsive to price movements

Finally, we investigate how the responsiveness of petrol consumption to price varies by geography.[17] We find that demand is significantly more elastic in urban areas than in rural regions (Figure 8). This pattern is consistent with urban consumers having a stronger outside option to vehicle use. Consumers in urban areas are likely to be better serviced by public transport, commute shorter distances, and have greater access to flexible work arrangements, making it easier for them to substitute away from petrol consumption when prices rise.

Simple urban-rural splits may mask important heterogeneity. Indeed it seems likely that motor vehicle reliance will differ substantially within Australia's urban environments. To further study these differences, we first group Statistical Area 4s (SA4s) into quintiles of population density. We might expect households in denser areas to have stronger outside options to motor vehicle use. Consistent with this, we find that demand is more elastic in areas with greater population density.

Next, we take the set of urban SA4s in our sample, and split them based on measures of motor vehicle reliance from Census data.[18] We find suggestive evidence that areas with lower car commuting shares, and higher public transport use have more elastic demand, consistent with residents in these areas having better outside options to petrol use (Figure 8).[19]

Figure 8: Differences in the elasticity of demand by geography

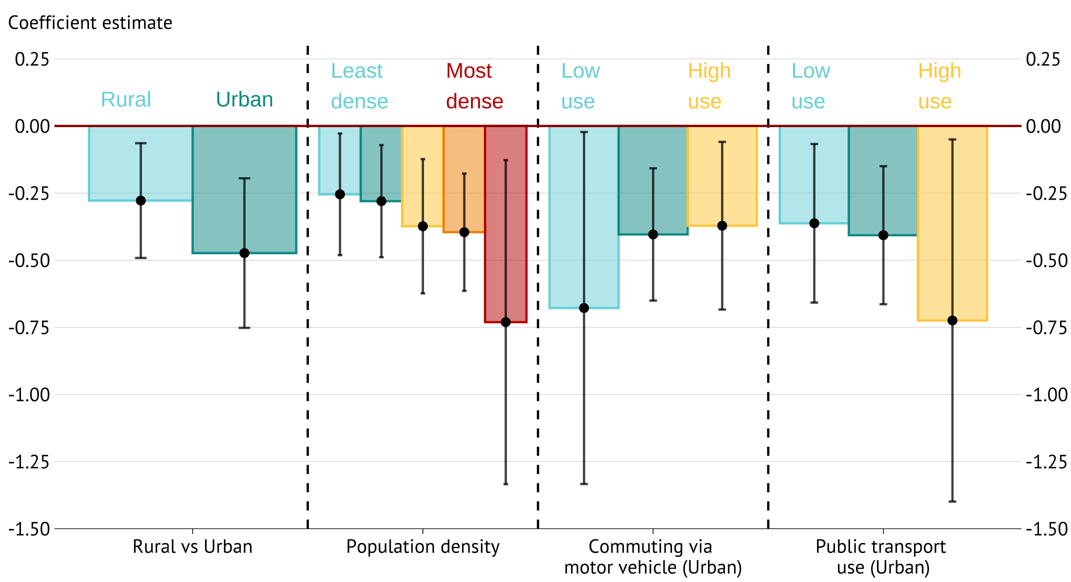

* Figure shows heterogeneity in demand elasticity estimates by geographic groupings. All geographic groupings are done at the SA4 level, and accounts are assigned to a group based on their SA4. Panel A groups SA4s into urban vs rural Panel B groups SA4s into quintiles of population density. Panels c and D group only urban SA4s into terciles based on methods of travel to work (vehicle and public transport). The 95% confidence interval shows if a coefficient estimate is statistically significantly different from 0. In Appendix C.4. we report differences and their significance in estimates across groups.
Source: e61 Institute

17 In Appendix C.4 we provide more information on the geographic variables and econometric specification used to test for heterogeneity. We also report the statistical significance of the difference in estimates across groups.
18 SA4 regions in metropolitan areas typically have populations of between 300,000 and 500,000 people. For each SA4 in Sydney, Perth and Brisbane, we use Census data to compute the share of residents commuting to work via car, and the share commuting to work via public transport. We then assign SA4s into terciles based on these measures.
19 However it's important to note that we find large standard errors on our estimates, and the differences between groups are not statistically significant at the 5% level.

Policy implications

This note has three main findings. First, we estimate that Australian households have a price elasticity of demand for petrol of −0.38, which is three times larger than previous Australian estimates. Second, lower-income consumers are more responsive to price movements than higher-income ones, but price shocks remain regressive, because petrol spending takes up a larger share of incomes at the bottom of the income distribution. Third, the responsiveness of demand varies by geography, with consumers in areas with more alternatives to car use being more responsive to movements in price.

Each of these findings has implications for policy, both in terms of how best to support households through global oil market volatility, and in the trade-offs for using price-based instruments to target the negative externalities of petrol use.

Price based instruments for emissions reductions

Our larger aggregate elasticity estimate implies price-based instruments targeting the negative externalities of petrol use (e.g. fuel excise, carbon pricing and road user charges) are likely to be more effective than previous evidence suggests. This matters for efforts to decarbonise Australia's transport use, where such instruments are one of the main options for cutting emissions. A larger elasticity means that a given rise in the carbon price or fuel excise drives more behavioural adjustment, and hence greater emissions reductions.

At the same time, our distributional findings imply that instruments of this kind place a larger cost on lower-income earners. Governments using price-based approaches to cut emissions could offset these regressive impacts by pairing them with compensatory transfers targeted at low-income earners.

Household support around global oil shocks

Perhaps more pressingly, our results also speak to the trade-offs in how governments support households through global oil price shocks. That petrol price rises fall hardest on low-income earners strengthens the case for some form of support during a crisis.

Figure 9: Effects and targeting of temporary fuel excise reductions

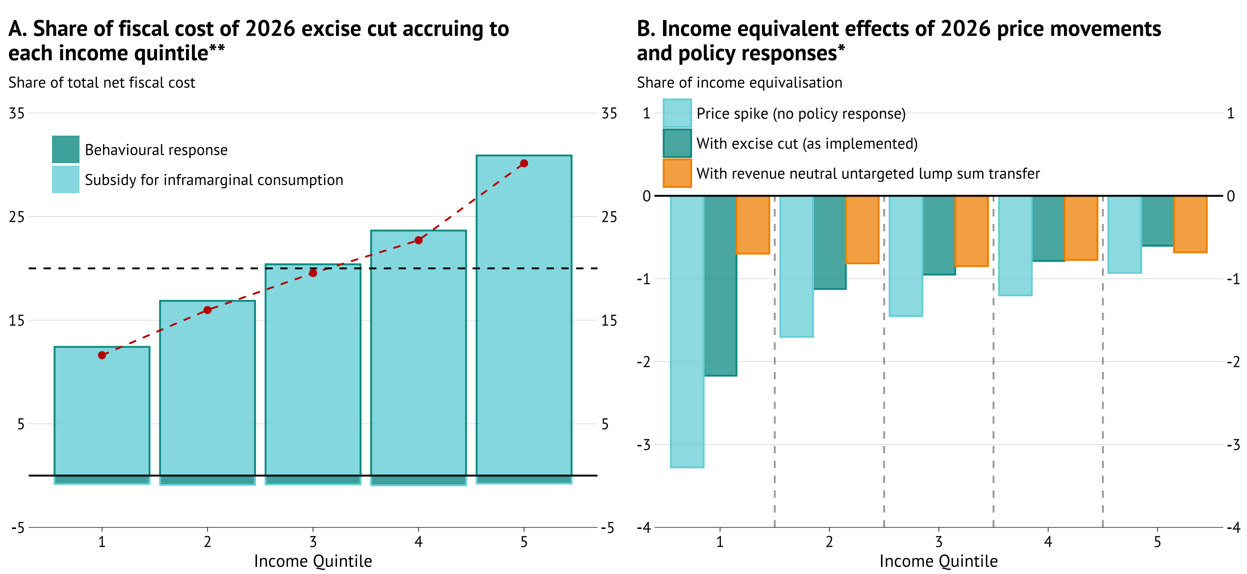

* Panel A shows the share of the excise cut's fiscal cost accruing to each income quintile. The change in revenue from the cut has two components. The first is the greater revenue caused by increased quantities sold under the new lower price. We call this the behavioural response component. Because it increases revenue it is represented as a negative to the fiscal cost. The second is the loss in revenue on every unit that would anyway have been purchased at the old (higher) price level. We call this the subsidy for the inframarginal units of consumption. The points and dashed line give the net effect between the two. The horizontal dashed line gives the equal share benchmark.
** Panel B shows the share of income equivalent effects of movements in price following the outbreak of conflict in Iran, along with associated policy responses. The first bar shows the effect of the peak of the price spike with no policy response. The second bar shows the effect with the excise cut implemented by the Federal Government. The third bar shows the effect of the price spike with no excise cut but instead a revenue neutral untargeted lump sum transfer. Specifically, we simulate the fiscal cost of the excise cut per consumer, and instead distribute it equally amongst drivers as a lump sum transfer.
Source: e61 Institute

In Australia and overseas, the most common tool to deliver relief during the recent shock has been temporary fuel excise reductions. Just as price rises hurt low earners most, price falls benefit them most, so an excise cut is progressive – it provides the greatest benefit (as a proportion of income) to the lowest income earners.

But a progressive policy is not necessarily a well targeted one. In dollar terms, a large share of the fiscal outlay of an excise cut flows to higher-income consumers because they simply buy more petrol (Figure 9 panel A).[20]

The poor targeting of an excise cut is clear when compared to an equally costly (revenue neutral) flat lump-sum transfer. Figure 9 panel B steps through an example that shows this. It plots the income-equivalent effects of the 2026 price shock and policy responses across the income distribution under three scenarios.

The first scenario models the impact of the price shock alone (light blue bar). At its March peak, the shock was equivalent to a 3.28% real income hit at the bottom of the income distribution, which was 3.5 times the shock at the top (0.93%).

The second scenario adds the offsetting effect of the excise cut as implemented (dark teal bar). This reduces the effect of the shock for all groups, and by more for low-income households. But the combined shock remains regressive. Accounting for the excise cut's effects, the bottom quintile still faces a real income hit of 2.17%.

The third scenario pairs the price shock with a flat lump-sum transfer paying every driver the same dollar amount regardless of petrol consumed (dark orange bar).[21] The total cost of the transfer is held equal to that of the modelled excise cut, so differences in its impact reflect targeting alone.[22] Because more of the outlay of a flat lump-sum transfer (relative to an excise cut) goes to lower-income households, who drive less, the regressivity of the shock is reduced. Under the flat lump-sum transfer the income-equivalent shock is almost exactly the same across the distribution, and only drivers in the top income quintile, who consume the most petrol, are worse off than under the excise cut.

Poor targeting is not the only drawback of an excise cut. Because we find a larger, substitution-driven elasticity of demand, holding the price of petrol down through an excise cut also carries a larger efficiency cost than if the demand response was lower. In particular, it suppresses the behavioural adjustment that consumers would otherwise make, likely working against other government measures aimed at shoring up the supply of petrol.

