Disentangling sign and size non-linearities[†]

Tomás E. Caravello[‡] Pedro Martínez Bruera[§]
MIT MIT

June 2024

Abstract: We study the identification of non-linear causal effects of macroeconomic shocks using local projections augmented by a non-linear function of the shock of interest. Our analysis focuses on two types of non-linearities: by size and sign of the shock. We characterize the estimand of our non-linear local projections. Our main result is that the local projections identify pure sign non-linearities when they include an even non-linear function of the shock and vice-versa for size non-linearities and odd functions. We illustrate our method with an application to oil supply news shocks, documenting evidence in favor of size (but not sign) non-linearities.

Keywords: Local Projection, non-linear impulse responses, sign non-linearity, size non-linearity, asymmetry, oil shocks, non-parametric model.
JEL codes: C14, C22, C32, C52

[†]We especially thank Christian K. Wolf for helpful comments, suggestions and overall guidance. For helpful comments and suggestions, we also thank Fernando Alvarez, Isaiah Andrews, Javier García-Cicco, Jackson Mejia, Anna Mikusheva, Victor Orestes, Mikkel Plagborg-Møller, Ashesh Rambachan, Martin Solá and Iván Werning, as well as seminar participants at Universidad Torcuato di Tella and Universidad de San Andrés.

[‡]tomasec@mit.edu

[§]pmbruera@mit.edu

1 Introduction

Traditionally, researchers in the applied macroeconomics literature have worked under the assumption of linearity of the data-generating process. However, there is an extensive body of work that documents the presence of non-linearities in macroeconomic data. The importance of detecting the presence of non-linearities goes beyond just improving statistical fit for two main reasons. First, the presence of non-linearities is of first-order concern for policy design. For example, the familiar metaphor of "pushing on a string" theorizes that monetary policy is effective to cool down economic activity, but not to stimulate it. Understanding whether the empirical evidence backs this theory provides guidance to policymakers. Detecting potential non-linearities in other shocks (for example, oil shocks) is also relevant to stabilization policy design, in order to properly react to those shocks. Second, disentangling different types of non-linearities in the data has implications for structural modeling. Properly identified non-linear impulse response functions could serve as identified moments that can help to distinguish between different classes of models, in the spirit of Nakamura and Steinsson (2018).

Several papers have documented evidence of the non-linear propagation of shocks related to the size or sign of the shock. Recent examples include government spending shocks (Barnichon et al., 2022a; Ben Zeev et al., 2023; Brinca et al., 2023), monetary policy shocks (Barnichon and Brownlees, 2019; Ben Zeev, 2020; Ascari and Haber, 2021), and risk premium shocks (Barnichon et al., 2022b; Forni et al., 2023). An usual procedure in the literature is to estimate a local projection augmented by a non-linear transformation of the shock of interest. While this procedure is simple to apply, it is unclear how much the results depend on the structural form assumptions. If the functional form is misspecified then there can be substantial bias. This observation leads to the following question: what can we learn from these types of regressions? Can we test for the presence of non-linearities in a robust manner?

In this paper, we fill this gap and provide an answer to both of these questions without making any functional form assumptions on the data-generating process. We study two types of non-linearities: by sign and by size. Sign non-linearities arise if the response of the outcome of interest depends on the sign of the shock, whereas size non-linearities arise if the response of the outcome of interest depends on the absolute magnitude of the shock. We formally define these concepts when the data-generating process is arbitrary. Our definition exploits that any function can be decomposed into the sum of an even function and an odd function.

We argue that this decomposition captures sign and size non-linearities respectively.

The main contribution of this paper is a separation result: under the assumption of a symmetric distribution of the shock, it is possible to test for the presence of sign and size non-linearities separately by estimating a local projection augmented by a non-linear transformation of the shock of interest. If the non-linear transformation is an even function then the estimand only captures sign non-linearities; whereas if the non-linear transformation is an odd function then the estimand captures size non-linearities. This result does not rely on any parametric assumptions on the data-generating process. Our result leverages on an insight from Angrist et al. (2000): under an arbitrary data generating process the population estimates are weighted averages of the marginal effects, regardless of the exact functional form assumption. Changes in the functional form of the non-linear transformation only affect the implicit weighting scheme. Within this framework, we show that the researcher can capture only sign or size non-linearities by appropriately choosing the non-linear transformation of the shock of interest.

We show that if the underlying distribution is not symmetric, then specifications previously thought to capture only sign or size non-linearities actually mix both kinds of non-linearities.[1] Intuitively, if the distribution of the shock is not symmetric then, for example, positive shocks are also on average larger. Therefore, the econometrician is not sure whether the differential effect of positive shocks is due to their sign or size. We propose a method to assess how much the asymmetry affects the results of typical specifications used in the literature.

A secondary contribution has to do with the causal interpretation of the magnitudes. We argue that using piecewise linear functions (such as $|x|$) as the nonlinear transformation is recommended over other types of non-linear transformations (such as x^2).[2] Piecewise linear functions simplify reporting and interpretation: since all our estimands are weighted averages, the estimated coefficients are directly related to the effects of, for example, positive and negative shocks. We argue that the alternative approach (using another non-linear function and then choosing a size of the shock to report, traditionally a one standard deviation shock) may be misleading: the implicit weighting scheme may put large weight in values far from one standard deviation.

[1] For example, by using $\max(x, 0)$ and $\min(x, 0)$ to allow for different coefficients for positive and negative values of the shock as in Barnichon et al. (2022a).

[2] Importantly, using $|x|$ instead x^2 does not add any implicit assumptions about the smoothness of the underlying DGP once we interpret results in a non-parametric way.

In short, there are three main takeaways. First, if the researcher is interested in non-linearities by sign, she should fit a local projection adding an even transformation of the shock of interest, in particular $|x|$. Second, if the researcher is interested in non-linearities by size, she should use an odd function as the transformation. Although x^3 is an option, we argue that using a piecewise linear function is preferable. In both cases, the researcher then proceeds to test if the coefficients on the non-linear terms are different from zero. Third, echoing Plagborg-Møller (2023), one should always report the implicit weighting scheme implicit in the regression. Furthermore, reporting weights is crucial to assess how much departures from symmetry affect the interpretation of results.

We also present a simulation exercise to explore the finite sample properties of this procedure. First, if the data-generating process is linear, the size of the test lines up with the theoretical level. Second, we consider the case where the data-generating process has only one type of non-linearity (either size or sign). Our simulations suggest that all functions within the class of odd or even functions that we consider have similar power for testing for the presence of size or sign non-linearities. Thus, getting the exact functional form right is not crucial for power. For instance, if the true data generating process has $|x|$ as the non-linear term then fitting a linear projection using x^2 has roughly the same power as fitting the correct model. Finally, we show that if one tests for one kind of non-linearity but only the other one is present in the data, non-trivial size distortions appear when using polynomials of high order, such as x^3. Size distortions are much smaller when using piecewise linear functions, which strengthens the argument for their usage.

Finally, we present a novel empirical application using data from Känzig (2021) to address the long-standing question about the presence of non-linear responses in oil shocks (Mork, 1989; Hamilton, 2003; Kilian and Vigfusson, 2011). Our results indicate the presence of non-linearities by size, but we show little evidence of non-linearities by sign. Large shocks have a proportionally larger impact in industrial production and trigger a stronger monetary policy response. In contrast, when comparing positive and negative shocks, the responses are quite similar. Further exploration of the underlying mechanisms for this non-linearity is left as an interesting avenue for future research.

The paper relates to several strands of literature. First, the paper builds on the pioneering work of Rambachan and Shephard (2021), which bring non-parametric interpretation of estimands (as in Angrist et al. (2000)) to the time series context.[3] Thus, we relate to the literature on using non-parametric identification tools to interpret estimands of linear

[3]See also Plagborg-Møller (2023).

regressions.

Second, recent literature estimates non-linearities in the response to externally identified macroeconomic shocks. Barnichon and Matthes (2019) and Ben Zeev (2020) both focus on sign dependence for monetary policy shocks. Ascari and Haber (2021) finds size dependence for monetary policy shock, consistent with menu cost models. Barnichon et al. (2022b) and Forni et al. (2023) focus on sign and size risk premium shocks. Auerbach and Gorodnichenko (2012), Ramey and Zubairy (2018), Ben Zeev et al. (2023), Barnichon et al. (2022a) look at non-linearities in government spending shocks.[4] There is also a vast, long-standing literature that estimates non-linearities in macroeconomic relationships without focusing on externally identified shocks. For example, Ravn and Sola (2004) uses a regime-switching model to estimate non-linear responses to monetary policy in the US, focusing on sign and size non-linearities.

This paper is especially related to Gonçalves et al. (2021) and Gonçalves et al. (2024). Gonçalves et al. (2021) proposes a particular class of non-linear DGPs, focuses on identifying unconditional IRFs for a given size of the shock, and shows what kinds of estimators correctly recover this object of interest. Instead, we define the two types of non-linearities of interest for a more general class of DGPs and show how we can identify them separately using local projections. Thus, we view our results as complementary to the ones in Gonçalves et al. (2021): we rely on weaker assumptions about the DGP because we are interested in identifying something less demanding than the unconditional IRF. Put differently, Gonçalves et al. (2021) defines an object of interest and shows what estimators recover it. Instead, we take a standard procedure in the literature (local projections with non-linear terms) and ask: i) what does this procedure identify? ii) when is this informative about sign and size non-linearities? In this regard, our approach is closer to Gonçalves et al. (2024). That paper focuses on state-dependent impulse responses and shows how different estimation methods identify different estimands. In particular, it shows that, as opposed to conventional VAR and LPs, state-dependent VARs and LPs do not recover the same estimand, and therefore applied researchers should carefully pick between the two depending on what is their estimand of interests. Our paper focuses on sign and size non-linearities instead of state-dependence and asks how different estimands are related to the presence of either kind of non-linearity in the data.

The paper is also related to Ben Zeev (2020) and Forni et al. (2023). Ben Zeev (2020) shows that there is bias in estimating a specification with x and $|x|$ as regressors if the true

[4]We thank Valerie A. Ramey for kindly providing an early draft of Ben Zeev et al. (2023).

DGP actually features quadratic terms. Forni et al. (2023) proposes a horse race between regressions including $|x|$, x^2 and both to test for size vs. sign asymmetry. Since our work does not assume any particular functional form for the DGP and we also present results for size asymmetries, the present paper extends, generalizes and corrects the results contained in these two papers. In particular, in contrast to Forni et al. (2023), we show that as $|x|$ and x^2 are both even functions, they both capture sign asymmetry only, with no power to test for size non-linearities. This point is also recognized in Tenreyro and Thwaites (2016), who uses a cubic term in order to test for size non-linearities, and Ascari and Haber (2021), who use $x|x|$.

Our paper is also related to the literature estimating non-parametric impulse response functions. Recent examples in time series include Gourieroux and Lee (2023) and Gonçalves et al. (2024). Relative to that literature, we focus on the non-parametric interpretation of typical estimands rather than full non-parametric estimation of objects of interest. Given that time series data is usually limited, full non-parametric estimation can be unpractical, compared with the simpler local projections estimates that are the focus of this paper.

The empirical application is related to a large body of literature that studies the impact of oil shocks on the economy, pioneered by Hamilton (1983). Känzig (2021) is a recent important contribution: that paper applies high-frequency identification to oil prices around OPEC meetings in order to obtain plausibly exogenous changes in oil prices, and shows that the macroeconomic impact can be relevant. Researchers have looked at potential non-linearities in the macroeconomic response to oil shocks. Early contributors point to an asymmetric response on output, employment (Mork (1989), Davis and Haltiwanger (2001), Balke et al. (2002), Hamilton (2003), Hamilton (2011)). However, the evidence for asymmetry has also been questioned (Kilian and Vigfusson (2011), Herrera et al. (2011), Herrera et al. (2015), Kilian and Zhou (2020)). In particular, Kilian and Vigfusson (2011) stress how the particular non-linear transformations used could generate spurious evidence of non-linearity if the true underlying DGP was linear. Misspecification concerns are at the core of our paper, thus addressing the points raised in Kilian and Vigfusson (2011) in a general way. As in Kilian and Vigfusson (2011), Herrera et al. (2011), Herrera et al. (2015), we find little evidence of sign non-linearities. However, we find significant evidence of size non-linearities.