The natural alternative to an excise cut is direct income support through the tax and transfer system. This avoids both of the main problems of an excise cut. First, it can be targeted more towards lower-income households, who are more exposed to and affected by the shock. Second, it doesn't distort the price signal, preserving the behavioural adjustment that an excise cut suppresses. Perhaps there are challenges to targeting or implementing this support in a timely manner. In this case, our estimates suggest that even an untargeted flat lump-sum transfer to all drivers might be preferable to an excise cut.

20 In Appendix C.5 we lay out a framework to calculate the total cost of the 2026 fuel excise cut, as well as the shares of this cost accruing to each income group.

21 This is similar in principle to the flat dollar cuts to vehicle registration fees implemented by states such as Victoria and New South Wales.

22 We provide more details on how we implement this lump-sum transfer in Appendix C.5. Conceptually, for any given level of fiscal cost, an excise cut will 'distribute' the cost in proportion to petrol consumption. A flat transfer will instead 'distribute' this cost evenly across drivers.

References

Breunig, R. V., & Gisz, C. (2009). An exploration of australian petrol demand: Unobservable habits, irreversibility and some updated estimates. *Economic Record*, *85*(268), 73–91.

Chen, J., & Roth, J. (2024). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics*, *139*(2), 891–936.

Coglianese, J., Davis, L. W., Kilian, L., & Stock, J. H. (2017). Anticipation, tax avoidance, and the price elasticity of gasoline demand. *Journal of Applied Econometrics*, *32*(1), 1–15.

Havranek, T., & Kokes, O. (2015). Income elasticity of gasoline demand: A meta-analysis. *Energy Economics*, *47*, 77–86.

Hendel, I., & Nevo, A. (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica*, *74*(6), 1637–1673.

Hughes, J. E., Knittel, C. R., & Sperling, D. (2008). Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal*, *29*(1), 113–134.

Kilian, L., & Zhou, X. (2024). Heterogeneity in the pass-through from oil to gasoline prices: A new instrument for estimating the price elasticity of gasoline demand. *Journal of Public Economics*, *232*, 105099.

Knittel, C. R., & Tanaka, S. (2021). Fuel economy and the price of gasoline: Evidence from fueling-level micro data. *Journal of Public Economics*, *202*, 104496.

Levin, L., Lewis, M. S., & Wolak, F. A. (2017). High frequency evidence on the demand for gasoline. *American Economic Journal: Economic Policy*, *9*(3), 314–347.

Li, S., Linn, J., & Muehlegger, E. (2014). Gasoline taxes and consumer behavior. *American Economic Journal: Economic Policy*, *6*(4), 302–342.

Maitra, A. (2025). Fuel excise cuts - relief today, risks tomorrow? [e61 Institute, plus61 series].

Maitra, A. (2026). Policy at the pump: Lessons from 2022's oil price shock [e61 Institute, plus61 series].

Park, S. Y., & Zhao, G. (2010). An estimation of us gasoline demand: A smooth time-varying cointegration approach. *Energy Economics*, *32*(1), 110–120.

Silva, J. S., & Tenreyro, S. (2006). The log of gravity. *The Review of Economics and statistics*, 641–658.

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.

A.1. Price and cost data

The empirical work in this research note is based on a rolling, account level panel of petrol purchases, linked to petrol prices, wholesale petrol costs, and account characteristics. We construct this panel by combining data from multiple distinct sources.

We collect data on wholesale petrol costs (known as the Terminal Gate Price and hereafter called TGP) from the Australian Institute of Petroleum (AIP).[23] The AIP records the daily average TGP for unleaded fuel in each capital city. We use the AIP data to construct daily – capital city level panels for TGP in Sydney, Brisbane and Perth between 1 January 2020 and 30 April 2026.

We obtain daily retail petrol price data from government run information platforms in New South Wales, Queensland and Western Australia.[24] We select these three states because their state-run information platforms allow us to collect publicly available price data over the entirety of our sample period. Combined, these three states have a population of 17.3 million people, or about 60% of Australia's total population.

In New South Wales, we collect this data from FuelCheck. Service stations in NSW are required by law to register with FuelCheck and notify the platform of the retail price of each fuel type available at their service station. State legislation requires that the price registered with FuelCheck must match the prices offered by the station. FuelCheck maintains historical records of prices submitted by stations since August 2016. These records include the station's name, brand and location, as well as the date, time and new price of a submission. Prices are reported in cents per litre. From November 2022, the FuelCheck records also include price submissions for stations in the Australian Capital Territory.

In Queensland, we collect this data from the Queensland Government Open Data Portal. Service stations in QLD are required to report their fuel prices within 30 minutes of a price change. The Queensland government then publishes historical records of the prices submitted by stations since December 2018. Like the FuelCheck scheme, these records include the station's name, brand and location, as well as the date, time and new price of a submission.

Finally, in Western Australia, we collect this data from FuelWatch. Per state law in WA, petrol retailers adhere to a '24-hour' rule whereby every day at 2pm they upload prices to the government platform for all of their stations. Stations are required to set and keep fixed these prices for 24 hours starting at 6am the next day. The FuelWatch platform publishes these daily prices dating back to February 2001. The same core set of information is observed in these historical records as in NSW and QLD.

In each state, we use the collected data to construct daily, station level panels of posted petrol prices. Each panel contains the universe of stations in their respective state, and covers each day from 1 January 2020 to 30 April 2026. Throughout, we focus on prices for Unleaded 91 as it is the most popular grade sold.

A.2. Transaction data

We construct information on petrol purchasing from a large dataset of aggregated, consented and de-identified bank transaction records. The data are collected by a third-party provider during routine consumer finance applications.[25] The transaction data provides a backwards looking window of financial history (typically 90-180 days of transactions across relevant accounts).

23 Available at: https://www.aip.com.au/historical-ulp-and-diesel-tgp-data (last accessed on 04 June 2026).
24 The price datasets are available at (all accessed June 30, 2026):

- New South Wales: https://data.nsw.gov.au/data/dataset/fuel-check
- Queensland: https://www.data.qld.gov.au/organization/treasury
- Western Australia: https://catalogue.data.wa.gov.au/dataset/fuelwatch-historic-fuel-prices

25 These include, for instance, applications for a mortgage, credit card, buy-now-pay-later services, personal finance management apps, and a change in utility providers.

The data provider cleans and aggregates the data: categorising transactions, constructing demographic variables, and removing personally identifiable information.

The sample we use for our empirical work includes accounts observed for at least one week over the period from 1 January 2020 to 30 April 2026.[26] We restrict this sample in a set of important ways for our analysis.

First, we select accounts categorised as being located in either New South Wales, Queensland or Western Australia. Second, we select accounts from a subset of institutions that reliably report both the settlement date and transaction date of a transaction. Finally, we filter to accounts that have at least one petrol transaction over their period of observation, and have an average real weekly income of at least $50 (measured in 2026 dollars).

For this final sample of aggregated accounts, we collect their full set of petrol transactions and income flows over the period they are observed over. We identify these transactions based on categorisation by the data-provider, as well as the third party reported on the transaction.

A.3. Panel construction

To construct the final panel used for our empirical work, we link the data from the various sources described above. There are a number of important steps in this process.

First, we impute petrol purchase quantities by linking expenditures to prices paid. In the transaction data we observe the date of the transaction, the total expenditure in dollars, and a set of information on where the transaction takes place.[27] We map each transaction to the Statistical Area 4 (SA4) it takes place in.[28] We then construct the average retail price in the SA4 on the date of the transaction $p_{sa4,t}$. We drop all transactions when expenditure e_{it} is less than $20, as these are likely to reflect convenience store purchases rather than petrol expenditures.[29] With the remaining sample, we impute the quantity of fuel q_{it} purchased by account i on date t as

$$q_{it} = \frac{e_{it}}{p_{sa4,t}} \tag{1}$$

Second, we construct average real weekly income at the account level by summing all observed income flows, adjusting for inflation, and normalising by the number of weeks the account is observed for. Using this measure we construct quintiles and buckets of the income distribution.

Third, we proxy for the price level facing an account with the average price in the SA4 the account is located in. Likewise, we proxy for the cost level with the TGP for the state the account is located in.

Finally, we aggregate up to the monthly level. The final panel is at the account-month level, and includes the following variables:

- Total petrol expenditure in dollars e_{it} by account i in month t.
- Total petrol quantity purchased in litres q_{it} by account i in month t.
- The average price level p_{it} faced by account i in month t.
- The average cost level c_{it} faced by account i in month t.
- A dummy for if account i purchases petrol in month t, $d_{it} = \mathbf{1}\{q_{it} > 0\}$.
- The number of petrol transactions n_{it} made by account i in month t.
- The average size of each transaction made by account i in month t, $s_{it} = q_{it}/n_{it}$
- Nominal y_{it}^N and real y_{it}^R income earned by account i in month t.

26 In our baseline results we cut off the sample at the end of 2025.
27 This information generally includes the suburb of the transaction and the brand of the service station. In some cases it will also include a store code, which allows us to match the transaction to the precise station it takes place at.
28 In principle, with complete price data we can do better and map transactions to the station they occur at. In practice, the set of information contained in the transaction data varies considerably and is often incomplete. We take a conservative and consistent approach with our SA4 mapping.
29 By relying on a measure of expenditure that potentially includes some non-petrol-related spending, we are implicitly assuming that the size of these non-petrol expenditures is uncorrelated with petrol prices. If, when petrol prices rise, people also spend less at petrol station convenience stores while filling up, this could lead to a larger decrease in measured spending that is unrelated to fuel demand.

- The budget share of petrol b_{it} for account i in month t. We compute this budget share by taking petrol expenditure e_{it} in the month, and dividing by the account's average level of monthly income (i.e. within an account, across time, movements in b_{it} come from movements in petrol expenditures rather than movements in income).
- A vector of account level demographics X_i that include the average real weekly income earned by the account x_i, and measures of geography – most granularly the postcode pc_i that the account is located in.

We also construct this panel at weekly and fortnightly frequency. The final panels cover approximately 25 million petrol purchases by 800,000 unique accounts over the period from 1 January 2020 to 31 December 2025.

B.1. Econometric framework

Let q_{it} be the quantity of petrol purchased by account i in period t. The elasticity of demand for petrol is defined as the change in this quantity with respect to a change in price.

(1) $$\varepsilon = \frac{\partial \log(q_{it})}{\partial \log(p_{it})}$$

When working with aggregate data, it is common to regress the log of purchase quantity on the log of the price level to estimate the elasticity ε. However, in our setting, we occasionally observe accounts that make no purchases over a period. As such, this common approach would select on the event of a purchase. To overcome this issue, we estimate the elasticity using the Poisson quasi-maximum likelihood estimator (PQMLE). The PQMLE deals with zeros whilst directly estimating the percentage change in the outcome variable. Formally, the PQMLE assumes the data generating process below

(2) $$q_{it} = \exp\left(\alpha_i + \lambda_{g(t)} + \beta \log(p_{it}) + \xi_{it}\right)$$

where p_{it} is the average price level faced by account i in period t, α_i is an account specific demand shifter, $\lambda_{g(t)}$ are time specific demand shifters and ξ_{it} capture idiosyncratic unobserved shocks to demand for account i in period t. We restrict the time fixed effects $\lambda_{g(t)}$ such that they are more aggregated than the price level time-series p_{it}.[30]

The parameter we seek to identify is β. Under this specification, the coefficient β directly recovers the elasticity of demand for petrol.