Section 2 lays out the general framework used throughout the paper. It also presents the estimands of interest and the weighted average representation of such estimands, which is the key tool used throughout the paper. Section 3 presents a simple example to build intuition and preview the mechanics of the main separation result. Section 4 is the core of the paper: it states and proves the main separation result. It also discusses what happens if the symmetry assumption does not hold, the causal interpretation of the estimands, and presents practical takeaways. Section 5 presents evidence on the finite sample performance of the procedure in simulations. Section 6 describes the empirical application. Section 7 concludes.

2 Econometric Framework

2.1 Model

Let $\mathbf{y}_t \in \mathbb{R}^{n_Y}$ be a vector of outcome variables of interest at time t, for example inflation and GDP growth. Let $\{\boldsymbol{\varepsilon}_t\}_{t=-\infty}^{\infty}$ be an n_ε-dimensional process of mean zero structural shocks. Elements of this vector can include monetary policy shocks and oil price shocks. We specify restrictions on this process later on. Let $g : \varkappa \to \mathbb{R}^{n_Y}$ be a structural function that maps sequences of shocks to outcomes.[5]

Assumption 1 (General DGP)**.** *The n_y-dimensional vector $\boldsymbol{y}_t$ is a function of the realization of the (present and past) n_ε-dimensional vector of structural shocks*

$$\boldsymbol{y}_t = g\left(\{\boldsymbol{\varepsilon}_{t-i}\}_{i=0}^{\infty}\right) \tag{1}$$

Assumption 1 is a generalization of the linear SVMA representation assumed in the literature. Virtually all macro models can be represented as (1). If we couple Assumption 1 with linearity in g we obtain the usual SVMA representation: if g is linear then (1) becomes $\boldsymbol{y}_t = \sum_{i=0}^{\infty} \boldsymbol{\Theta}_i \boldsymbol{\varepsilon}_{t-i}$, where $\boldsymbol{\Theta}_j$ are $n_y \times n_\varepsilon$ matrices.

Since our estimands of interest are functions of unconditional moments of the data, we assume stationarity and ergodicity.[6][7]

Assumption 2 (Stationarity and erogdicity)**.** *The processes $(\boldsymbol{y}_t, \boldsymbol{\varepsilon}_t)$ is strictly stationary and ergodic.*

[5] Where $\varkappa$ is the set of two-sided (deterministic) sequences with typical element $\{\boldsymbol{x}_t\}_{t=-\infty}^{\infty}$, where $x_t \in \mathbb{R}$ for all t. Thus, the structural function $g(\cdot)$ maps any potential realization of $\{\boldsymbol{\varepsilon}_t\}_{t=-\infty}^{\infty}$ into outcomes.

[6] We impose these as high-level assumptions. Given the little structure we impose, providing general conditions on $g(\cdot)$ that imply stationarity and ergodicity is beyond the scope of this paper.

[7] A drawback of this assumption is that it rules out conditional heteroskedasticity of the shocks.

We now turn to assumptions about the underlying causal structure of the economy. Let $\varepsilon_{1,t} = \varepsilon_t$ be the shock of interest. We assume that the shock of interest ε_t is independent of all other shocks at all leads and lags, as well as lags and leads of itself.

Assumption 3 (Independence)**.** $\varepsilon_{1,t} \perp \boldsymbol{\varepsilon}_{-1,t}, \{\boldsymbol{\varepsilon}_{t+j}\}_{j\in\mathbb{Z}-\{0\}}$.

Assumption 3 parallels two assumptions made in linear models. First, $\boldsymbol{\varepsilon}_t$ is uncorrelated over time. Second, the elements of $\boldsymbol{\varepsilon}_t$ are mutually uncorrelated. Under linearity these assumptions are sufficient. In contrast, if we want to allow for arbitrary non-linearities we need to assume independence. However, the underlying economics of these assumptions are the same as in the linear case.

We make the following technical assumptions.

Assumption 4 (Regularity)**.** *The unconditional distribution of ε_t has density $F'(x)$ and supp$(\varepsilon_t) \subset [-\bar{\varepsilon}, \bar{\varepsilon}]$.*

The bounded support assumption is for expositional simplicity. The continuous distribution of the shock is somewhat stronger than what we need: the main result goes through if there is a mass point at $\varepsilon_t = 0$. However, it cannot be fully dispensed, see Appendix A.2.

The object of interest of the paper is potential non-linearities in the causal effect of the shock of interest $\varepsilon_{1,t}$ on the outcome h periods ahead in the future $\mathbf{y}_{t+h}$. We focus on a single outcome $y_{j,t}$. From (1) we have

$$y_{j,t+h} = g^j(\underbrace{\{\boldsymbol{\varepsilon}_{t+i}\}_{i=1}^{h}}_{\text{shocks after } t}, \varepsilon_{1,t}, \underbrace{\boldsymbol{\varepsilon}_{-1,t},}_{\text{other shocks at } t} \underbrace{\{\boldsymbol{\varepsilon}_{t-j}\}_{i=0}^{\infty}}_{\text{lags}}) \tag{2}$$

For notational simplicity, we suppress explicit mention of shocks after t, lags, and shocks at time t other than the shock of interest. For a given value of the shock of interest $\varepsilon_{1,t} = a$, define $g^{j,h}(a, \bullet)$ as.

$$g^{j,h}(a, \bullet) = g^j\left(\{\boldsymbol{\varepsilon}_{t+i}\}_{i=1}^{h}, a, \boldsymbol{\varepsilon}_{-1,t}, \{\boldsymbol{\varepsilon}_{t-j}\}_{i=1}^{\infty}\right)$$

The partial derivatives of the function $g^{j,h}(\varepsilon_t, \bullet)$ are a key object of our analysis. Fixing an outcome of interest j and horizon h define

$$g_\varepsilon(a, \bullet) \equiv \frac{\partial}{\partial \varepsilon_{1,t}} g^{j,h}(\varepsilon_{1,t}, \bullet)|_{\varepsilon_{1,t}=a} \tag{3}$$

as the marginal effect of shock ε_t on $y_{j,t+h}$, evaluated at $\varepsilon_t = a$. We also assume this derivative always exists, except for a set of measure zero.

Assumption 5 (Differentiability)**.** *$g_\varepsilon(a, \bullet)$ is well defined for all $a \in [\underline{\varepsilon}, \overline{\varepsilon}]$ except for a finite (possibly empty) set.*

As we show later in the paper, our estimands depend on $g_\varepsilon(a, \bullet)$, which is why we need it to exist. This assumption allows the derivative to be not defined at some points if the DGP has kinks: we only require that the number of kinks is finite. Functions with kinks are frequently assumed in the literature. A recent example is Barnichon et al. (2022a).

2.2 Estimands of Interest

We assume the shock of interest, ε_t, to be observed by the econometrician. In practice, this means that the econometrician has access to an external measure of the shock, as typically assumed in the proxy-SVAR literature. For example, in our empirical application we use the shock identified in Känzig (2021) as a measure of oil supply news shocks.[8]

We study local projections (Jorda, 2005) of the form

$$y_{j,t+h} = \alpha + \beta_1 \varepsilon_t + \beta_2 f(\varepsilon_t) + u_{j,t+h} \tag{4}$$

where $f(\cdot)$ is a known function proposed by the researcher. The estimands of interest are β_1 and β_2, the coefficients of the linear projection.[9] This approach has been extensively used in the literature (see, e.g. Ben Zeev et al. (2023)).[10] Some specifications used in the literature are equivalent to (4), for example using $\min(x, 0)$ and $\max(x, 0)$ as in Barnichon et al. (2022a).[11]

This paper is concerned with two types of non-linearities[12]:

(i) *Sign non-linearity.* There is a sign dependence if the magnitude of the effect depends on the sign of the shock. For example, consider the effect of government spending

[8]Alternatively, one could also impose additional assumptions about the DGP to recover the shock, as in Gonçalves et al. (2021) or Forni et al. (2023).

[9]Notice that this is not a definition of the DGP, which is given by (1), but rather a definition of the estimands of interest β_1, β_2.

[10]Empirical implementations assume different functional forms for $f(\cdot)$ and then analyze the sign of β_2 to test for the presence of nonlinearities. For example, Tenreyro and Thwaites (2016) uses $f(x) = x^3$ to test for size non-linearities in the context of monetary policy shocks. Ben Zeev (2020) compares $f(x) = |x|$ and $f(x) = x^2$.

[11]Note that the fit of such a regression is the same as using x and $|x|$, since $(\min(x, 0), \max(x, 0)) = \left(\frac{x-|x|}{2}, \frac{x+|x|}{2}\right)$

[12]We do not focus on identifying state-dependence in causal effects. However, note that the general DGP (1) allows for general state-dependence.

on output. There is sign dependence if the effect on output of increasing government spending by one percentage point is not equal to the effect on output of decreasing government spending by one percentage point.

(ii) *Size non-linearity.* There is size dependence if the proportional effect, i.e., $(\mathbb{E}[y_{t+h}|\,\varepsilon_t = a] - \mathbb{E}[y_{t+h}|\,\varepsilon_t = 0])/a$, depends on the size of the shock, a. Turning back to the example of government spending, there would be size dependence if the total effect of increasing the government spending by two percentage points is *different* from two times the effect of increasing government spending by one percentage point. Put it differently, large and small shocks (in absolute magnitude) have different (scaled) effects.

Importantly, the DGP assumed in (1) allows also for more general kinds of non-linearities, such as state dependence. Those are not the subject of study of the present paper, see Gonçalves et al. (2024) for a treatment on the subject.

2.3 Estimands as weighted averages

The main results of this paper build on the following extension of a well-known result due to (Yitzhaki, 1996) and (Angrist et al., 2000).

Lemma 1. *Under Assumptions 1-5, the β_1 and β_2 are given by*

$$\beta_i = \int_{\underline{\varepsilon}}^{\overline{\varepsilon}} \omega_i(a)\mathbb{E}\left[g_\varepsilon\left(a, \bullet\right)\right] da$$

where

$$\omega_1(a) = \frac{Cov\left(\mathbf{1}_{\{a \leqslant \varepsilon_t\}}, \varepsilon_t\right)}{Var\left(\varepsilon_t\right)}, \qquad \omega_2(a) = \frac{Cov\left(\mathbf{1}_{\{a \leqslant \varepsilon_t\}}, f^{\perp}\left(\varepsilon_t\right)\right)}{Var\left(f^{\perp}(\varepsilon_t)\right)}$$

and

$$\int_{\underline{\varepsilon}}^{\overline{\varepsilon}} \omega_1(a)da = 1, \qquad \int_{\underline{\varepsilon}}^{\overline{\varepsilon}} \omega_2(a)da = 0$$

where $f^{\perp}(x) = f(x) - \frac{Cov(\varepsilon_t, f(\varepsilon_t))}{Var(\varepsilon_t)}x$.

The proof is shown in Appendix A.1. Lemma 1 echoes the result in Yitzhaki (1996), Angrist et al. (2000), Rambachan and Shephard (2021): the estimands are weighted averages of the derivatives of the g function, $\mathbb{E}\left[g_\varepsilon\left(a, \bullet\right)\right]$.[13] That property extends to multiple

[13]In full rigor, $\mathbb{E}\left[g_\varepsilon\left(a, \bullet\right)\right]$ denotes the *average marginal effect* of shock ε_t at a value a. The "average" part comes from the fact that we first compute the partial derivative $g_\varepsilon\left(a, \bullet\right)$ for all possible values of other variables summarized in $\bullet$. The expectation averages across those, maintaining $\varepsilon_t = a$ fixed.

regressions, where we use different non-linear transformations of the shock. Lemma 1 shows that the weights can be recovered from the distribution of ε_t and the functional form of f, and that they do not depend on the structural function $g(\cdot)$. In addition, the weights integrate to zero and one respectively. Therefore, if the DGP is linear then $\beta_2 = 0$:

Corollary 1. *Assume that $\mathbb{E}[g_\varepsilon(a, \bullet)]$ is a constant function of a. Then $\beta_2 = 0$.*

Corollary 1 restates that in terms of weights: since a linear term in $g(\cdot)$ will show up as a constant in the derivative, all non-linear regressors will be unaffected by that: as the weights integrate to zero, constant terms cancel out. Thus, β_1 captures some *average* of the marginal effects, whereas β_2 captures how marginal effects differ for different values of ε_t.

It is important to note that the weights $\omega_2(a)$ will be negative for some values of a, and positive for others. As we show later on, this is a feature: by putting positive and negative weight in different values of the shock, the term β_2 captures different kinds of non-linearities.

3 Example and preview of results

In order to illustrate the mechanics behind the non-linear local projections, we present a simple model with a non-linear Phillips Curve. Readers interested in the general identification result may proceed to Section 4.

Example 1 (NK model with Non-linear Phillips curve)**.** *Let y_t denote the output gap, and π_t to be inflation. Consider a version of the standard New Keynesian model with a non-linear Phillips curve given by:*

$$\begin{aligned} y_t =& \varepsilon_t^d \\ \pi_t =& c\left(y_t\right) + \beta \mathbb{E}_t\left[\pi_{t+1}\right] + \varepsilon_t^s \end{aligned}$$

where $\varepsilon_t^d \sim \mathcal{U}\left(-b, b\right)$, $\mathbb{E}[\varepsilon_t^s] = 0$, both $\varepsilon_t^d, \varepsilon_t^s$ are iid, and $c(y_t) = \begin{cases} ky_t & y_t \geqslant -\frac{\gamma}{\kappa} \\ -\gamma & y_t < -\frac{\gamma}{\kappa} \end{cases}$ with $\gamma \geqslant 0$. Note also that since the economy is iid, $\mathbb{E}_t\left[\pi_{t+1}\right] = \frac{\mathbb{E}[c(\varepsilon_t^d)]}{1-\beta}$, constant.

This can be microfounded from the textbook three-equation New Keynesian model, for a specific monetary policy rule, and adding a downward real wage rigidity, see Online Appendix B.

This model exhibits a Phillips curve with a kink at $-\frac{\gamma}{\kappa}$. When $y_t < -\frac{\gamma}{\kappa}$ changes in output gap do not affect inflation. When output gap is sufficiently low, the downward wage rigidity binds. Thus, lower output does not lead to lower marginal costs: real wages stay fixed, and the adjustment occurs through quantities.

Assume that the researcher has access to a time series ε_t^d, and they are interested in possible non-linearities in the response of inflation to demand shocks. However, the researcher does not know the true DGP. Thus, she decides to estimate the following local projections

$$\pi_t = \beta_0 + \beta_1 \varepsilon_t^d + u_{0,t} \tag{5}$$

$$\pi_t = \alpha_0 + \alpha_1 \varepsilon_t^d + \alpha_2 (\varepsilon_t^d)^2 + u_{1,t} \tag{6}$$

$$\pi_t = \theta_0 + \theta_1 \varepsilon_t^d + \theta_2 (\varepsilon_t^d)^3 + u_{2,t} \tag{7}$$

We only consider the contemporaneous effect of shocks on inflation because this economy is static. We characterize the estimands for each of the proposed regressions. We focus on the case of $\gamma = 0$ below. The case of $\gamma > 0$ is explained in detail in Online Appendix B.3.

Only Sign Non-linearities

Assume $\gamma = 0$ so that there is a kink in the Phillips curve at $y_t = 0$. In this economy, only a positive output gap has an impact on inflation. Since the kink is exactly at 0, the function $f(y_t)$ features a pure sign non-linearity: all that matters for the response is whether the shock is negative or positive, independently of the size of the shock. In particular, the marginal effect of a demand shock in inflation, $g_\varepsilon(a)$, is given by:

$$g_\varepsilon(a) = \begin{cases} \kappa & \text{if } a > 0 \\ 0 & \text{if } a < 0 \end{cases} \tag{8}$$

Note that this function is not defined at 0 (since $c(x)$ is not differentiable at 0), but this does not pose a problem to computing the weighted averages since no value of a has a positive mass in terms of weights.

Only linear terms. Assume that we fit the linear projection (5). What is the value of β_1? Using Lemma 1

$$\beta_1 = \int_{-b}^{b} \omega_1(a) \mathbb{E}\left[g_\varepsilon\left(a, \bullet\right)\right] da = \frac{\kappa}{2}.$$

Figure 1 shows each component of the weighted average representation of the estimand. The first panel depicts the marginal effect of a demand shock on inflation, which is not constant: positive shocks have a constant positive marginal effect, whereas negative shocks have no effect at all. The second panel shows the implicit weighting scheme when the distribution of the shocks is uniform. Finally, the last panel shows the product of the marginal effect and the weight. Note that our estimate is the integral of the curve in the third panel.

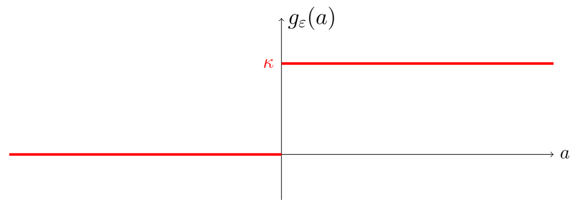

(a) Marginal effect of a demand shock on inflation, $g_\varepsilon(a)$

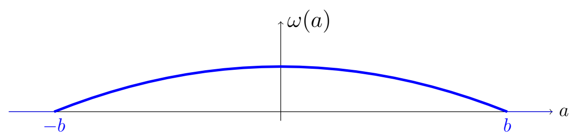

(b) Weighting scheme implied by a uniform distribution over $[-b, b]$

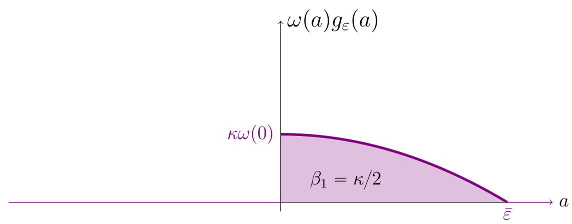

(c) Product of the marginal effect and the weight for a given value of a, i.e $\omega_1(a)g_\varepsilon(a)$

Figure 1

The first panel shows the marginal effect of a demand shock in inflation. The second panel shows the implicit weighting scheme in a uniform distribution. The last panel shows the product of the two. Therefore, the estimand $\beta_0 = \int_{-b}^{b} \omega(a)g_\varepsilon(a)da$ is given by the area under the curve.

Adding an even transformation. What happens if we include some even transformation of the shock as a regressor, as in (6)? The left panel in Figure 2 shows the different components of the estimand α_2. The first panel depicts the "demeaned" marginal effect of the shock. The marginal effect is below average for negative demand shocks, and above average for positive shocks. The second panel depicts the implicit weighting scheme when we use x^2 as a regressor. In this case, weights are negative for negative values of a and positive

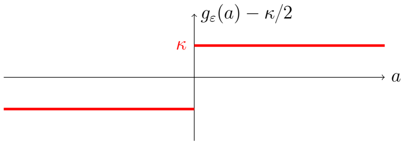

(a) Marginal effect of a demand shock on inflation minus its average

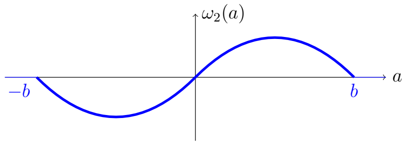

(b) Weighting scheme for $f(x) = x^2$

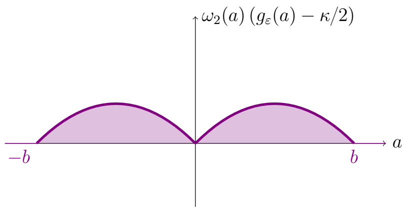

(c) The curve shows the product of lines in the above two panels. The estimand is the area below the curve.

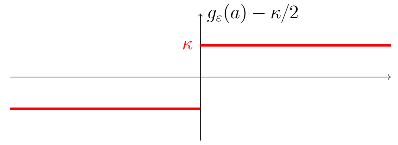

(d) Marginal effect of a demand shock on inflation minus its average

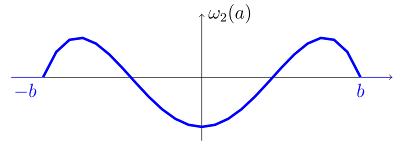

(e) Weighting scheme for $f(x) = x^3$

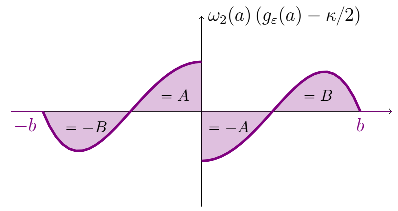

(f) The curve shows the product of lines in the above two panels. Different areas cancel out when integrating over $[-b, b]$.

Figure 2

The true DGP features sign non-linearities only. Left figure: even function as a regressor. Right figure: odd function as a regressor.

for positive values.[14] This is the reason why we get a non-zero estimate.[15] The third panel depicts the product of the two. For the negative region, we have negative weights multiplied by below-average marginal effects, which are also negative. Consequently, their product is positive, captured by the first hump. For positive values, both curves are positive, so the product of the two components is also positive. This generates the second hump. Finally, the estimand is the area below the curve, which is positive. We argue in Section 4 that, under some conditions, something similar will happen regardless of the specific even regressor used.

Adding an odd transformation. We now consider what happens if the researcher runs (7). In this case, the estimand is zero. The right panel in Figure 2 shows what happens in terms of weights when using an odd transformation but there are only pure sign non-linearities. The center, right panel shows the weighting scheme when $f(x) = x^3$ is used. In this case, the weights are symmetric around zero. The bottom right panel shows the product

[14] Furthermore, the weights are an odd function. In Section 4 we argue that this is the key property.

[15] Online Appendix B computes the estimands corresponding to the non-linear regressors.

of the lines in the first two panels. The area that corresponds to small, negative shocks (A) perfectly cancels out with the area that corresponds to small, positive shocks ($-A$). The same happens with large, negative and large, positive shocks. This symmetry property of the weights is crucial to obtain the described cancellations. A similar argument would apply for any symmetric weights $w_2(a)$ that integrate to zero. We argue in Section 4 that this symmetry will arise whenever we use an odd function as the non-linear transformation.

Takeaways. There are three main takeaways from this exercise. First, the shape of the weights has direct implications on whether the estimand of interest is zero or not. Symmetry of the weights is the key property of the generated $\theta_2 = 0$. On the other hand, when the weights are odd then positive and negative values of the shock get weights of different signs. Finally, the type of function used (even or odd) had direct implications on the symmetry of the weights. In the example, even transformations (such as x^2) generated odd weights, and odd transformations (such as x^3) generated even weights. In the following section, we provide conditions under which this insight generalizes to a general DGP.

4 Identification Results

4.1 Main separation result

This subsection presents the main separation result. First, we formally define sign and size non-linearities and present some examples. Building on these definitions, we present the result, and discuss the necessary assumptions and connection to the literature.

Formal definition of sign and size non-linearities. First, we need to formally define non-linearities by sign and size for any arbitrary function $m(x)$.

Definition 1 (Pure size dependence)**.** *The (non-constant) function $m(x)$ has pure size dependence but not sign dependence if it is an even function, i.e. $m(a) = m(-a)$ for any a.*

Definition 2 (Sign dependence)**.** *The (non-constant) function $m(x)$ has sign but not pure size dependence if it is an odd function, i.e $m(a) = -m(-a)$ for any a.*

If m is even, then the sign of the running variable a does not matter. Only the magnitude a matters in order to determine $m(a)$. Therefore, in terms of intuition, even functions do not show sign dependence. For odd functions, in principle there is both dependence of sign and size: after all, larger a can be different than smaller a. However, the odd nature of the function makes it such that if we average the value of the function at a and $-a$, we get zero, i.e $(m(a) + m(-a))/2 = 0$ by definition. In that sense, *on average*, size does not matter. Positive and negative values cancel out if we average (with symmetric weights) over any symmetric interval around zero. Therefore, the even-odd decomposition as the natural definition of size and sign non-linearities for a general function $m(a)$.[16]

Since any function has a unique decomposition as a sum of an even and an odd function, we can decompose the average marginal effect $\mathbb{E}\left[g_{\varepsilon}\left(a, \bullet\right)\right]$ into its even and odd components:

Definition 3. *Define $h(a) = \mathbb{E}[g_{\varepsilon}(a, \bullet)] - \beta$ and decompose it in its even and odd components as:*

$$h(a) = h^{odd}(a) + h^{even}(a)$$

where $h^{odd}(a)$ is the odd part and $h^{even}(a)$ is the even part; and β is the OLS estimand of regressing y_{t+h} on a constant and ε_t. We denote the odd part as the one related to sign-only non-linearities, and the even part as the one related to size-only non-linearities.