(3) $$\varepsilon = \beta$$

The key identifying assumption made by the PQMLE is that the conditional mean of quantity is correctly specified. In equation 2, this means that, conditional on fixed effects, the unobserved shocks ξ_{it} should be uncorrelated with price such that the conditional expectation of quantity takes the form below.

(4) $$E(q_{it}|p_{it}, \alpha_i, \lambda_{g(t)}) = \exp\left(\alpha_i + \lambda_{g(t)} + \beta \log(p_{it})\right)$$

Importantly, there is no requirement that the outcome variable follow a Poisson process, nor any assumption on the variance of quantities. The consistency of the estimator requires only that this conditional mean in equation 4 is correct (Silva and Tenreyro 2006; Wooldridge 2010, Chapter 18.2).

In addition, the PQMLE assumes that the conditional mean takes log-linear functional form. Effectively, this means that in expectation, a given percentage change in price induces the same percentage change in quantities across accounts and time.

B.2. Identification and inference

Within our econometric framework laid out above, the elasticity of demand is identified using within-account variation in petrol purchases in response to the price level. There are two important threats to identification in this setting. The first is endogeneity from unobserved demand shocks, and the second is consumer inventory dynamics.

30 For instance, if we estimate equation 4 at the monthly frequency, then λ_t is able to include month of year and year level fixed effects, but cannot include month of sample fixed effects.

Price endogeneity

The canonical threat to the identification of demand comes from simultaneity between price and quantity. Observed prices and quantities under market equilibrium are jointly determined by demand and supply equations. As a result, if sufficient controls are not included in the demand equation, unobserved shifts in demand will trace out the supply curve rather than demand. The resulting endogeneity of price will bias estimates of ε towards 0.[31] In general, the severity of this bias will depend on two factors: (1) the inelasticity of the supply curve, and (2) the prevalence and magnitude of unobserved demand shifters relative to unobserved supply shifters (Levin et al., 2017).[32]

To identify demand, we want to control for these unobserved demand shifters so that the resulting variation in prices is driven by supply shifters, allowing us to trace out the demand curve. We adopt a twofold approach to identification.

First, we include extensive fixed effects by accounts and by time. Much of the variation in petrol demand across the sample is likely to come from either seasonal factors, long run trends, or persistent differences in petrol consumption across accounts. To control for these factors we include account-level fixed effects, calendar-month fixed effects, and year fixed effects. These controls mitigate the unobserved shifts in demand that would otherwise introduce endogeneity in price.

Second, we use an instrumental-variables approach to address any remaining endogeneity in price. We use a standard instrument from the demand estimation literature. A valid instrument for price must be relevant and exogenous, meaning that it must have explanatory power over price and not be correlated with any unobserved demand shocks captured by ξ_{it}. In practice, credible instruments for price will take the form of exogenous cost shifters.

In line with this, we use the wholesale petrol spot price (TGP) as an instrument for price. TGP constitutes the largest component of a station's marginal cost and is therefore highly relevant. The key identifying assumption is that the wholesale spot price is exogenous – uncorrelated with unobserved demand shocks ξ_{it}. There are two reasons why this assumption is likely to be credible in the Australian context. First, consumers are unlikely to condition their purchasing decisions on wholesale prices. Second, Australia, as a price taker in global markets, is unlikely to experience aggregate demand shocks that affect the globally determined wholesale spot price.

Demand accumulation

A second threat to identification in our context comes from consumer inventory dynamics. Petrol is a storable good, meaning that the timing of purchase is separated from the timing of consumption. We wish to identify the elasticity of consumption with respect to price changes; however, we observe purchases rather than consumption. In these settings, static analysis of demand may be biased by dynamic consumer behaviour (Hendel & Nevo, 2006).

If consumers respond to price movements by changing both how much they consume and when they purchase, then this intertemporal substitution effect will bias estimates of the demand elasticity. This bias could work in either direction.

If prices rise but consumers expect them to come back down, they may delay purchases, shifting them away from expensive periods and towards cheaper ones. In this scenario, static analysis would overestimate demand responsiveness. Alternatively, if prices rise and consumers expect them to keep rising, then they may bring forward their future purchases. In the short run, these 'demand-run' style dynamics would cause static estimates to understate the true demand elasticity (or perhaps even induce an incorrect positive slope in the demand relationship).

Hendel and Nevo (2006) show that in principle, it is possible to account for these dynamics either through additional structure or through aggregation. We exploit the fact that petrol's storability is constrained by tank capacity, making it difficult for consumers to intertemporally substitute over longer time horizons.[33] Following this logic, we estimate equation 4 at a longer time horizon (monthly), over which time inter-temporal substitution style behaviour is more likely to aggregate out.

Inference

When estimating equation 4, we construct clustered standard errors at the account-period level. This allows for possible persistence in the unobservables ξ_{it} within account i over time, and across all accounts in a given period t.

31 This bias is commonly referred to in the literature as attenuation bias. In an intuitive sense, applying OLS to observed market outcomes will map out some combination of the demand and supply curves. The result will be less downward sloping than the true demand curve.
32 This attenuation bias will be smaller if the supply curve is fairly elastic or unobserved demand shocks are small compared to unobserved supply shocks.
33 Intuitively, if a consumer knows they require fuel, they could choose to fill up tomorrow rather than today. But it would be difficult to wait four weeks from now to fill up, or to purchase enough today to last multiple months' worth of consumption

C.1. Descriptive evidence

Figures 2 and 4 present descriptive evidence on the co-movement of wholesale costs, retail prices and purchase quantities over 2022 and 2026. We generate this descriptive evidence using the account-weekly panel of prices, costs and quantities. For the time-series movements of price and costs, we regress the log of the cost and price levels on weekly-dummies, and plot these dummies. Each weekly dummy gives the percentage change in cost/price relative to the base week (the omitted weekly dummy). Formally, we estimate the equations below

$$\log(c_{iw}) = \mu_w + v_{iw} \tag{1}$$

$$\log(p_{iw}) = \mu_w + v_{iw} \tag{2}$$

and plot the resulting estimates of μ_w.

For the time-series movements in quantities, we regress account weekly quantity (in levels) on weekly dummies and account level-fixed effects. We then plot the weekly dummies, divided by the average weekly purchase quantity. Formally we estimate

$$q_{iw} = \alpha_i + \beta_w + v_{iw} \tag{3}$$

and plot the resulting estimates of $\beta_w/\bar{q}$. The results from each of these descriptive regressions can be found in Figure 2 and Figure 4, in the main body of the note.

C.2. Baseline demand estimates

We estimate the elasticity of petrol demand using the log-linear specification set out in Appendix B. We include account, calendar month, and yearly fixed effects. Our baseline estimating equation is

$$q_{it} = \exp\left(\alpha_i + \beta \log(p_{it}) + \delta_{m(t)} + \lambda_{y(t)} + \xi_{it}\right) \tag{4}$$

One practical difficulty with the PQMLE is that the non-linear structure means we can no longer use the 2SLS estimator to instrument for an endogenous regressor. As such, we address the endogeneity of price through a control function approach. In the first stage we regress the log of price on the log of cost as well as a full set of fixed effects and recover the residuals. We then include these first stage residuals as additional controls in the second stage. Under the identifying assumption that wholesale cost movements are uncorrelated with unobserved demand shocks conditional on fixed effects, the control function approach consistently estimates β.

In Table C.1 we report estimates of the demand elasticity using both the simple one stage PQMLE estimator, as well as the control function approach. In the columns we include progressively richer fixed-effect structures. The control function estimate in column (4) of panel B contains our preferred elasticity estimate of −0.38, implying that a 10% increase in prices yields a 3.8% decrease in petrol demand.

Two patterns from Table C.1 are worth noting explicitly. First, the estimated elasticity grows in magnitude with the inclusion of fixed-effects, particularly year-of-sample fixed effects. This pattern is consistent with these controls absorbing demand-side confounds that would otherwise attenuate the estimate toward zero. Second, the control function estimates are slightly smaller in magnitude than the PQMLE results. One interpretation is that the instrument strips out the component of retail price variation driven by the predictable retail price cycle, which consumers may respond to through intertemporal substitution.

Table C.1: Baseline demand elasticity estimates

	Dependent variable: Quantity (litres)			
	No FE (1)	Account FE (2)	Account + Month FE (3)	Account + Month + Year FE (4)
Panel A: Poisson QMLE Results				
Elasticity (ε)	-0.222^{***}	-0.261^{***}	-0.252^{***}	-0.459^{***}
	(0.08)	(0.08)	(0.06)	(0.11)
Observations	7,288,941	6,985,334	6,985,334	6,985,334
Panel B: Poisson QMLE Control Function results				
Elasticity (ε)	-0.154^{**}	-0.239^{***}	-0.211^{***}	-0.383^{***}
	(0.08)	(0.07)	(0.06)	(0.12)
Observations	7,288,941	6,985,334	6,985,334	6,985,334

Notes: Panel A reports Poisson QMLE estimates. Panel B reports Poisson control-function estimates using log(cost) as an instrument for log(price). The dependent variable is quantity purchased measured in litres. Standard errors are reported in parentheses. All specifications cluster standard errors two-way by account and period. Significance is given by *p<0.1; **p<0.05; ***p<0.01.

We estimate a set of additional models to decompose the demand elasticity into extensive and intensive margin components. First, we model how the number of transactions within a month, and the average size of a transaction vary with price.

$$n_{it} = \alpha_i + \lambda_{g(t)} + \theta_1 \log(p_{it}) + \xi_{it} \tag{5}$$

$$\log(s_{it}) = \alpha_i + \lambda_{g(t)} + \theta_2 \log(p_{it}) + \xi_{it}, \quad n_{it} > 0 \tag{6}$$

We estimate equation 6 conditional on $n_{it} > 0$, and use the same empirical approach as in the baseline specification, with account-level, calendar month and year fixed effects, and using wholesale cost to instrument for price. Under this specification, θ_2 is the elasticity of transaction size with respect to price, whilst θ_1 gives the effect of price on the number of transactions made in a month. We divide θ_1 by the average number of monthly transactions $\bar{n}$ to recover the elasticity of transaction counts with respect to price. We report estimates of each of these coefficients in Table C.2, Panel A. The estimates imply that a 10% increase in price leads to a 4.6% increase in the number of transactions made in a month, and a 7.7% fall in the average size of each of those transactions. We refer to this pattern of higher prices leading to more frequent, smaller transactions as transaction size smoothing.

Table C.2: Decomposing petrol demand elasticity

	Panel A: Transaction size				**Panel B**: Purchase probability			
	θ_1	θ_2	$\bar{n}$	$\theta_1/\bar{n}$	γ_1	γ_2	$\bar{d}$	$\gamma_1/\bar{d}$
Estimate	1.523^{***}	-0.766^{***}	3.305	0.461^{***}	0.076^{***}	-0.393^{***}	0.678	0.112^{***}
	(0.270)	(0.022)		(0.082)	(0.022)	(0.078)		(0.033)
Observations	7,267,566	4,868,784	–	–	7,267,566	4,868,784	–	–

Notes: The table decomposes petrol demand responses across behavioural margins. Panel A reports estimates from the transaction size decomposition model, where θ_1 captures the effect of prices on the number of purchases in a month, and θ_2 captures the effect of prices on average fill-up size. Scaling θ_1 by the sample mean of the dependent variable $\bar{n}$ gives the elasticity of the number of transactions in a month with respect to price. Panel B reports estimates from the two-part model, where γ_1 captures the effect of prices on the probability of making a purchase, and γ_2 captures the effect of prices on quantity conditional on purchase. Again, scaling γ_1 by $\bar{d}$ gives the elasticity of the purchase probability with respect to price. All specifications include account fixed effects, month-of-year fixed effects, and year-of-sample fixed effects. Retail price is instrumented using wholesale fuel cost. Standard errors are reported in parentheses, and clustered at the account–month level. Significance is given by *p<0.1; **p<0.05; ***p<0.01.

Second, we separately model the probability that an account makes a purchase in month t, as well as the quantity they purchase conditional on making a purchase.

$$d_{it} = \alpha_i + \lambda_{g(t)} + \gamma_1 \log(p_{it}) + \xi_{it} \tag{7}$$

$$\log(q_{it}) = \alpha_i + \lambda_{g(t)} + \gamma_2 \log(p_{it}) + \xi_{it}, \quad d_{it} = 1 \tag{8}$$

We estimate equation 8 conditional on $d_{it} = 1$. Here, γ_1 captures how the probability of making a purchase changes with movements in price, whilst γ_2 captures how the quantity purchased, conditional on making a purchase, changes with price. We report estimates of these coefficients in Table C.2, Panel B. At the monthly level we find a small but significant positive relationship between price and the probability of making a purchase. This relationship is driven by the smoothing behaviour documented above – as prices rise, consumers make more frequent purchases, making them more likely to record a purchase over any length of time.

C.3. Heterogeneity by income

Heterogeneity in the demand response

We next relax our baseline demand model in equation 4 to flexibly allow the elasticity of demand to vary across the income distribution. Specifically, we estimate the following interaction model

$$q_{it} = \exp\left(\alpha_i + \sum_{k=1}^{5}\left[\beta_k\left(\log(p_{it}) \times Inc_{ik}\right)\right] + \delta_{m(t)} + \lambda_{y(t)} + \xi_{it}\right) \tag{9}$$

(10)

where Inc_{ik} is a dummy variable equaling 1 if account $i's$ average real weekly income falls within quintile k of the income distribution across the accounts in our sample. All other coefficients in 9 are the same as in 4 above. To identify the $\beta_1, ..., \beta_5$ coefficients in 9, we instrument for each price income interaction term $\log(p_{it}) \times Inc_{ik}$ with the corresponding cost income interaction $\log(c_{it}) \times Inc_{ik}$.We report the results by income in Table C.3. We also report a set of results omitting the first income quintile to assess statistical significance in differences.

Table C.3: Demand elasticity estimates by income quintile

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Panel A: Direct estimates					
Elasticity (ε)	-0.620^{***}	-0.506^{***}	-0.398^{***}	-0.379^{***}	-0.235
	(0.117)	(0.111)	(0.108)	(0.107)	(0.201)
Observations	6,676,600	6,676,600	6,676,600	6,676,600	6,676,600
Panel B: Estimates relative to quintile 1					
Difference	–	0.114^{***}	0.222^{***}	0.241^{***}	0.385^{**}
		(0.038)	(0.039)	(0.039)	(0.168)
Observations	6,676,600	6,676,600	6,676,600	6,676,600	6,676,600

Notes: Panel A reports direct Poisson control-function estimates of petrol demand elasticities by income quintile. Panel B reports coefficient differences relative to quintile 1. All specifications include account fixed effects, month-of-year fixed effects, and year fixed effects. Retail price is instrumented using wholesale fuel cost. Standard errors are reported in parentheses. Significance is given by *p<0.1; **p<0.05; ***p<0.01.

The results in Table C.3 suggest that the elasticity of demand for petrol is monotonically decreasing across the income distribution. Consumers in the lowest income quintiles are significantly more elastic than consumers in the top quintile.

Income and substitution effects

When the price of a good changes, the demand response reflects a combination of income and substitution effects. The Slutsky equation decomposes the total demand change between these components.[34] Specifically, the equation (in elasticity form) states that

34 The standard Slutsky equation in levels is given by:

$$\frac{\partial x}{\partial p} = \frac{\partial h}{\partial p} - x\frac{\partial x}{\partial m}$$

where x is the Marshallian demand function, h is the Hicksian demand function, p is price and m is income. Dividing both sides of the equation by x/p yields the decomposition in elasticity form.

(11) $$\varepsilon^M = \varepsilon^H - b \times \eta$$

The left-hand side of the equation ε^M is the Marshallian (or uncompensated) price elasticity. It captures both the tendency to switch away from petrol as it becomes relatively more expensive than other goods (substitution effect), and the tendency to cut consumption because a price rise makes the household effectively poorer (income effect).

The first term on the right-hand side ε_H is the Hicksian (or compensated) price elasticity. It captures the demand response that would occur if a consumer was simultaneously compensated for the loss in purchasing power caused by the price rise. It isolates the pure substitution effect (i.e. how much less petrol would a household consume if relative prices changed but their utility was held constant).

The second term on the right-hand side $b \times \eta$ is the income effect component. The budget share of petrol $b = px/m$ is the fraction of income spent on fuel. η is the income elasticity of demand, capturing how consumption of petrol responds to changes in income. Their product measures how much of the observed demand response is driven by the implicit reduction in real income that a price rise entails. For a normal good ($\eta > 0$), a price rise reduces real purchasing power and further reduces consumption on top of the substitution effect, meaning that ε^M is more negative than ε^H.

We estimate the Marshallian elasticity of demand across income quintiles ε_k^M using equation 9 above. We also directly estimate the petrol budget share b_k for each income quintile, using its conditional mean across accounts within a quintile. Namely, we compute these conditional means directly from the transaction data using the b_{it} variable. We construct b_{it} by taking account i's petrol expenditure e_{it} in the month, and dividing by the account's average monthly income.[35]

With estimates of ε_k^M and b_k in hand, an estimate of the income elasticity of demand η would allow us to back out the implied income and substitution effect contributions to the demand response. The existing literature provides a fairly large range of estimates for this parameter. Surveying the literature, Havranek and Kokes (2015) find mean short-run and long-run η estimates of 0.28 and 0.66.[36] Direct estimates from Hughes et al. (2008) suggest that η ranges from 0.21 to 0.75.

Table C.4: Income and substitution effect decomposition of demand elasticities

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Panel A: Marshallian elasticities and budget shares					
Marshallian elasticity (ε_M)	−0.620	−0.506	−0.398	−0.379	−0.235
Budget share (s_y)	0.090	0.046	0.039	0.032	0.024
Panel B: Baseline income elasticity assumption (η = 0.47)					
Income effect (IE)	−0.042	−0.022	−0.018	−0.015	−0.011
Substitution effect (SE)	−0.578	−0.484	−0.380	−0.364	−0.224
Panel C: Alternative income elasticity assumptions					
Income effect (IE), $\eta = 0.1$	−0.009	−0.005	−0.004	−0.003	−0.002
Substitution effect (SE), $\eta = 0.1$	−0.611	−0.501	−0.394	−0.376	−0.233
Income effect (IE), $\eta = 0.3$	−0.027	−0.014	−0.012	−0.010	−0.007
Substitution effect (SE), $\eta = 0.3$	−0.593	−0.492	−0.386	−0.370	−0.228
Income effect (IE), $\eta = 0.6$	−0.054	−0.028	−0.023	−0.019	−0.014
Substitution effect (SE), $\eta = 0.6$	−0.566	−0.478	−0.375	−0.360	−0.221
Income effect (IE), $\eta = 1.0$	−0.090	−0.046	−0.039	−0.032	−0.024
Substitution effect (SE), $\eta = 1.0$	−0.530	−0.459	−0.359	−0.347	−0.211

Notes: The table decomposes Marshallian demand elasticities into income and substitution effect components using the Slutsky equation, $\varepsilon_M = \varepsilon_H - s_y\eta$, where s_y is the budget share of fuel expenditure and η is the income elasticity of demand. Panel A reports the estimated Marshallian elasticities and fuel budget shares by income quintile. Panel B reports the decomposition using the estimated income elasticity of demand, $\eta = 0.47$. Panel C reports the decomposition under alternative assumptions for the income elasticity of demand. Negative values indicate reductions in fuel demand in response to higher prices.

We directly estimate the income elasticity of demand using our sample and the following model

35 This means that within an account, across time, variation in b_{it} comes from variation in expenditures rather than variation in income. Across accounts, variation in the b_{it} measure comes from variation in both the numerator and the denominator.

36 Havranek and Kokes (2015) argue that adjusting for publication bias in these estimates gives mean short-run and long-run η estimates of 0.1 and 0.46.

$$q_{it} = \exp\left(\eta \log(x_i) + \omega_t + \xi_{it}\right), \tag{12}$$

where (recall) x_i is mean real monthly income for individual i, ω_t is a month-of-sample fixed effect that controls for all temporal variation in petrol consumption, including responses to price fluctuations, and ξ_{it} is the econometric error. The coefficient of interest, η, quantifies the income elasticity. We obtain an estimate (standard error) of $\hat{\eta}$ =0.47 (0.03) using the PQMLE.

The specification in equation 12, effectively recovers η via the average cross-sectional gradient in petrol consumption across the income distribution. We argue that η in this specification is likely to capture the long-run income elasticity as it captures all the differential investments made by households across the income distribution that lead to consume different levels of petrol.

The estimate of η we obtain implies that a 1% increase in income causes a 0.47% increase in petrol consumption. Then, we can back out the income effect across each income quintile as $\bar{b}_k \times 0.47$, and the substitution effect as $\hat{\varepsilon}_k^M + \bar{b}_k \times 0.47$. Using this estimated value of η, as well as the estimates described above, we find that the majority of the demand response in levels, as well as the gradient in this response across income quintiles, is driven by substitution rather than income effects (Table C.4, Panel B).

It's worth noting that the relative importance of income effects in explaining the level of the demand response, as well as the differential across groups, is increasing in the choice of η. Nevertheless, we find that our results are qualitatively robust to the estimated value of η. Namely, we check our result under different conceivable values of η, and show that the results described above remain true. In Table C.4 panel C we take a highly conservative estimate of $\eta = 1$. This would imply that a 1% increase in income leads to a 1% increase in petrol consumption. Even under this parameterisation, we still see that the majority of the demand response, as well as the differential across groups, is driven by substitution rather than income effects.