This way, we have a formal sense in which $\mathbb{E}\left[g_{\varepsilon}(a, \bullet)\right]$ has sign and/or size non-linearities. If $\mathbb{E}\left[g_{\varepsilon}(a, \bullet)\right]$ is constant (i.e g is linear in that shock) and $h(a) = 0$ for all a. Thus, if the underlying DGP is linear there are neither sign nor size non-linearities. We subtract β so that $h(a)$ captures only non-linearities, otherwise $h(a)$ would be constant under a linear DGP.

Examples for specific functional forms. We now present some examples to clarify how the definition works. These examples should not be seen as coming from any structural economic model, although some of them closely resemble the solution of DSGE models for different orders of approximation.[17]

Example 2. *Assume that $\mathbf{y}_t$ is one dimensional and that there is only one structural shock, ε_t. Assume that $g(\{\varepsilon_t\}) = \sum_{j=0}^{\infty} a_j^+ \max(\varepsilon_{t-j}, 0) + a_j^- \min(\varepsilon_{t-j}, 0)$.*

[16] The use of language here is not casual: we name these two properties "pure size dependence" and "sign dependence". We make this distinction because of the following: the value of an odd function can depend on a, so in that sense it is also "size dependent". However, as it was previously explained when averaging across any symmetric interval around zero, the values for positive and negative x cancel out.

[17] See Lan and Meyer-Gohde (2013) for a thorough discussion of solution methods for DSGE in the sequence space for higher-order approximations.

In this case, $h(x) = 1_{\{x>0\}}a_h^+ + (1 - 1_{\{x>0\}})a_h^-$ *(it is not defined when* $x = 0$*). After demeaning, our definition labels this case as sign asymmetry with no pure size asymmetry.*

The first case shows that our definition matches standard intuition. When the literature considered different slopes for positive and negative shocks (as Barnichon et al. (2022a)), the intuition is that all that matters is the sign of the shock. The next example shows that some extra care has to be put in place to think about size and sign asymmetries in the quadratic case.

Example 3. *Assume that* $\mathbf{y}_t$ *is one dimensional and that there is only one structural shock,* ε_t*. Assume that* $g(\{\varepsilon_t\}) = \sum_{j=0}^{\infty} a_j\varepsilon_{t-j} + \frac{1}{2}\sum_{j=0}^{\infty} b_j\varepsilon_{t-j}^2$.

In this case, $h(x) = a_h + b_h x$*. If we run a regression with a linear and a quadratic term, the linear term will capture* a_h*, and the quadratic term will capture* b_h*. Note that our definition labels this case as sign asymmetry with no pure size asymmetry.*

The last sentence may sound puzzling to the reader, especially since some empirical literature tries to tell apart a quadratic specification from a kinked specification.[18] *Our definition labels both as sign asymmetries with no pure size asymmetries.* Assume for the sake of the argument that $b_j > 0$. Then, under a quadratic specification, negative shocks have on average lower marginal effect than positive shocks. For example, assume that the shocks are standard normal.[19] Then $\mathbb{E}\left[a_h + b_h\varepsilon_t \middle| \varepsilon_t > 0\right] = a_h\frac{2b_h}{\sqrt{2\pi}}$ and $\mathbb{E}\left[a_h + b_h\varepsilon_t \middle| \varepsilon_t < 0\right] = a_h - \frac{2b_h}{\sqrt{2\pi}}$, so on average positive shocks have larger effects than negative shocks: there is a sign asymmetry. Our definition also says that in this case there is no "pure" size asymmetry. However, it is clear that the size of the shock also matters: small shocks in absolute value have different marginal effects than larger shocks. But under this specification, negative, large shocks have very low marginal effect; whereas positive, large magnitude shocks have a very high marginal effect. Once averaged over a range of values both cancel, so it is not size that matters, but sign.

The next example shows a pure size non-linearity.

Example 4. *Assume that* $\mathbf{y}_t$ *is one-dimensional and that there is only one structural shock,* ε_t*. Assume that* $g(\{\varepsilon_t\}) = \sum_{j=0}^{\infty} a_j\varepsilon_{t-j} + \frac{1}{6}\sum_{j=0}^{\infty} c_j\varepsilon_{t-j}^3$.

[18] For example, Forni et al. (2023) and Ben Zeev et al. (2023)

[19] Of course, this distribution does not have bounded support, but our derivations can be adapted to account for that.

In this case, $h(x) = a_h + \frac{c_h}{2}x^2$. If we run a regression with a linear and a cubic term, the linear term will capture a_h, and the cubic term will capture c_h. Note that our definition labels this case as pure size asymmetry with no sign asymmetry.

In this case, all that matters is whether the shock is "large" or "small" in absolute value. The quadratic term in h implies that the sign of the shock does not matter at all. If y_t is output and ε_t are government spending shocks, the response is completely symmetric to a 5 percent increase or decrease. However, if $c_j > 0$ the response is proportionally larger for large magnitude shocks: the impact of a 10 percent shock is larger than 2 times the impact of a 5 percent shock.

Separation result. We would like to use the weighting schemes defined above in order to keep only one of the components at the time, thus addressing the amount of each type of non-linearity present in the data. The following lemmas show that, if we had weights that are either even or odd, we can average out the odd and even parts (respectively) of $h(a)$.

Lemma 2 (Odd (even) weights wipe out pure size (sign) non-linearities)**.** *The following hold:*

(i) Assume that $m(a)$ is an even function (i.e $m(a) = m(-a)$). Assume that $\omega(a)$ is an odd function. Let $\bar{x} \in \mathbb{R}_+$. Then, $\int_{-\bar{x}}^{\bar{x}} \omega(a)m(a)da = 0$

(ii) Assume that $m(a)$ is an odd function (i.e $m(a) = -m(-a)$). Assume that $\omega(a)$ is an even function. Let $\bar{x} \in \mathbb{R}_+$. Then, $\int_{-\bar{x}}^{\bar{x}} \omega(a)m(a)da = 0$

The Lemma follows from the orthogonality between odd and even functions when the inner product is the integral of the product over a symmetric interval. Careful inspection of the formula for β_2 in Lemma 1 shows that, if one can obtain either even or odd weights, then each of those weighting schemes would wipe out only one type of non-linearity. The following lemma shows that, if the distribution of the shock is symmetric then using odd (even) functions generates even (odd) weights.

Assumption 6. *The unconditional distribution of ε_t is symmetric around zero. Formally, if we let $F(x)$ be the unconditional CDF of the shock ε_t, the distribution is symmetric if $G(x) = F(x) - 1/2$ is an odd function.*

Lemma 3 (Even (odd) functions yield odd (even) weights)**.** *Under Assumptions 1-6* *(i) If $f(x) = f(-x)$, i.e f is an even function, then $\omega(a) = -\omega(-a)$, i.e the weights are an odd function of a.*

(ii) If $f(x) = -f(-x)$, i.e f is an odd function, then $w(a) = w(-a)$, i.e the weights are an even function of a.

The proof is shown in Appendix A.2. The next proposition summarizes the practical implications of the above lemmas.

Proposition 1 (Separation of sign and size non-linearities)**.** *Consider the coefficient θ_h of the following linear projection:*

$$y_{t+h} = \alpha_h + \beta_h \varepsilon_t + \theta_h f(\varepsilon_t) + u_{t+h}$$

Let $h(a) = \mathbb{E}[g(a, \bullet)] - \beta$ and consider the even-odd decomposition of $h(a) = h^{odd}(a) + h^{even}(a)$. Under Assumptions 1-6, if $f(\cdot)$ is an even function then θ_h does not depend on $h^{even}(a)$, i.e it only captures sign non-linearities. Analogously, if $f(\cdot)$ is an odd function, θ_h does not depend on $h^{odd}(a)$, i.e. it only captures size non-linearities.

To see this at work, assume that we are interested in sign non-linearities. That is, $\mathbb{E}[g_\varepsilon(a, \bullet)]$ is larger (on average) for positive shocks than for negative shocks. Thus, if there were sign non-linearities, we would expect $(\mathbb{E}[g(a, \bullet)] - \beta)$ to be negative for negative shocks, and $(\mathbb{E}[g(a, \bullet)] - \beta)$ to be positive for positive shocks. Therefore, we should use weights $\omega_2(a)$ such that they have one sign when the shock is positive, and the opposite sign when the shock is negative to detect sign non-linearities: negative $(\mathbb{E}[g_\varepsilon(a, \bullet)] - \beta)$ would have negative weights, and positive $(\mathbb{E}[g_\varepsilon(a, \bullet)] - \beta)$ would have positive weight, so the estimand would be positive. However, we might be worried about confusing sign with size non-linearities. In that case, we want the weight of the shock to be identical for shocks with the same absolute size, but with different sign. That is, we would like $\omega_2(a) = -\omega_2(-a)$. That way, even if there are size non-linearities, large positive shocks would cancel with large negative shocks of the same magnitude, given zero impact of the size non-linearity in β_2.

Analogously, we might be interested in size non-linearities: larger shocks in absolute value having different impact that smaller shocks. In that case, we should use weights $\omega_2(a)$ that have one sign when the shock is small in absolute value, and the opposite sign when the shock is large. Moreover, in order not to confuse size with sign non-linearities, we the weights should $\omega_2(a) = \omega_2(-a)$ so that positive and negative shocks have the same weights. This would wipe out any sign non-linearity. If $\omega_2(a) = \omega_2(-a)$ and there is sign asymmetry, then the shock with positive a will get higher than average $\mathbb{E}[g_\varepsilon(a,\bullet)]$, and the shock with negative a will get lower than average $\mathbb{E}[g_\varepsilon(a,\bullet)]$. As both have the same weight, they will cancel out and the estimand will not be contaminated by any sign non-linearities.

Comparison to the literature. Some empirical work estimates specifications including a linear term and either $f(x) = x^2$ or $f(x) = |x|$ (or transformations of these) to disentangle size and sign non-linearities (see Ben Zeev (2020), Forni et al. (2023) and Ben Zeev et al. (2023)). The idea is that x^2 captures size effects, whereas $|x|$ captures sign only. However, the above lemmas imply that this approach is only captures asymmetries by sign since both $f(x) = x^2$ and $f(x) = |x|$ are even functions. On the other hand, in order to capture pure size asymmetries, one must include odd non-linear transformations. This is in line with the procedure followed, for example, in Ascari and Haber (2021).

4.2 Non-symmetric distributions

At this point, a natural question to ask is: what happens if the distribution is not symmetric? It is always possible to decompose weights into an even and odd part. Using the weighted average representation of the estimand:

$$\beta_2 = \int \omega_2(a)h(a)da = \int \omega_2^{odd}(a)h^{odd}(a)da + \int \omega_2^{even}(a)h^{even}(a)da.$$

The result follows because odd and even functions are orthogonal, all cross terms cancel. Therefore, the estimand can be decomposed into a sign and size component. The even-odd decomposition of the weights can be estimated from the observed distribution of the shock ε_t (Plagborg-Møller, 2023). Although the exact result separation relies on symmetry, small deviations from this assumption only generate small amounts of contamination. In our empirical application, we show that, as the distribution of the shocks is close to symmetric, even (odd) weights become close to zero when we use an odd (even) transformation of the shock.

On the other hand, our analysis also shows what happens if the underlying distribution is far from symmetric. In that case, regressors such as $|x|$, which are thought to capture sign non-linearities only, could actually mix both kinds of non-linearities. Assume, for example, that the distribution of the shock is right-skewed. In that case, positive values are also larger on average. If we fit a local projection with $|x|$ as a regressor, we are implicitly comparing the response of the economy to positive and negative shocks. If $\beta_2 \neq 0$, that intuitively means that the causal effect of positive shocks is not the same as the one of negative shocks. However, as positive shocks are also larger on average, a priori we do not know whether the differential response is due to the shocks being positive or large.

Comparison to the literature. These results stress that the common practice of adding $|x|$ as a regressor and interpreting the results as evidence of sign non-linearities can be very misleading if the underlying distribution is not symmetric. This is of particular relevance for the case of government spending shocks (Barnichon et al. (2022a), Ben Zeev et al. (2023)). The distribution the government spending news shocks (Ramey and Zubairy (2018)) used by these papers is far from symmetric, casting doubt on the interpretation of the results in Barnichon et al. (2022a) as coming from sign non-linearities.