An alternative way to think about our results here is to ask: what income elasticity of demand would rationalise flat substitution effects across the income distribution? To answer this question, we set the income elasticity of demand in quintile 5 to zero, such that the entire demand response for this group is driven by substitution effects. Then we solve for the value of η in each quintile that would set $\varepsilon_k^H = \varepsilon_5^H$. Formally, for $k \in \{1, 2, 3, 4\}$, we solve the following equation

$$\eta_k = \frac{\varepsilon_5^M - \varepsilon_k^M}{\bar{b}_k} \tag{13}$$

We report the resulting estimates of η_k in Table C.5. The results show that the income elasticity values that would rationalise flat substitution effects across the income distribution are implausibly large. The most intuitive way to think about this result is that relative to the observed gradient of budget shares, the gradient in the Marshallian demand elasticities is too steep to be explained solely by differential income effects across the income distribution.

Table C.5: Income elasticities required for equal substitution effects

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Imputed income elasticity (η_k)	4.27	5.82	4.19	4.49	-

Notes: The table reports the income elasticity of demand required to rationalise equal Hicksian demand elasticities across the income distribution. Quintile 5 is treated as the reference group and assigned an income elasticity of zero. For each quintile $k \in \{1, 2, 3, 4\}$, the reported value solves $\eta_k = (\varepsilon_5^M - \varepsilon_k^M)/\bar{b}_k$, where ε_k^M is the Marshallian elasticity and $\bar{b}_k$ is the average budget share of fuel expenditure. Larger values imply that increasingly implausible income elasticities are required to reconcile the observed gradient in Marshallian elasticities with equal substitution effects across income groups.

In sum, we find that under reasonable parameterisations of the income elasticity of demand η, the demand response we estimate, as well as its gradient with respect to income, is predominantly explained by substitution rather than income effects.

Incidence of price changes

We measure the distributional incidence of a change in the price of petrol in two ways. The first measure is the budget share incidence of a price change. Namely, we take our baseline empirical specification and use it to estimate how a change in price affects petrol's budget share. The equation we estimate is given by

$$b_{it} = \alpha_i + \sum_{k=1}^{5} \left[\omega_k \left(\log(p_{it}) \times Inc_{ik}\right)\right] + \delta_{m(t)} + \lambda_{y(t)} + \xi_{it} \tag{14}$$

where all coefficients are the same as in equation 9, and we again use interactions between income quintiles and cost to instrument for the endogenous income quintile price interactions. Because equation 14 is linear, we do this using 2SLS. We report results by income quintile and pooled (i.e. without the income quintile interactions) in Table C.6. On average a 1% increase in price leads to a 0.025 percentage point increase in the share of income being spent on petrol. For the bottom income quintile this number is 0.041, roughly twice as large as for the top quintile (0.020).

Table C.6: Budget share incidence of fuel price increases by income quintile

	Pooled	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Panel A: Direct estimates						
Price effect	0.025***	0.041***	0.023***	0.023***	0.020***	0.020***
	(0.005)	(0.007)	(0.005)	(0.004)	(0.004)	(0.005)
Panel B: Estimates relative to quintile 1						
Difference from quintile 1	–	–	−0.019***	−0.019***	−0.022***	−0.022***
			(0.005)	(0.005)	(0.005)	(0.006)

Notes: The table reports estimates from regressions of the fuel expenditure share of income on fuel prices. The pooled specification estimates a common response across all households, while the quintile specifications allow responses to vary across the income distribution. Panel A reports direct estimates for each income quintile. Panel B reports differences relative to quintile 1. All specifications include account fixed effects, month-of-year fixed effects, and year fixed effects. Retail fuel prices are instrumented using wholesale fuel costs. Standard errors are reported in parentheses and are clustered two-way by account and period. Significance is given by *p<0.1; **p<0.05; ***p<0.01.

Our second measure is the compensating variation from a price change. The compensating variation (CV) is defined as the income transfer required to keep a consumer at their original level of utility following a price increase (i.e. it is an income equivalised measure of the welfare cost of changes in price). In theory the CV is the 'correct' measure to think about welfare effects, as it accounts for both the direct cost of higher prices, as well as the welfare effects of behavioural adjustments.

Formally, for a price change from p_0 to p_1, the compensating variation is defined as the area under the Hicksian demand curve between the two prices:

$$CV = \int_{p_0}^{p_1} h(p, u_0), dp \tag{15}$$

where $h(p, u_0)$ is the Hicksian demand function evaluated at the original utility level u_0. Taking a second order Taylor expansion around p_0, we can approximate the compensating variation as a share of consumer income cv as

$$cv = \frac{CV}{m} \approx b \times \left(\frac{p_1 - p_0}{p_0}\right) \times \left(1 + \frac{1}{2} \times \varepsilon^H \times \left(\frac{p_1 - p_0}{p_0}\right)\right). \tag{16}$$

The formula in 16, implies that if we know the Hicksian elasticity ε^H and petrol budget share b, then we can approximate the compensating variation from any arbitrary percentage change in the price level. Under our baseline results, we back out the Hicksian elasticity ε_i^H at the account level, using account level mean budget shares $\bar{b}_i$, the estimated Marshallian elasticity of the income group the account belongs to, and the estimated value of $\eta = 0.47$. Namely, we estimate the account-level Hicksian elasticity as

$$\varepsilon_i^H = \varepsilon_{k(i)}^M + \bar{b}_i \times \eta \tag{17}$$

Then we use the estimated budget shares $\bar{b}_i$ and Hicksian elasticities $\hat{\varepsilon}_i^H$ to estimate each accounts compensating variation following a 20% increase in price. We then aggregate to the income quintile level by taking the mean compensating variation across accounts in the quintile.

We report the results of this exercise in Table C.7. The baseline numbers imply that a 20% increase in petrol prices is equivalent to roughly a 1.7% income shock for consumers in the bottom income quintile, relative to a 0.47% shock for consumers in the top quintile. We also report these same results under alternative assumptions on the income elasticity of demand η, and show they are highly robust to this choice.[37] In the body we repeat this exercise with different changes in the price level.[38]

Table C.7: Compensating variation from a 20% increase in fuel prices

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Panel A: Baseline estimate (η = 0.47)					
Compensating variation (% income)	1.713%	0.886%	0.747%	0.618%	0.471%
Panel B: Alternative income elasticity assumptions					
Compensating variation (% income), η = 0.1	1.694%	0.882%	0.745%	0.616%	0.470%
Compensating variation (% income), η = 0.3	1.704%	0.885%	0.746%	0.617%	0.470%
Compensating variation (% income), η = 0.6	1.720%	0.888%	0.748%	0.618%	0.471%
Compensating variation (% income), η = 1.0	1.741%	0.892%	0.751%	0.620%	0.473%

Notes: The table reports compensating variation estimates from a 20% increase in fuel prices, expressed as a percentage of annual household income. Estimates are computed using the second-order approximation to the expenditure function and incorporate both income and substitution effects. Panel A reports the baseline calibration using the estimated income elasticity of demand, η = 0.47. Panel B reports robustness to alternative assumptions regarding the income elasticity of demand. The similarity of the estimates across calibrations reflects the relatively small role of income effects in explaining observed demand responses.

C.4. Heterogeneity by geography

To study how the elasticity of demand for petrol may vary by geography, we estimate the following general interaction model

$$q_{it} = \exp\left(\alpha_i + \sum_{k=1}^{K}\left[\beta_k\left(\log(p_{it}) \times G_{ik}\right)\right] + \delta_{m(t)} + \lambda_{y(t)} + \xi_{it}\right), \tag{18}$$

where G_{ik} is a dummy variable for if account i is in geographic group k. We consider a range of geographic groupings. First, we assign each account as being either in an urban region or a rural region, depending on the Greater Capital City they reside in. Second we group SA4s by their quintile of population density (1 is the most sparse, and 5 is the most dense). Then we assign accounts into a quintile of density based on the SA4 they reside in.[39] Third, we take all urban SA4s (those in Greater Sydney, Greater Brisbane or Greater Perth). We then use census data at the SA4 level to compute the proportion of residents in each SA4 who (1) commute to work via a motor vehicle, and (2) commute to work via public transport. We assign urban SA4s into terciles according to each of these metrics. Finally, we assign accounts into the terciles based on the SA4 they reside in. We report the results from all of these geographic groupings in Table C.8 below.

C.5. Excise cut and policy evaluation

In this section we outline how we simulate the effects of different policy responses to the 2026 oil price shock. In particular we consider three policy settings. The first is no policy response to the shock. The second is the cut to the fuel excise that was implemented by the Federal Government. The third is a revenue-neutral untargeted lump sum transfer. Specifically, we take the fiscal cost of the excise cut, and instead distribute it to all accounts observed making fuel purchases in the form of an equal lump sum amount. This transfer is similar in spirit to a vehicle registration rebate, as implemented by the Victorian government.

37 In particular, the choice of income elasticity enters via the imputed Hicksian elasticity ϵ^H. This elasticity is part of the second order piece of the Taylor series expansion (i.e. in equation 16 is multiplied by the square of the percentage change in price). As such, its contribution to the overall compensating variation is small relative to the budget shares.
38 It's worth noting that the percentage change in the price level enters equation 16 quadratically. This means, for instance, that the compensating variation under a 20% increase in the price level will not exactly equal twice the compensating variation under a 10% increase.
39 We get information on population density from the ABS's data by regions.

Table C.8: Geographic heterogeneity in petrol demand elasticities

	Demand elasticity estimates				
Panel A: Urban vs rural					
	Rural	Urban			
Elasticity	-0.278^{**}	-0.473^{***}			
	(0.109)	(0.142)			
Difference relative to rural	–	-0.196^{**}			
		(0.087)			
Panel B: Quintiles of population density					
	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Elasticity	-0.255^{**}	-0.280^{***}	-0.373^{***}	-0.396^{***}	-0.731^{***}
	(0.116)	(0.107)	(0.128)	(0.111)	(0.308)
Difference relative to quintile 1	–	-0.025	-0.119^{*}	-0.141^{***}	-0.476
		(0.035)	(0.065)	(0.047)	(0.293)
Panel C: Urban regions - Motor vehicle use terciles					
	Tercile 1	Tercile 2	Tercile 3		
Elasticity	-0.678^{**}	-0.404^{***}	-0.372^{**}		
	(0.334)	(0.126)	(0.159)		
Difference relative to tercile 1	–	0.274	0.306		
		(0.295)	(0.288)		
Panel D: Urban regions - Public transport use terciles					
	Tercile 1	Tercile 2	Tercile 3		
Elasticity	-0.362^{**}	-0.407^{***}	-0.724^{**}		
	(0.151)	(0.131)	(0.344)		
Difference relative to tercile 1	–	-0.044	-0.362		
		(0.063)	(0.305)		

Notes: The table reports Poisson control-function estimates of petrol demand elasticities by geographic and transport-use groups. Panel A compares rural and urban regions. Panel B reports estimates by population density quintile. Panel C reports estimates by terciles of motor vehicle use for travel to work, restricting to urban regions. Panel D reports estimates by terciles of public transport use for travel to work, restricting to urban regions. Difference rows report coefficient differences relative to the omitted reference group in each panel: rural regions in Panel A, the lowest-density quintile in Panel B, and the lowest tercile in Panels C and D. All specifications include account fixed effects, month-of-year fixed effects, and year fixed effects. Retail fuel prices are instrumented using wholesale fuel costs. Standard errors are reported in parentheses and are clustered two-way by account and period. Significance is given by *p<0.1; **p<0.05; ***p<0.01.