4.3 A pragmatic defense of piecewise linear functions

The separation result derived above shows how, by picking non-linear transformations of the appropriate kind (even or odd), it is possible to test for the presence of size and sign non-linearities separately. However, it is unclear what specific functional form to use. Furthermore, it is also unclear if we can give any interpretation to the estimands in terms of magnitudes. In this subsection, we recommend the use of piecewise linear functions because it facilitates reporting. Simulations in Section 5 show that they also have better finite sample properties.[20]

Recall the estimands are weighted averages of marginal effects (Angrist et al. (2000)). Then, the intuitive interpretation in the context of sign and/or size non-linearities must be in terms of "average (marginal) responses". For example, if we are interested in sign, we should aim at estimands that could be interpreted as some weighted average of (marginal) effect of positive or negative shocks. In what follows, we argue that this observation makes piecewise linear functions particularly attractive among the class of non-linear functions. In a nutshell, other kinds of functions (for example, a quadratic) require the researcher to choose a given size for the shock in order to report. As we show in the following example, this can be potentially misleading.

Example 5. *Assume the distribution of the shock ε_t is standard normal, and the researcher*

[20] Part of this argument is conceptually related to the discussion in Gelman and Imbens (2019).

fits three different local projections:

$$y_{j,t+h} = a_0 + a_1\varepsilon_t + a_2|\varepsilon_t| + u_{j,t+h} \tag{9}$$

$$y_{j,t+h} = b_0 + b_1\varepsilon_t + b_2\varepsilon_t^2 + v_{j,t+h} \tag{10}$$

$$y_{j,t+h} = c_0 + c_1\varepsilon_t + r_{j,t+h} \tag{11}$$

The researcher is interested in the effect of positive shocks. According to the functional form used, the researcher computes the effect of a positive shock from a linear specification as c_1, the one from a quadratic specification as $b_1 + b_2 x$ for $x \in \{1, 2, 3\}$ and the one from a kinked specification as $a_1 + a_2$.

The way in which results are reported in the example mirrors the standard practice in the literature (see Barnichon et al. (2022a), Ben Zeev et al. (2023), Caramp and Feilich (2022) for some recent examples). Usually, just one standard deviation is used. Here we use a wider range of values for better visualization.

As the estimands are linear in the weights, we can always compute the implicit weights in any linear combination of coefficients by simply taking the linear combination of the weights —this follows directly from the weighted average representation of the estimands and linearity of the integral. Thus, whenever we do a transformation for reporting as proposed above, we can compute the implicit weights on that transformation in order to understand what are the implicit weights on the transformation chosen. This is of particular importance whenever the researcher chooses to use a functional form in which the impulse response, scaled by the size of the shock, depends on the value of x, as in the case of x^2.

Turning the actual weights in Example 5, those are depicted in Figure 3. The main takeaway is that using the assumed functional form to interpret the results is potentially *very* misleading. When reporting treatment effects under the assumption of a particular functional form assumption, we are actually computing weighted averages of marginal effects across the whole support of the distribution. These may hold little resemblance with the actual effect of interest if the functional form is misspecified, as it likely is. Take, for example, the weights implicit in a shock of size 2 (dashed blue line). As we can see in Figure 3, this value is actually determined mainly by shocks of size $x \in [-0.5, 2]$. Furthermore, there are negative weights given to the marginal effects of negative shocks to the left of -0.5. It is clear how problematic can this be to interpret: what, in theory, should be a measure of $\mathbb{E}\left[y_{t+h}|\,\varepsilon_t = 2\right] - \mathbb{E}\left[y_{t+h}|\,\varepsilon_t = 0\right]$ is actually a weighted average that places the most weight on values around 0 and 1.5, and that also uses negative weights. If the true functional form

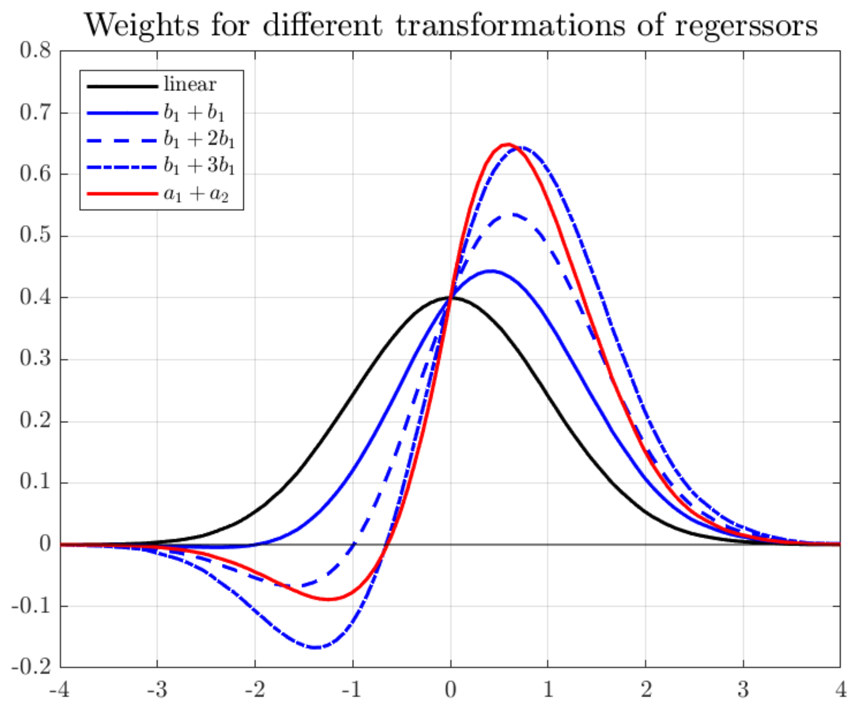

Figure 3: Implicit weights on different combinations of regressors.
The black line shows the marginal effects computed assuming linearity, i.e. the weights implicit in c_1. The blue lines denote the total effects with the assumption of a quadratic function, i.e. the weights in $b_1 + b_2x$ for $x \in \{1, 2, 3\}$. Finally, the red line shows the weights on $a_1 + a_2$, the implicit marginal effect under the assumption using a linear term, and an absolute value as regressors.

was indeed quadratic, then this weighted average would yield $b_1 + b_2x$. However, if the functional form is misspecified, the alleged effect of a shock of size 2 is actually a weighted average, with non-trivial weight below 0.[21]

On the other hand, piecewise linear functions have weights that align more with the desired interpretation. The red line in Figure 3 shows that most of the weight is put in positive values, with different weights on different sizes. There are also mildly negative weights. But conceptually, it is correct to interpret $a_1 + a_2$ as a measure of the effect of a positive shock: this estimand is a weighted average of marginal effects, with almost all weight on positive values.

In a nutshell, since all our estimands are averages of marginal effects, we should use a non-linear functional form that has an easy interpretation in terms of averages. Piecewise

[21]Online Appendix C contains an extended discussion on the implicit weighting schemes of different non-linear transformations for the case of a Gaussian regressor.

linear functions serve this purpose. A secondary advantage, to be discussed in Section 5, is that piecewise linear functions show better finite sample performance in simulations.

4.4 Practical takeaways

There are several practical takeaways from the previous discussion.

(i) For testing, estimate a local projection augmented with a non-linear transformation of the shock as in (4). Then, one can test for sign or size non-linearities at horizon h by testing $H_0 : \beta_{2,h} = 0$.[22] We stress that being unable to reject $\beta_{2,h} = 0$ for both size and sign regressors does not imply that non-linearities are not present at all. For example, the data could feature state-dependence.

(ii) If the researcher is interested in sign non-linearities, they should choose a non-linear even function of the shock. Two intuitive choices are $|x|$ and x^2, although $|x|$ is preferred for the reasons explained in Section 4.3. $(\beta_{1,h} + \beta_{2,h})$ is the estimated effect of positive shocks, and $(\beta_{1,h} - \beta_{2,h})$ is the estimated affect of negative shocks.[23]

(iii) If the researcher is interested in size non-linearities, they should choose a non-linear odd function of the shock. In the empirical application in Section 6 we use a piecewise linear function.[24] This is preferred to $x|x|$ and x^3 by the arguments presented in Section 4.3.

(iv) Following Plagborg-Møller (2023), we recommend to compute and plot the weights. Including a non-linear odd transformation may lead to negative weights on the coefficient of the linear term. Thus, one should be careful when interpreting $\beta_{1,h}$ when another odd regressor is added. It is advisable to benchmark results from this regression against the case in which only a linear regressor is used.

(v) Shocks may have an asymmetric distribution. The assumption of symmetry can and should be formally tested. In our empirical implementation, we use the test devel-

[22] Since we propose to fit regressions that are linear in parameters, standard guidance for inference in local projections applies (Montiel Olea and Plagborg-Møller (2021), Montiel Olea and Plagborg-Møller (2019)). See D for details.

[23] Note furthermore that, in terms of fit and testing, using x and $|x|$ as regressors and test $\beta_2 = 0$ is equivalent to using $\min(x, 0)$ and $\max(x, 0)$ and test for differences in the coefficients.

[24] In particular, we use $f(x) = 1_{\{x \leqslant -\bar{b}\}}(x + \bar{b}) + 1_{\{x \geqslant \bar{b}\}}(x - \bar{b})$. This function is similar to the one depicted in the bottom panel of Figure C.2 in Online Appendix C. In our empirical application, we choose $\bar{b}$ to be the 80-th quantile of the distribution of the shock.

oped by Randles et al. (1980). We also recommend computing and plotting even-odd decomposition of the weights, in order to assess the degree of contamination in the estimand in the sample.

In Online Appendix D we develop extensions on further topics of practical relevance, such as imperfect measurement of the shock, inclusion of controls, different estimation methods, and inference.

5 Simulation Study

In this section, we present simulation evidence from a stylized DGP to study the finite sample properties of the procedure outlined above. We use the following DGP, which is a variant of the one used by Gonçalves et al. (2021).

$$\begin{pmatrix} y_{1,t} \\ y_{2,t} \end{pmatrix} = \begin{pmatrix} \rho_y & 0 \\ 0.5 & \rho_y \end{pmatrix} \begin{pmatrix} y_{1,t-1} \\ y_{2,t-1} \end{pmatrix} + \begin{pmatrix} 1 & 0 \\ \rho_\varepsilon & \sqrt{1-\rho_\varepsilon^2} \end{pmatrix} \begin{pmatrix} f(\varepsilon_{1,t}) \\ \varepsilon_{2,t} \end{pmatrix} \tag{12}$$

The nonlinearity in the DGP is introduced via the $f(x)$ term, so the contemporary response of the shock is non-linear. Conditional on the initial impulse, the dynamic behavior of the system is linear. This allows us to cleanly control what kinds of non-linearities are present in the true DGP. The structural shocks $\varepsilon_{1,t}, \varepsilon_{2,t}$ have iid standard normal distribution. We consider several choices of $f(x)$ below. In all cases, we scale $f(x)$ so that guarantees that the process $f(\varepsilon_{1,t})$ has unit variance. Under this parametrization, $corr(f(\varepsilon_{1,t}), \varepsilon_{2,t}) = \rho_\varepsilon$. We calibrate $\rho_\varepsilon = 0.3$ and $\rho_y = 0.95$. For $y_{2,t}$, this calibration implies that the impulse response of $y_{2,t}$ to $\varepsilon_{1,t}$ peaks at horizon 20, which is consistent with what is observed in monthly data.

We consider different choices for f in the true DGP. First, we consider $f(x) = x$ as a benchmark. For sign non-linearities, we specify $f(x) = c(x + |x|^\alpha)$ where $\alpha \in \{1, 2, 3\}$ is the degree of the non-linear function. For size non-linearities, we use $f(x) = c(x + k(x))$. For degree 1, $k(x)$ is piecewise linear, $k(x) = 1_{\{x \leqslant -\bar{b}\}}(x + \bar{b}) + 1_{\{x \geqslant \bar{b}\}}(x - \bar{b})$ where $\bar{b}$ is set such that 40% of the shocks fall outside of $[-\bar{b}, \bar{b}]$. For degrees 2 and 3, $k(x) = x|x|^{\alpha-1}$ where α is the degree. In all cases, c is chosen such that $Var(f(\varepsilon_t)) = 1$.