Fiscal cost of excise cut

Before moving to policy effects, we lay out a framework to simulate the fiscal cost of a cut to the fuel excise, as well its incidence across the income distribution. Suppose that originally, the government charges an excise of t_0 cents per litre (cpl) on each unit of fuel purchased. The retail price is p_0, and the quantity of fuel purchased and consumed is q_0. The government then cuts the excise by $\Delta t > 0$, such that the new excise is given by $t_1 = t_0 - \Delta_t < t_0$. Assume that changes to the excise pass through to retail prices at a rate of $\rho > 0$. Then the new retail price post excise cut is $p_1 = p_0 - \rho\Delta t$, and the new quantity post excise cut can be expressed as $q_1 = q_0 + \Delta q$, where Δq is the increase in consumption induced by the lower price level. The original revenue earned by the government from the excise is $R_0 = t_0 q_0$, whilst the new revenue is $R_1 = t_1 q_1$. We can write the change in revenue caused by the cut to the excise as:

$$\begin{aligned} \Delta R = R_1 - R_0 &= (t_0 - \Delta t)(q_0 + \Delta q) - t_0 q_0 \\ &= (t_0 - \Delta t)\Delta q - \Delta t q_0 \end{aligned} \tag{19}$$

The first term $(t_0 - \Delta t)\Delta q$ reflects the increase in revenue due to greater consumption under the lower price level. The magnitude of this increase depends on the size of the behavioural response Δq. The second term $\Delta t q_0$ is the revenue loss from the lower excise rate being charged on all the inframarginal units of petrol purchased (i.e. the units that would have been purchased at the old higher price level, and continue to be purchased at the new lower price level). We refer to the first component as the behavioural response component BR, and the second as the inframarginal unit component IM.

Next recall that the elasticity can be approximated by the percent change in quantity divided by the percent change in price. Therefore we can express the change in quantity induced by the excise cut as

$$\Delta q = q_0 \varepsilon \frac{\Delta p}{p_0} = -q_0 \varepsilon \frac{\rho \Delta t}{p_0}. \tag{20}$$

Substituting this into equation 19, we can express the change in revenue caused by the cut to the excise as:

$$\Delta R = (t_0 - \Delta t) q_0 \epsilon \frac{-\rho \Delta t}{p_0} - \Delta t q_0 \tag{21}$$

To simulate this change in revenue we make some additional assumptions. First we assume that there is perfect pass through in levels from changes in the excise to retail prices $\rho = 1$.[40] Second, we assume that the original quantity q_0 is given by an account's sample mean of monthly quantity, and that an account's Marshallian elasticity of demand is given by the income group the account belongs to. Under these assumptions, we can write the change in revenue from the excise cut attributable to account i as

$$\Delta R_i = BR_i + IM_i = (t_0 - \Delta t) \bar{q}_i \hat{\epsilon}^M_{k(i)} \frac{-\Delta t}{p_0} - \Delta t \bar{q}_i \tag{22}$$

Then, from the expressions above, we can derive the total fiscal cost of the excise cut by summing the cost across all accounts ($\Delta R = \sum_i \Delta R_i$). Likewise, we can get the total cost to each quintile by summing across accounts in that quintile.

$$\Delta R_k = BR_k + IM_k = \sum_{i \in k} BR_i + \sum_{i \in k} IM_i \tag{23}$$

Finally, we can recover the share of the fiscal cost attributable to each income quintile k as

$$\Delta r_k = \frac{\Delta R_k}{\sum_{k=1}^{5} \Delta R_k} = \frac{BR_k}{\sum_{k=1}^{5} \Delta R_k} + \frac{IM_k}{\sum_{k=1}^{5} \Delta R_k} \tag{24}$$

where this final expression decomposes this share by the behavioural response component and the inframarginal units component. We show these results in Table C.9.

Table C.9: Distribution of the fiscal cost of a temporary fuel excise cut

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Behavioural share (%)	−0.803	−0.889	−0.846	−0.936	−0.759
Inframarginal share (%)	12.424	16.868	20.394	23.652	30.896
Total share (%)	11.620	15.978	19.548	22.716	30.137

Notes: The table reports the distribution of the fiscal cost of a temporary fuel excise cut across the income distribution. The inframarginal component reflects revenue foregone on fuel that would have been purchased even in the absence of the tax cut. The behavioural component reflects the additional fiscal cost arising from increased fuel consumption induced by lower prices. The total share is the sum of the two components and reports the percentage of the overall fiscal cost accruing to each income quintile. Calculations assume a fuel excise reduction of 26.3 cents per litre from a pre-policy rate of 52.6 cents per litre.

Policy evaluation

We consider the effects of three policy settings: (1) no policy response to the price spike, (2) the excise cut as implemented by the Federal Government, and (3) a revenue neutral untargeted lump sum transfer. Recall the general formula for the income equivalised effect of a price change.

40 In reality, the rate of pass through ρ is endogenous to the demand environment and competitive structure of the market. Previous work provides suggestive evidence that in the past similar cuts to the excise have been fully passed through to retail prices (Maitra, 2025).

$$cv = \frac{CV}{m} \approx b \times \left(\frac{p_1 - p_0}{p_0}\right) \times \left(1 + \frac{1}{2} \times \varepsilon^H \times \left(\frac{p_1 - p_0}{p_0}\right)\right). \tag{25}$$

We simulate this at the account level using the approach laid out in equations 16 and 17, under (1) the price change induced by the full price spike amount, and (2) the price change with the cut to the excise (fully passed through).[41]

To evaluate the effects of the lump-sum transfer, we take the total fiscal cost of the excise cut ΔR, estimated above. This number then allows us to think through alternative revenue-neutral policy choices in response to the shock. The excise cut essentially 'distributes' the support out in proportion to levels of quantities purchased. We implement the untargeted lump-sum transfer by instead 'distributing' out the fiscal cost of the excise cut equally to all accounts we observe purchasing fuel. i.e. the transfer amount to account i is given by $c_i = \Delta R/N$ where N is the total number of accounts observed making fuel purchases. The income equivalised effect under this alternative policy setting is

$$cv = \frac{CV}{m} \approx b \times \left(\frac{p_1 - p_0}{p_0}\right) \times \left(1 + \frac{1}{2} \times \varepsilon^H \times \left(\frac{p_1 - p_0}{p_0}\right)\right) + \frac{c}{m}. \tag{26}$$

In each case, we compute the income equivalised effect of the shock and policy response at the account level. We then aggregate up by taking the sample mean of the effects across accounts in a given income quintile. We report the results for each income quintile in table C.10.

Table C.10: Effects of the 2026 petrol price shock and policy responses

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Welfare effects from the 2026 petrol price shock					
(A) Without any government intervention					
CV (weekly dollars)	−\$9.68	−\$13.49	−\$16.79	−\$19.64	−\$26.81
CV (as a share of income)	−3.276%	−1.701%	−1.450%	−1.201%	−0.930%
(B) With the fuel excise cut					
CV (weekly dollars)	−\$6.43	−\$8.91	−\$11.01	−\$12.86	−\$17.35
CV (as a share of income)	−2.169%	−1.123%	−0.950%	−0.786%	−0.602%
(C) With a revenue-neutral unconditional lump-sum transfer					
CV (weekly dollars)	−\$2.73	−\$6.54	−\$9.84	−\$12.69	−\$19.86
CV (as a share of income)	−0.697%	−0.814%	−0.848%	−0.773%	−0.681%

Notes: The table reports income-equivalent welfare effects of the 2026 petrol price shock and alternative policy responses across the income distribution. Panel A reports compensating variation (*CV*) in the absence of government intervention. Panel B reports welfare effects when the government responds through a temporary reduction in fuel excise. Panel C reports welfare effects when the fiscal cost of the intervention is instead redistributed through an unconditional lump-sum transfer. Dollar values are expressed in weekly terms. Share values are expressed as a percentage of household income. Compensating variation estimates are calculated using the second-order approximation to the expenditure function and an estimated income elasticity of demand of $\hat{\eta} = 0.47$.

41 Namely, for case (1) we take the average change in price between the end of February 2026 (before the outbreak of conflict) to the end of March (when prices were at their peak). For case (2) we do the same thing, but subtract the full amount of the excise decrease from this higher price, meaning that the percent increase in price is lower under case (2) than case (1) - which is the exact aim of an excise cut.

D.1. Alternative functional forms

In our baseline specification, presented in the note, we model demand in linear-log functional form and estimate the model using the Poisson quasi-maximum-likelihood estimator (PQMLE). We outline this baseline model formally in Appendix B.1.

We use this non-standard estimation method due to the 'Logs with zeros' problem (Chen & Roth, 2024). An elasticity corresponds to the percentage change in an outcome variable caused by a percentage change in the independent variable. Our outcome variable, the quantity of petrol purchased by account *i* in month *t*, takes zero values. As such, the standard approach of taking logs to estimate percentage changes selects on the event of purchase. This would induce bias in our elasticity estimate.

Estimating the model using the PQMLE is one way to circumvent this log with zeros problem. However, the non-linear form of the estimator implies we can no longer use 2SLS to instrument for an endogenous regressor, necessitating the less familiar control function approach. In addition, the estimator is likely to be less familiar to readers than more standard linear models.

In this section, we discuss other methods of estimating percentage changes in settings where the outcome variable takes zero values. We use the results from these alternative methods to discuss the robustness of our baseline approach. Our discussion in this section is heavily motivated by Chen and Roth (2024).

Level-log models

A common alternative in settings with zeros is to estimate the model in level-log form and recover an elasticity by scaling by the sample mean of the dependent variable. Namely, we estimate the following linear model

$$q_{it} = \alpha_i + \lambda_{g(t)} + \beta log(p_{it}) + \xi_{it}, \tag{1}$$

where q_{it} is the quantity in levels of petrol purchased by account *i* in time *t*. In equation 1, β gives the effect of a percentage change in price on purchase quantities in levels. Notice that via the chain rule, we can rewrite the demand elasticity ε as the below.

$$\varepsilon = \frac{\partial \log(q_{it})}{\partial q_{it}} \frac{\partial q_{it}}{\partial \log(p_{it})} = \frac{1}{q_{it}} \frac{\partial q_{it}}{\partial \log(p_{it})} \tag{2}$$

Then, we can obtain an estimate for the mean elasticity of demand based on our estimate of equation 1, and using the sample mean to estimate the expected purchase quantity.

$$\hat{\varepsilon} = \frac{1}{\bar{q}} \hat{\beta} \tag{3}$$

Table D.1 (Panel B) shows that the level-log model returns very similar estimates to the baseline. The reason we opt for the PQMLE over the level-log model comes from the restrictions imposed by the functional form assumptions. The level-log model assumes that consumers who purchase very different amounts of petrol all reduce their consumption by the same number of litres in response to a given percentage change in price. In contrast the PQMLE assumes that the quantity response is the same in percentages.[42] Whilst the results in Table D.1 suggest that in practice the difference between estimators is small, in our setting the restrictions imposed by the PQMLE appear more appropriate.

42 In other words, the level-log model estimates a model where the elasticity of demand varies along the demand curve. It estimates the elasticity by taking its value at the mean of the curve. In contrast the linear-log model (under the PQMLE) estimates a constant elasticity demand curve and hence directly estimates the average elasticity of demand.