For each DGP, we run specifications of the form:

$$y_{j,t+h} = \alpha_{j,h} + \beta_{j,1,h}\varepsilon_{1,t} + \beta_{j,2,h}m(\varepsilon_{1,t}) + \text{controls} + u_{j,t+h}. \tag{13}$$

All specifications include a constant and one lag of $[y_{1,t}, y_{2,t}]$ as a controls. First, as a benchmark, we only include a linear term. For to specifications with even transformations, we use $m(x) = |x|^\alpha$ with $\alpha = 1, 3$.[25] We refer to α as the order. For specifications with odd transformations, we use $m(x) = k(x)$ with the same functions described in the above section.[26] The sample size is $T = 400$, and we take 5000 Monte-Carlo draws for each DGP.

We focus only on the response to variable $y_{2,t}$ because it exhibits richer dynamics. For each regression specification, we report the fraction of samples in which the null $H_0 : \beta_{2,h} = 0$ is rejected, for horizons $h = 0, 6, 12, 24, 48$. We use a t-test with heteroskedasticity-consistent standard errors as suggested by Montiel Olea and Plagborg-Møller (2021).[27] For the specification that only includes a linear term, we present the rejection frequencies of the null $\beta_{1,h} = 0$ instead.

Results. Table 5.1 displays the results. The first column shows the results when the true DGP is linear. The first set of rows (labeled "Linear") shows the proportion of rejections for the null hypothesis of $\beta_{1,h} = 0$. Under this parametrization, local projections show good power to detect the effect of a $\varepsilon_{1,t}$ shock. The second set of rows shows what happens when we add a non-linear even transformation, which should detect only sign non-linearities. Overall, the size of the test remains close to the theoretical 10%, with some over-rejection when we use higher-order transformations. A similar picture emerges when we consider the last set of rows, corresponding to odd transformations (labeled "Size").

Columns 2, 3, and 4 present the results when the true DGP has sign non-linearities, for different even functions. Focusing on the second set of rows, the power to detect the underlying sign non-linearity is close to invariant to the type of non-linear term used in the regressions: for a given column, the numbers reported in rows "Sign -1" and "Sign - 3" are almost equal. In order words, as long as we include an even transformation, the power to detect the non-linearity does not depend on the correct specification of the model. The only (partial) exception to this occurs when the true DGP is cubic, in which case the correct specification shows better power for horizon 24. Turning to the last set of rows, there are non-trivial size distortions when we use regressors of higher order. For example, when the true DGP features x^2 as the non-linear term but we use x^3 as a regressor, at horizon 0 we falsely detect a size non-linearity 28.1 percent of the samples, far more than the theoretical 10

[25] We do not report $\alpha = 2$, results in this case lie in between the ones for $\alpha = 1$ and $\alpha = 3$.

[26] In particular, for $\alpha = 1$, we use a piecewise linear function that is correctly specified, i.e it coincides with the one of the DGP.

[27] Under the assumed DGP, the regressions scores are uncorrelated, so heteroskedasticity-only SEs are consistent (Montiel Olea and Plagborg-Møller, 2021).

Reg.	Horizon	DGP						
		Linear	Sign			Size		
			1	2	3	1	2	3
Linear	0	1.000	1.000	1.000	0.694	1.000	1.000	1.000
	6	1.000	1.000	0.991	0.523	1.000	1.000	1.000
	12	1.000	0.999	0.881	0.352	1.000	1.000	0.999
	24	0.611	0.493	0.286	0.155	0.624	0.627	0.542
	48	0.105	0.118	0.103	0.106	0.111	0.102	0.101
Sign - 1	0	0.104	1.000	1.000	1.000	0.105	0.117	0.140
	6	0.106	0.998	1.000	1.000	0.105	0.110	0.114
	12	0.101	0.860	0.997	0.996	0.106	0.104	0.107
	24	0.107	0.261	0.441	0.467	0.107	0.104	0.104
	48	0.106	0.111	0.108	0.110	0.105	0.113	0.108
Sign - 3	0	0.153	0.999	1.000	1.000	0.147	0.229	0.267
	6	0.140	0.973	1.000	1.000	0.132	0.166	0.224
	12	0.132	0.717	0.999	1.000	0.135	0.145	0.182
	24	0.137	0.248	0.506	0.687	0.136	0.138	0.142
	48	0.138	0.148	0.148	0.147	0.141	0.150	0.135
Size - 1	0	0.112	0.122	0.145	0.168	0.902	1.000	1.000
	6	0.116	0.109	0.124	0.150	0.416	0.710	0.993
	12	0.115	0.111	0.109	0.128	0.224	0.366	0.763
	24	0.108	0.105	0.111	0.109	0.135	0.141	0.225
	48	0.103	0.104	0.103	0.111	0.105	0.108	0.110
Size - 3	0	0.157	0.202	0.281	0.319	0.806	1.000	1.000
	6	0.154	0.165	0.249	0.305	0.386	0.758	0.998
	12	0.162	0.153	0.204	0.271	0.238	0.429	0.886
	24	0.149	0.137	0.157	0.175	0.168	0.177	0.320
	48	0.145	0.137	0.155	0.139	0.145	0.151	0.141

Table 5.1: Rejection probabilities under different regression specifications (rows) and DGPs (columns). Regression specification is given by (13). The outcome variable is $y_{2,t+h}$. For the row "linear" the null hypothesis is $\beta_{1,h} = 0$ for the linear regressor. For all other rows, the null hypothesis is $\beta_{2,h} = 0$. All tests are conducted at 90% confidence using heteroskedasticity-consistent SEs. The sample size is $T = 400$, based on 5000 Monte-Carlo draws for each DGP. See text for details on the DGPs considered.

percent level. The size distortion is much smaller (but not zero) when considering regressors of lower order (piecewise linear). Columns 5,6 and 7 present the results for DGPs with size non-linearities. Essentially the same conclusions hold.

Takeaways. There are three main takeaways from the presented simulation exercise. First, if the true DGP is linear, the size of the test for non-linearities in finite samples is close to the theoretical size. Some small-size distortions (in particular, over-rejection) do appear when we use regressors of higher order. Second, if the true DGP features a given kind of non-linearity, the actual function used as a regressor matters little in terms of power. In other words, once we get the kind of non-linearity right, misspecification of the actual functional form does not meaningfully affect the power to reject the null of linearity. Third, if some kind of non-linearity is present (for example, sign) and we test for a different kind (for example, size), non-trivial size distortions appear when we use regressors of higher order. Overall, the results suggest that using piecewise linear functions (regardless of the kind of non-linearity) achieves both high power to reject the null when needed, as well as controlled size when the null is true.

6 Empirical Application: Oil supply news shocks

6.1 Data

For comparability reasons, we use the exact same data contained in the replication package of Känzig (2021), both for the shocks and for the outcomes of interest. The baseline specification uses monthly data from 1983M01 to 2017M12. Figure 4 shows the histogram of the series of oil supply news shocks, conditional on the shock being observed in the period. The distribution is fairly close to symmetric. In order to formally test this, we use the triples test developed by Randles et al. (1980). We fail to reject the null of symmetry. However, in our empirical estimates, we plot the implicit weights in the regression to assess how much the departure from symmetry matters in the sample.

6.2 Empirical estimates

We estimate the following specifications:

$$y_{j,t+h} = \alpha_{j,h} + \beta_{1,j,h}\varepsilon_t^{oil} + \beta_{2,j,h}f\left(\varepsilon_t^{oil}\right) + \text{controls} + u_{j,t+h}, \tag{14}$$

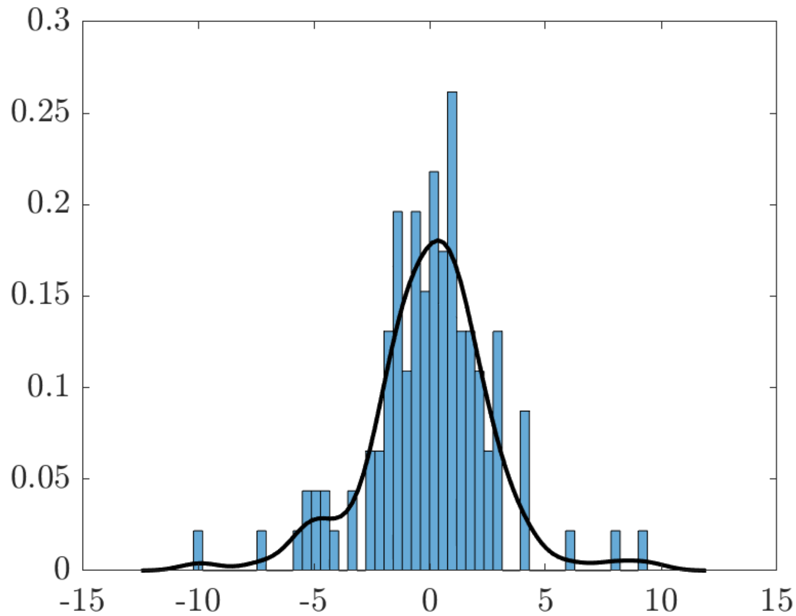

Figure 4: Distribution of oil supply news shocks, conditional on a shock occurring in the period. The x-axis is measured in percentage points. Bins correspond to the histogram, whereas the black line is the estimated density via kernel.

where $y_{j,t+h}$ is outcome variable j at horizon h, $\beta_{j,h,1}, \beta_{j,h,2}$ denote the coefficients on the linear and non-linear term respectively, and ε_t^{oil} is the oil supply news shock.[28] The baseline specification is the same as in Känzig (2021): i) we run the local projections in levels, with a constant ii) as controls, we use 18 lags of real oil price, oil production, oil stocks and world industrial production in all equations, and we add 18 lags of the outcome variable in each regression.[29,30,31] We use Newey and West (1987) standard errors with lag-length set

[28]Note that we are *not* using the shock as IV for the real oil price, but rather directly as a regressor in its own right. A similar approach is used in Tenreyro and Thwaites (2016), Ben Zeev (2020) and Ben Zeev et al. (2023). We expand on this point in Online Appendix D.

[29]The main analysis in Känzig (2021) is carried out using a SVAR-IV approach. However, several specifications using LP-IV are discussed in the appendix. We take the baseline specification used in Appendix A.2.2 of that paper.

[30]Even though the shock series is not predictable using lags, we include lags to mirror the benchmark specification in Känzig (2021). As it is known in the literature (e.g (Hamilton and Herrera, 2004)), the oil market features persistent cycles, which makes it important to control for sufficient lags.

[31]Our method assumes stationarity, which does not hold for all series considered in levels. For comparability with Känzig (2021), we nevertheless run the regressions in levels. In the presence of a unit root, the fact that we augment the specification with lags allows us to use standard inference on the coefficient on the regressor of interest (Montiel Olea and Plagborg-Møller (2021)). Alternatively, we also considered a specification where the data for CPI and industrial production is log-differences. The results regarding

to the horizon of the local projection.[32] Results are robust to using a smaller set of controls, different lag lengths, and accounting for multiple testing.[33]

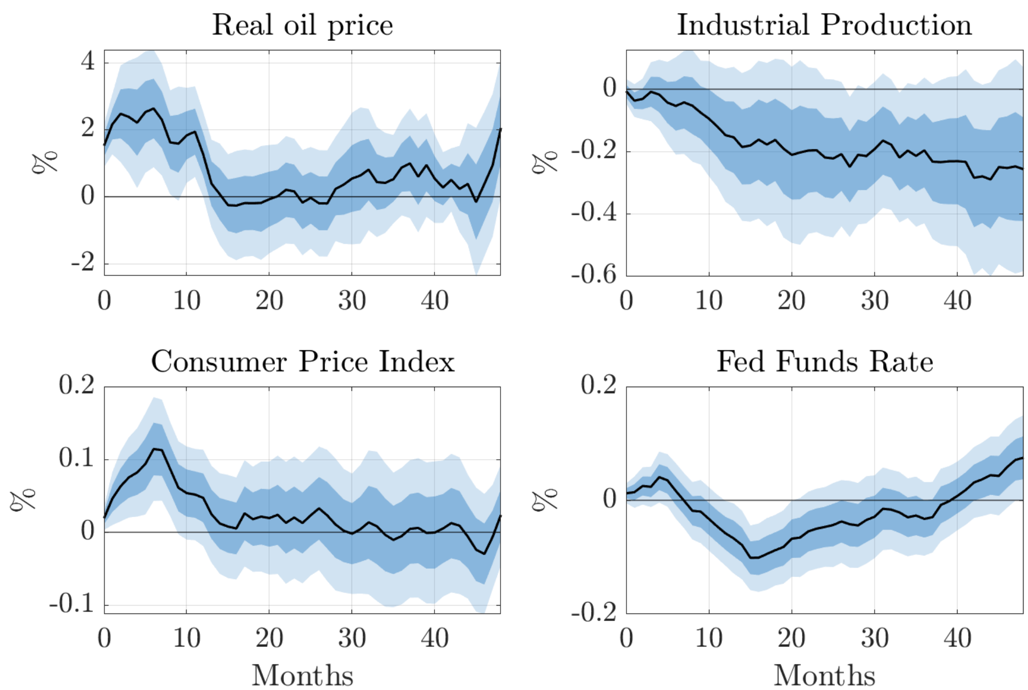

Figure 5: IRFs obtained via Local Projections

Each panel depicts the sequence $\hat{\beta}_{j,h}$ that arises from estimating:
$y_{j,t+h} = \alpha_{j,h} + \hat{\beta}_{j,h}\varepsilon_t^{oil} +$ controls $+ u_{j,t+h}$ for each different outcome variable j, and for horizons $h = 0, \ldots, 48$. Shaded areas indicate 68 and 95 percent confidence bands respectively.