Table D.1: Demand elasticity estimates under alternative functional forms

	Dependent variable: Quantity (litres)			
	No FE	Account FE	Account + Month FE	Account + Month + Year FE
	(1)	(2)	(3)	(4)
Panel A: Poisson QMLE control-function estimates				
Elasticity (ε)	-0.154^{**}	-0.239^{***}	-0.211^{***}	-0.383^{***}
	(0.078)	(0.075)	(0.056)	(0.118)
Panel B: Level-log estimates				
Level effect (β)	-19.396^{*}	-30.611^{***}	-27.346^{***}	-47.293^{***}
	(10.068)	(10.294)	(8.267)	(15.374)
Mean quantity ($\bar{q}$)	122.739	122.739	122.739	122.739
Elasticity ($\varepsilon = \beta/\bar{q}$)	-0.158^{*}	-0.249^{***}	-0.223^{***}	-0.385^{***}
	(0.082)	(0.084)	(0.067)	(0.125)
Panel C: Purchase probability two-part model				
Purchase probability (γ_1)	0.008	-0.012	-0.010	0.076^{***}
	(0.013)	(0.013)	(0.012)	(0.022)
Conditional quantity (γ_2)	-0.287^{***}	-0.303^{***}	-0.286^{***}	-0.393^{***}
	(0.039)	(0.044)	(0.029)	(0.078)
Mean purchase probability ($\bar{d}$)	0.678	0.678	0.678	0.678
Elasticity ($\varepsilon = \gamma_2 + \gamma_1/\bar{d}$)	-0.275^{***}	-0.321^{***}	-0.300^{***}	-0.281^{***}
	(0.043)	(0.048)	(0.033)	(0.085)
Panel D: Transaction count-size two-part model				
Transaction count (θ_1)	0.897^{***}	1.021^{***}	1.054^{***}	1.523^{***}
	(0.174)	(0.164)	(0.120)	(0.270)
Average transaction size (θ_2)	-0.532^{***}	-0.586^{***}	-0.577^{***}	-0.766^{***}
	(0.019)	(0.016)	(0.012)	(0.022)
Mean transactions ($\bar{n}$)	3.305	3.305	3.305	3.305
Elasticity ($\varepsilon = \theta_2 + \theta_1/\bar{n}$)	-0.261^{***}	-0.278^{***}	-0.258^{***}	-0.305^{***}
	(0.056)	(0.052)	(0.038)	(0.084)

Notes: This table compares estimates of the price elasticity of petrol demand under alternative functional forms. Panel A reports Poisson QMLE control-function estimates, using log(cost) as an instrument for log(price). Panel B reports level-log estimates, where the elasticity is recovered by dividing the estimated level effect by mean quantity, $\bar{q}$. Panel C reports a two-part model that decomposes demand into the probability of purchase and quantity conditional on purchase. Panel D reports an alternative two-part model that decomposes demand into the number of transactions and average transaction size. Standard errors are reported in parentheses. All specifications cluster standard errors two-way by account and period. Significance is given by *p<0.1; **p<0.05; ***p<0.01.

Two-part models

Another natural approach in settings with zeros is to model the intensive and extensive margins separately. Rather than specifying a single equation for total quantity, a two-part model separately estimates the probability of making a purchase in the period, and the quantity purchased conditional on doing so.[43] The two margins can then be combined to produce an elasticity for total quantity. To show this, we write the quantity purchased as the event that a purchase occurs multiplied by the quantity purchased conditional on purchase,

$$q_{it} \equiv d_{it} \times (q_{it} | d_{it} = 1) \tag{4}$$

Taking logs, differentiating with respect to the log of the price level, and using the chain rule, we can write the elasticity of demand as

$$\varepsilon = \frac{\partial \log(q_{it})}{\partial \log(p_{it})} = \frac{1}{d_{it}} \frac{\partial d_{it}}{\partial \log p_{it}} + \frac{\partial \log(q_{it} | d_{it} = 1)}{\partial \log p_{it}}. \tag{5}$$

To take this expression to the data, we estimate the following equations using the constructed panel of petrol purchases,

43 To simplify our discussion, in this section we treat the extensive margin as the decision to purchase or not purchase in the month, and the intensive margin as the amount purchased in the month conditional on purchase. We could equivalently treat the extensive margin as the number of transactions in the month, and the intensive margin as the average size of each transaction.

$$d_{it} = \alpha_i + \lambda_t + \gamma_1 \log(p_{it}) + \xi_{it} \quad (6)$$

$$\log(q_{it}) = \alpha_i + \lambda_t + \gamma_2 \log(p_{it}) + \xi_{it} \quad (7)$$

Then we can estimate the total elasticity of demand based on the estimates of the equations above, and using the sample mean to estimate the expected probability of purchase.

$$\hat{\varepsilon} = \frac{1}{\bar{d}}\hat{\gamma}_1 + \hat{\gamma}_2 \quad (8)$$

Table D.1 (Panel C) reports results from the two-part model. We find a negative elasticity on the intensive margin, and a positive elasticity on the extensive margin, consistent with the transaction size smoothing mechanism explained in Appendix C.2. When these extensive and intensive margin estimates are combined, the resulting elasticity estimate is substantially attenuated relative to the baseline.

This attenuation stems from the decomposition implied by equations 4 and 5. Intuitively, the two-part model estimates the extensive and intensive margin responses separately, and then adds them together. The key step in this 'combination' is scaling the extensive margin response (which is discrete) into a quantity effect (continuous). In the above, this is done by dividing $\hat{\gamma}_1$ by $\bar{d}$, the mean probability of purchase. Implicitly, this assumes that the marginal extensive margin purchaser, purchases the average quantity conditional on purchase.[44]

In practice, we observe that marginal purchasers (on the extensive margin) are likely negatively selected with respect to purchase amounts, meaning the assumption above will be violated. Table D.2 shows that accounts with lower purchase frequency also have smaller conditional purchase amounts, consistent with negative selection on the extensive margin. Under this negative selection, the decomposition in 4 over-scales the positive extensive margin response, attenuating the elasticity estimate.

Table D.2: Purchase frequency and conditional purchase quantities

	Proportion of months with a petrol purchase				
	0−0.2	0.2−0.4	0.4−0.6	0.6−0.8	0.8−1.0
Amount purchased conditional on purchase (litres)	77.326	94.488	99.041	135.982	193.097

Notes: This table reports mean petrol purchases conditional on purchasing, grouped by the proportion of months in which an account purchases petrol. The increasing relationship between purchase frequency and conditional purchase quantity is consistent with negative selection on the extensive margin: accounts that purchase less frequently also tend to purchase smaller amounts when they do purchase.

We see a similar pattern with the number of transactions − average transaction size two part model (Table D.1, Panel D). The combined elasticity is again attenuated relative to the baseline, however the degree of attenuation is smaller than under the purchase probability model. One plausible explanation for this is that the marginal transaction is less negatively selected than the marginal purchaser at the monthly level.

In contrast, the PQMLE and level-log models do not face this source of attenuation. Rather than estimating the two margins separately and combining them with an explicit scaling assumption, they directly estimate the total response accounting for zeros.[45]

Transformations of the outcome variable

Finally, in practice, a common approach to the logs with zeros problem is to apply a transformation to the outcome variable that maps zeros to a finite value, and then proceed with standard log-linear estimation. There are two commonly used transformations.

The first is the $\log(q_{it} + c)$ transformation, which adds a small constant $c > 0$ before taking logs. The issue with this transformation is that it implicitly places a weight on the extensive margin relative to the intensive margin that is determined

44 In other words, we scale up the transition from 0 to 1 (on the extensive margin), by the average value taken by the intensive margin.
45 In words, the model implicitly weights the extensive and intensive margins according to the data.

by the arbitrary choice of c (Chen & Roth, 2024). Because the log function approaches negative infinity at zero, very small choices of c imply very large changes in the transformed outcome under movements from zero to one, putting higher weight on the extensive margin response. On the other hand, because the log function is concave, larger values of c move the transformed outcome to the flatter section of the curve, attenuating the responsiveness of the outcome variable. The result is that estimates are highly sensitive to the choice of c.

The second is the inverse hyperbolic sine transformation

$$arcsinh(q_{it}) = \log\left(q_{it} + \sqrt{q_{it}^2 + 1}\right) \tag{9}$$

which approximates $\log(2q_{it})$ for large values of q_{it} but remains defined at zero. Chen and Roth (2024) show that the arcsinh coefficient converges to a weighted average of the extensive and intensive margin responses where the weights depend on the units in which q_{it} are measured. Rescaling the outcome variable - for instance converting quantities from litres to millilitres - changes the implicit weighting and hence the estimated coefficient, violating the basic requirement that an elasticity (or general percentage change) be unit free.

For these reasons, we do not use either transformation in our analysis, and instead rely on the PQMLE specification in our baseline.

D.2. Income groupings

Income buckets

In the baseline results, we group accounts into income quintiles based on their average real weekly income (in 2026 dollars). An alternative way to study differences across the income distribution is instead to annualise income amounts and look at differences across groups of annual income levels. This approach has the benefits of the group levels being more easily interpretable.

To implement this 'income buckets' approach, we take an account's average real weekly income, and use it to estimate average annual income (in 2026 dollars). Then we construct buckets based on this annual income amount. We group accounts according to the following thresholds: (1) those earning less than \$30k per year, (2) those earning between \$30k and \$50k per year, (3) those earning between \$50k and \$75k per year, (4) those earning between \$75k and \$100k per year, (5) those earning between \$100k and \$150k per year, and (6) those earning more than \$150k per year.[46]

We then reproduce the main results by income for these buckets. We show these results in Table D.3, and find that the qualitative interpretations of our baseline results are unchanged under this alternative approach. Namely, the lowest income group (those earning less than \$30,000 per year), spend a larger proportion of their budgets on petrol, are more elastic in response to price movements, and experience larger welfare losses from price spikes.

In the baseline results we find that the highest income quintile has a demand elasticity that is statistically indistinguishable from 0. The final two columns of Table D.3 suggest this is due to substantial noise in the estimates for those earning more than \$150,000 per year.

Treatment of non petrol purchasers

In the baseline results, we drop accounts who never purchase petrol. As such, all our results on heterogeneity and distributional effects should be interpreted as conditional on being both in the relevant income group and also purchasing petrol.[47] In the data, a large proportion of accounts make petrol purchases at some stage. Accounts that never purchase fuel are more likely to be in lower-income groups. This is consistent with higher-income households being more likely to own motor vehicles.

This pattern has two important implications for our baseline results. First, it mitigates the regressivity of price shocks (since a higher proportion of lower-income households have no first-order exposure to these shocks). In Table D.4 we scale the

46 Whilst not a perfect fit, buckets (1) to (4) roughly capture quintiles 1 to 4 in the baseline results. Buckets 5 and 6 split the fifth income quintile roughly in half.
47 In other words, we condition on exposure to petrol prices.