6.2.1 Baseline Specification: Linear

Suppose we estimate (14) using only a linear term of the shock as a regressor. Figure 5 shows the estimated impulse responses for the variables of interest, which mirrors the results in Känzig (2021). There is a significant impact on real oil prices and CPI. Most of the impact is over before 12 months. There are also some real effects, but not significantly different from

testing non-linearities are exactly the same.

[32]Using heteroskedasticity-only standard errors yields almost identical estimates for standard errors. This suggests that, although the assumptions in Montiel Olea and Plagborg-Møller (2021) may not be satisfied under the assumed DGP, in practice the difference in estimated standard errors is small.

[33]In particular, we consider using only real oil price as control instead of the full set of oil market variables. We consider lag lengths of 12, 18, and 24 lags as done in Känzig (2021). We account for multiple testing using the sup-t test proposed in Montiel Olea and Plagborg-Møller (2019). See Online Appendix E for details.

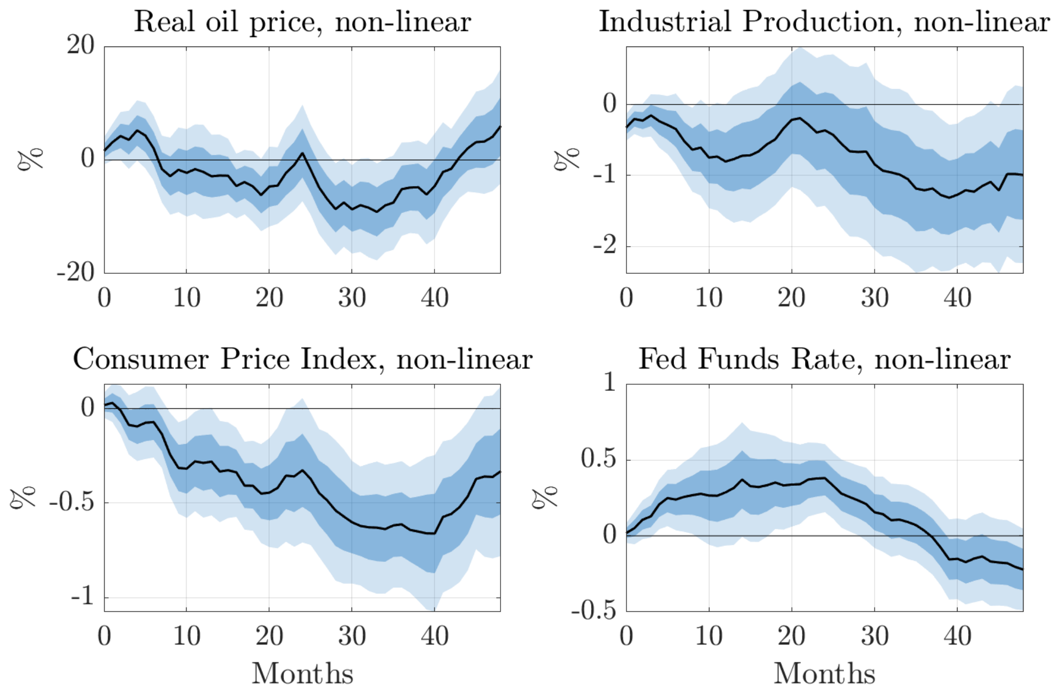

Figure 6: Estimates of the coefficients on the non-linear terms obtained via Local Projections
Each panel depicts the sequence $\hat{\beta}_{j,h,2}$ that arises from estimating:
$y_{j,t+h} = \alpha_{j,h} + \hat{\beta}_{j,h,1}\varepsilon_t^{oil} + \hat{\beta}_{j,h,2} f(\varepsilon_t^{oil}) +$ controls $+ u_{j,t+h}$ for each different outcome variable j, and for horizons $h = 0, \dots, 48$. Shaded areas indicate 68 and 95 percent confidence bands respectively.

zero at the 95 percent level. Finally, the interest rate path features a small rise followed by a reversal. The peak response is negative at around 10 basis points. Overall, oil supply news shocks have a non-trivial impact on prices and a moderate impact on industrial production and interest rates.

6.2.2 Size non-linearities

We now run (14) adding an non-linear regressor. In particular, we add a piecewise linear function $f(x) = 1_{\{x \leqslant -\bar{b}\}}(x + \bar{b}) + 1_{\{x \geqslant \bar{b}\}}(x - \bar{b})$ to capture non-linearities by size. We define "large" values of the shock in $[-\bar{b}, \bar{b}]$ as "small", and values outside as "large". The value of $\bar{b}$ is chosen as the maximum between the absolute value of the quantiles 0.2 and 0.8 of the distribution of the shocks conditional on no zeros.[34,35]

[34]This function is continuous, constant at zero when $x \in [-\bar{b}, \bar{b}]$, and is linear with a slope equal to 1 for all other values of x. The shape is the same as the piecewise linear function depicted in Figure C.2.

[35]We use the max between the absolute value of the quantiles because the regressor has to be and odd function. The 0.2 and 0.8 are arbitrary, although we choose percentiles such that there is a significant number

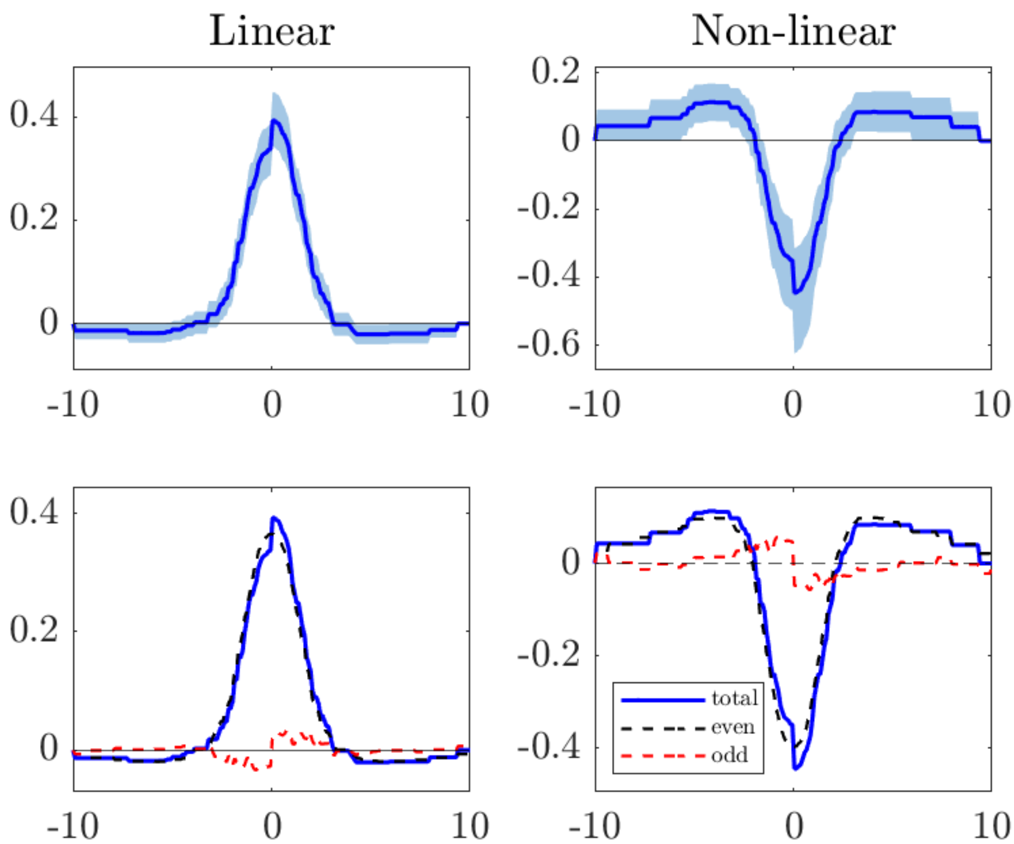

Figure 7: Implicit weights in the Local Projection $y_{j,t+h} = \alpha_{j,h} + \hat{\beta}_{j,h}\varepsilon_t^{oil} + \beta_{j,h,2} f(\varepsilon_t^{oil}) +$ controls $+ u_{j,t+h}$ when f is piecewise linear.
The shaded area indicates a 90 percent confidence band computed via bootstrap (See Online Appendix D).

Figure 6 shows the coefficients on the non-linear term ($\hat{\beta}_{2,j,h}$) for different horizons. We reject the null of linearity on impact for Industrial Production and at several horizons for Fed Funds Rate, and CPI. Figure 7 shows the estimated implicit weights on the linear and non-linear regressors on the estimands β_1 and β_2 in (14). The weights on the non-linear regressor are positive to "large" shocks, and negative weight to "small shocks". The linear regressor now features negative weights for large shocks, so the interpretation of it as a convexly-weighted average of marginal effects is less clear since some shocks have negative weights. The bottom panels in Figure 7 again show that the odd part of the weights is negligible. Thus, we are mainly capturing non-linearities by size and not by sign. This is

of observations in both groups. Results are robust to using other arbitrary bounds such as 1.5 and 2. Out of 115 non-zero observations, 45 are qualified as large shocks and 70 as small shocks. Online Appendix E.1 shows balance between the types of shocks.

due to the high degree of symmetry in the distribution of the shocks.

Since the coefficient on the linear term has negative weights, we use the IRF estimated only with a linear term (as in Section 6.2.1) as a benchmark. Figure 8a shows the estimated response using only a linear regressor and the response to a large shock.[36] Consistent with the non-linear term being significant, responses are quite different for large shocks. In particular, the effects on industrial production are proportionally larger. There is also a positive monetary policy response, whereas the linear case featured a mildly negative one. Finally, there is evidence of a reversal in prices. In particular, there is a negative response in prices starting a year after the shock. This mainly reflects that the IRF of real oil prices features a reversal, and it is not robust to the selection of controls and lags used.[37] Overall, there is significant evidence of size non-linearities: large oil shocks imply a proportionally larger real impact and difference in monetary policy response than small shocks.

In terms of models that can potentially account for this non-linearity, models with lumpy durable purchases are a possible candidate.[38] An important transmission channel for oil shocks are its impact on automobile purchases (Ramey and Vine, 2011). Since these purchases are lumpy, large shocks can trigger proportionally more adjustments than small shocks, implying size non-linearities.[39] Another possible explanation could be related to non-linearities in monetary policy response, with the Fed ignoring small shock but hiking rates aggressively for large shocks. Further study of the mechanisms underlying the observed non-linearity is left as an interesting avenue for future research.

6.2.3 Sign non-linearities

We now turn to analyze asymmetries by sign. The baseline implementation uses $f(x) = |x|$ as a regressor. Figure 9 shows the IRFs for positive and negative shocks. By construction, under linearity, those two should coincide. Figure 9 shows that there is some asymmetry in the response in CPI. However, this is driven by negative shocks having a reversal in the

[36] Since we are using the shock directly as a regressor, we do not scale the responses to coincide in their effect on impact.