Table D.3: Income heterogeneity results by annualised income buckets

	Buckets of annual income (2026 dollars)					
	<30k	30–50k	50–75k	75–100k	100–150k	150k+
Panel A: Descriptive statistics						
Mean monthly quantity	70.3	95.0	117.9	136.2	163.4	194.4
Mean budget share	0.094	0.047	0.038	0.032	0.028	0.020
Panel B: Demand elasticity estimates						
Direct estimate	-0.640^{***}	-0.511^{***}	-0.401^{***}	-0.380^{***}	-0.237^{**}	-0.246
	(0.116)	(0.112)	(0.108)	(0.107)	(0.108)	(0.387)
Difference relative to <30k	–	0.129^{***}	0.239^{***}	0.260^{***}	0.403^{***}	0.394
		(0.039)	(0.040)	(0.042)	(0.045)	(0.366)
Panel C: Slutsky decomposition						
Hicksian elasticity (ε_H)	−0.596	−0.489	−0.383	−0.365	−0.225	−0.237
Income effect ($s_y\eta$)	−0.044	−0.022	−0.018	−0.015	−0.013	−0.009
Panel D: Budget share incidence						
Direct estimate	0.043^{***}	0.023^{***}	0.023^{***}	0.020^{***}	0.020^{***}	0.019^{***}
	(0.008)	(0.005)	(0.004)	(0.004)	(0.004)	(0.006)
Difference relative to <30k	–	-0.020^{***}	-0.021^{***}	-0.023^{***}	-0.023^{***}	-0.024^{***}
		(0.006)	(0.006)	(0.006)	(0.006)	(0.007)
Panel E: Welfare effects and policy responses						
2026 price shock (% income)	−3.413	−1.723	−1.432	−1.190	−1.065	−0.757
With excise cut (% income)	−2.263	−1.139	−0.939	−0.779	−0.689	−0.490
With lump-sum transfer (% income)	−0.686	−0.803	−0.840	−0.768	−0.761	−0.574

Notes: The table reports robustness results using income buckets rather than income quintiles. Panel A reports descriptive statistics by income bucket. Panel B reports Poisson control-function estimates of petrol demand elasticities and differences relative to the lowest income bucket. Panel C reports the Slutsky decomposition using the estimated income elasticity of demand, $\hat{\eta}$ = 0.47. Panel D reports budget share incidence estimates and differences relative to the lowest income bucket. Panel E reports income-equivalent welfare effects of the 2026 petrol price shock and alternative policy responses, expressed as a percentage of income. Standard errors are reported in parentheses. Significance is given by *p<0.1; **p<0.05; ***p<0.01.

compensating variation from a 20% price increase by the proportion of non-purchasers. Under this scaling, the price increase remains regressive, but is less so than in the baseline case where we condition on purchasing.

Second, it means that the revenue-neutral untargeted lump sum transfer (whilst equal in amounts for all drivers) is not equal in amounts across the population. Namely, higher-income groups receive more than lower-income groups on average, because they are more likely to be a fuel purchaser. This is why we describe the transfer as being similar in spirit to a car registration rebate, rather than an untargeted transfer implemented for instance through the tax and transfer system.

Table D.4: Treatment of accounts not purchasing fuel

	<30k	30–50k	50–75k	75–100k	100–150k	150k+
Panel A: Petrol-purchase exposure						
Share purchasing petrol	79.6%	89.5%	92.4%	92.7%	92.1%	89.6%
Panel B: Welfare effects of a 20% price increase						
CV, conditional on petrol purchasers	−1.79%	−0.90%	−0.74%	−0.61%	−0.54%	−0.38%
CV, population average	−1.42%	−0.80%	−0.68%	−0.57%	−0.50%	−0.34%

Notes: The table reports how accounting for households with no observed petrol purchases changes the distributional incidence of petrol price shocks. Panel A reports the share of accounts in each income bucket that purchase petrol at least once. Panel B reports compensating variation from a 20% petrol price increase, first conditional on petrol purchasers and then scaled by the share of accounts purchasing petrol to obtain a population-average effect.

D.3. Income and geography interactions

In Table D.5, we present results splitting by both income and geography. Within both urban and rural areas, we see the same pattern of results by income as found in the baseline. Within an income group, consumers in rural regions tend to spend a larger budget share on petrol, be less elastic in their demand response, and suffer larger welfare costs from higher prices. Together these results (unsurprisingly) suggest that lower-income consumers in rural areas are most adversely affected by large price spikes.

Table D.5: Heterogeneity by geography and income

	<30k	30–50k	50–75k	75–100k	100–150k	150k+
Panel 1: Urban regions						
Panel A: Descriptive statistics						
Mean monthly quantity	69.2	92.7	113.4	127.6	150.1	175.5
Mean budget share	0.095	0.046	0.037	0.030	0.025	0.018
Panel B: Demand elasticity estimates						
Direct estimate	−0.644***	−0.561***	−0.442***	−0.443***	−0.290**	−0.434
	(0.140)	(0.136)	(0.129)	(0.129)	(0.132)	(0.593)
Difference relative to <30k	–	0.083	0.202***	0.201***	0.354***	0.210
		(0.052)	(0.053)	(0.055)	(0.056)	(0.548)
Panel C: Slutsky decomposition						
Hicksian elasticity (ε_H)	−0.600	−0.539	−0.425	−0.429	−0.278	−0.426
Income effect ($s_y\eta$)	−0.045	−0.022	−0.017	−0.014	−0.012	−0.008
Panel D: Budget share incidence						
Direct estimate	0.042***	0.020***	0.020***	0.017***	0.017***	0.016*
	(0.008)	(0.005)	(0.004)	(0.005)	(0.004)	(0.008)
Difference relative to <30k	–	−0.022***	−0.022***	−0.025***	−0.025***	−0.026***
		(0.007)	(0.007)	(0.007)	(0.007)	(0.010)
Panel E: Welfare effects and policy responses						
2026 price shock (% income)	−3.445	−1.666	−1.366	−1.099	−0.966	−0.653
With excise cut (% income)	−2.285	−1.105	−0.898	−0.723	−0.628	−0.430
With lump-sum transfer (% income)	−0.619	−0.746	−0.774	−0.677	−0.663	−0.472
Panel 2: Rural regions						
Panel A: Descriptive statistics						
Mean monthly quantity	71.7	97.3	123.2	147.3	181.5	225.0
Mean budget share	0.093	0.048	0.040	0.034	0.031	0.023
Panel B: Demand elasticity estimates						
Direct estimate	−0.663***	−0.460***	−0.367***	−0.317***	−0.182*	0.102
	(0.113)	(0.105)	(0.102)	(0.102)	(0.101)	(0.120)
Difference relative to <30k	–	0.204***	0.296***	0.346***	0.481***	0.765***
		(0.065)	(0.066)	(0.066)	(0.070)	(0.100)
Panel C: Slutsky decomposition						
Hicksian elasticity (ε_H)	−0.620	−0.437	−0.349	−0.301	−0.168	0.113
Income effect ($s_y\eta$)	−0.044	−0.023	−0.019	−0.016	−0.014	−0.011
Panel D: Budget share incidence						
Direct estimate	0.048***	0.026***	0.025***	0.023***	0.024***	0.023***
	(0.011)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Difference relative to <30k	–	−0.022**	−0.023**	−0.025**	−0.024**	−0.025**
		(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Panel E: Welfare effects and policy responses						
2026 price shock (% income)	−3.353	−1.786	−1.507	−1.306	−1.198	−0.946
With excise cut (% income)	−2.228	−1.175	−0.985	−0.850	−0.772	−0.597
With lump-sum transfer (% income)	−0.742	−0.864	−0.912	−0.883	−0.894	−0.761

Notes: The table reports robustness results using income buckets interacted with geography. Panel 1 reports estimates for urban regions and Panel 2 reports estimates for rural regions. Within each geography, Panel A reports descriptive statistics by income bucket. Panel B reports Poisson control-function estimates of petrol demand elasticities and differences relative to the lowest income bucket. Panel C reports the Slutsky decomposition using the estimated income elasticity of demand, $\hat{\eta} = 0.47$. Panel D reports budget share incidence estimates and differences relative to the lowest income bucket. Panel E reports income-equivalent welfare effects of the 2026 petrol price shock and alternative policy responses, expressed as a percentage of income. Standard errors are reported in parentheses. Significance is given by *p<0.1; **p<0.05; ***p<0.01.

A large empirical literature has sought to estimate the price elasticity of demand for petrol. Table E.1 summarises estimates from a selection of studies in this literature. The older generation of studies, generally using aggregate time series data, tend to find highly inelastic demand, with short-run elasticity estimates in the range of −0.05 to −0.15. Contrastingly, recent studies using more granular data find substantially larger elasticities, in the range of −0.3 to −0.4.

Table E.1: Petrol demand elasticity estimates from the literature

Paper	Location	Sample Period	Data Type	Elasticity Estimate
Breunig and Gisz (2009)	Australia	1966-2006	National - Quarterly	−0.13
Park and Zhao (2010)	US	1976-2008	National - Monthly	−0.2
Hughes et al. (2008)	US	1974-2006	National - Monthly	−0.03 to −0.08
Li et al. (2014)	US	1966-2008	State - Year	−0.192
Kilian and Zhou (2024)	US	1989-2008	State - Monthly	−0.3
Levin et al. (2017)	US	2006-2009	City - Daily	−0.35
Knittel and Tanaka (2021)	Japan	2005-2014	Individual - Daily	−0.37

Notes: Table reports demand elasticity estimates from the literature. Results are organised by levels of spatial and time aggregation.

Levin et al. (2017) provide an in-depth explanation of why more highly aggregated data tends to lead to attenuated elasticity estimates. First, higher levels of aggregation make it harder to include sufficient controls for demand shocks over time. For instance, over long periods of time, it might be true that growth in both prices and quantities comes from changes in demand. With sufficiently rich data, one can control for these shifts with time fixed effects, but this is harder to do if the data is at the yearly or quarterly level. Second, if the identifying variation comes from short run fluctuations in price, or differences in price across locations, then more highly aggregated data can wash out this source of identification.[48]

We find a price elasticity of demand of −0.38, larger than previous studies, but close to recent estimates using comparable data. The gap with the most recent Australian estimate from (Breunig & Gisz, 2009) reflects a combination of factors. The most important are likely structural change, consumers today may plausibly have more viable alternatives to driving than they did decades ago, and differences in data and identifying variation. Breunig and Gisz (2009) identify their elasticity using long-run time series variation in aggregate prices and quantities, making it more difficult to control for shocks to demand over time.

Our results do provide some evidence of the aggregation challenges documented by Levin et al. (2017). When we drop year-of-sample fixed effects from our specification, the estimated elasticity attenuates towards zero, consistent with the interpretation that over longer time horizons, common trends in prices and quantity may be more likely to reflect demand than supply shocks. However, because our sample period is dominated by large, exogenous shocks to global oil markets (most notably Russia's invasion of Ukraine in 2022), the long run demand trends that concern Levin et al. (2017) are less of a threat to our identification than they might be in a quieter sample period.

Finally, very few studies in the existing literature flexibly estimate how the price elasticity of demand varies across the income distribution. Kilian and Zhou (2024) separate US states into above and below median income groups, and find that below median income states have more elastic demand. This is directionally consistent with our results.

48 Levin et al. (2017) formalise these arguments.