[37] In particular, if we take the baseline specification but replace CPI with core CPI, reversal effects are much smaller. Changing the set of controls to only real oil prices eliminates the reversal.

[38] For example, Berger and Vavra (2015).

[39] Given that these models have a drift because of depreciation, they would also imply sign non-linearities. However, the relative importance of both kinds will depend on how skewed the distribution of adjustment gaps is. If the skeweness is low, then size non-linearities will be more likely to be detected than sign non-linearities.

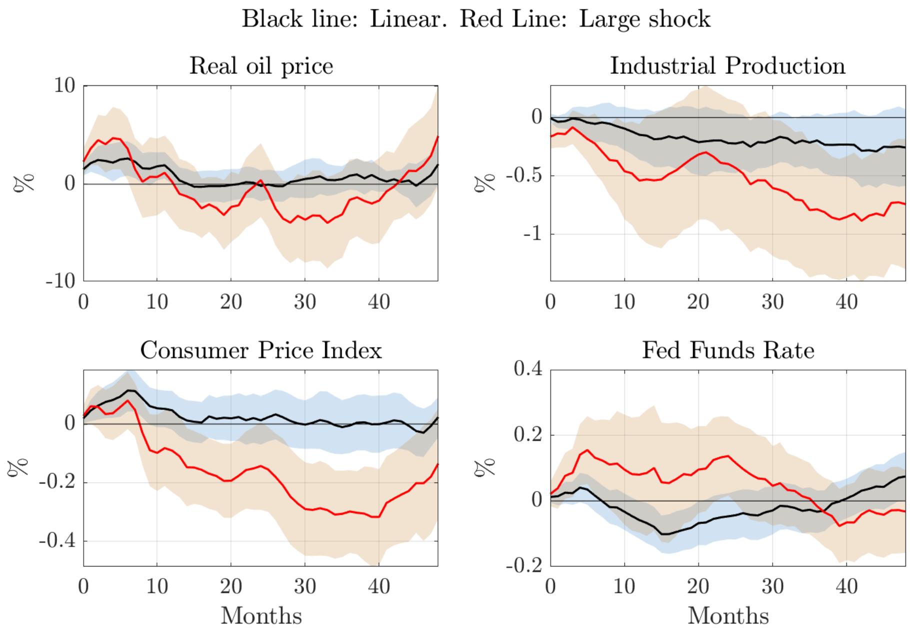

(a) IRFs estimated for the linear case (black) and for a large shock (red). Shaded areas indicate 95 percent confidence bands.

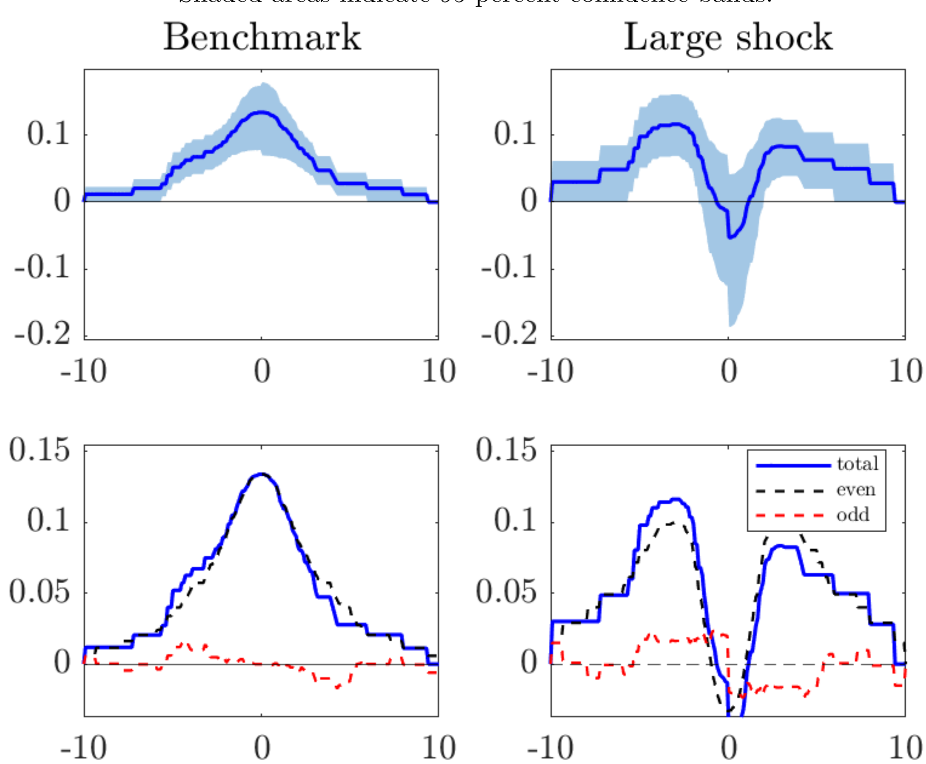

(b) Weights implicit in both the baseline case (black) and for a large shock (red).

Figure 8: The shaded area indicates a 90 percent confidence band compute via bootstrap, see Appendix D for details.

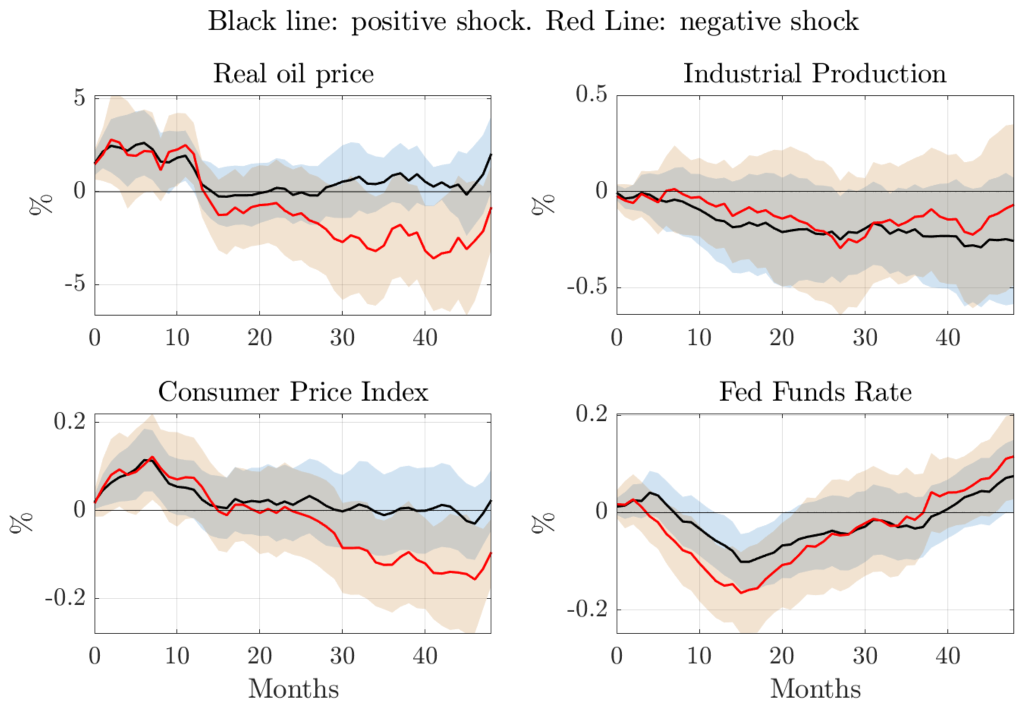

Figure 9: IRFs estimated for the positive (black) and negative (red) shocks. Shaded areas indicate 95 percent confidence bands.

path for real oil prices. Online Appendix E shows that if we used Core CPI instead of CPI (which includes energy prices directly), both IRFs are virtually identical. Focusing on the response of the Fed Funds rate, we can reject the null of linearity for some horizons between 6 and 12 months. However, this is not robust to the set of controls used nor to the use of critical values adjusted for multiple testing. Overall, both sets of IRFs look remarkably similar. There is little evidence of sign asymmetry, both in terms of statistical and economic significance. The fact that point estimates are remarkably close indicates that this is not due to low power.

7 Conclusion

The paper studies sign and size non-linearities in impulse responses when the true data-generating process is allowed to have an arbitrary functional form. The main takeaways of the presented findings can be summarized as follows: i) if the researcher is interested in non-linearities by sign, she should fit a local projection adding an even transformation of the shock of interest, such as $|x|$; ii) if the researcher is interested in non-linearities by size, she should use an odd function as the transformation. Although x^3 is an option, we argue that using a piecewise linear function is preferred; iii) one should always report the implicit weighting scheme that the linear regression is using Plagborg-Møller (2023), and its even-odd decomposition; iv) symmetry of the shock is required for disentangling both kinds of non-linearities, otherwise the procedure outlined above (and followed in some of the literature) would mix both kinds of non-linearities. Symmetry can be formally tested (Randles et al., 1980), and the plotting even-odd decomposition on weights is a simple way to report the degree of contamination in sample.

When applied to oil supply news shocks (Känzig (2021)), we find evidence of significant size non-linearities, but little evidence of sign non-linearities. Thus, the proposed method is capable of disentangling both types of non-linearities and therefore better inform modeling and policy.

Regarding potential avenues for future work, it is natural to extend these results to state-dependence. This is the most common kind of non-linearity studied in macroeconomics. Recent work on state-dependent local projections (Gonçalves et al. (2023)) shows that the actual specification of the DGP matters for interpreting the estimands of state-dependent local projections. Thus, a fully non-parametric approach such as the one presented in the present paper seems natural to think about this problem. Secondly, exploring the underlying mechanisms for the size non-linearity found in the empirical application is also a promising avenue of research.

References

Angrist, J. D., K. Graddy, and G. W. Imbens (2000): "The Interpretation of Instrumental Variables Estimators in Simultaneous Equations Models with an Application to the Demand for Fish," *The Review of Economic Studies*, 67, 499–527.

Ascari, G. and T. Haber (2021): "Non-Linearities, State-Dependent Prices and the Transmission Mechanism of Monetary Policy," *The Economic Journal*, 132, 37–57.

Auerbach, A. J. and Y. Gorodnichenko (2012): "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 4, 1–27.

Balke, N. S., S. P. Brown, and M. K. Yücel (2002): "Oil Price Shocks and the U.S. Economy: Where Does the Asymmetry Originate?" *The Energy Journal*, 23, 27–52.

Barnichon, R. and C. Brownlees (2019): "Impulse Response Estimation by Smooth Local Projections," *The Review of Economics and Statistics*, 101.

Barnichon, R., D. Debortoli, and C. Matthes (2022a): "Understanding the Size of the Government Spending Multiplier: It's in the Sign," *The Review of Economic Studies*, 89, 87–117.

Barnichon, R. and C. Matthes (2019): "Functional Approximation of Impulse Responses," *Journal of Monetary Economics*, 99, 41–55.

Barnichon, R., C. Matthes, and A. Ziegenbein (2022b): "Are the Effects of Financial Market Disruptions Big or Small?" *The Review of Economics and Statistics*, 104, 557–570.

Ben Zeev, N. (2020): "Identification of Sign-Dependency of Impulse Responses," *working paper*.

Ben Zeev, N., V. A. Ramey, and S. Zubairy (2023): "Do Government Spending Multipliers Depend on the Sign of the Shock?" *AEA Papers and Proceedings*, 113, 382–87.

Berger, D. and J. Vavra (2015): "Consumption Dynamics During Recessions," *Econometrica*, 83, 101–154.

Brinca, P., M. Faria-e Castro, M. H. Ferreira, H. A. Holter, and V. Nobrega (2023): "The Nonlinear Effects of Fiscal Policy," *Working Paper*.

Caramp, N. and E. Feilich (2022): "Monetary Policy and Government Debt," *UC Davis working papers*.

Davis, S. J. and J. Haltiwanger (2001): "Sectoral job creation and destruction responses to oil price changes," *Journal of Monetary Economics*, 48, 465–512.

Forni, M., L. Gambettu, N. Maffei-Faccioli, and L. Sala (2023): "Nonlinear Transmission of Financial Shocks: Some New Evidence," *Journal of Money, Credit and Banking*, forthcoming.

Gelman, A. and G. Imbens (2019): "Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs," *Journal of Business & Economic Statistics*, 37, 447–456.

Gonçalves, S., A. M. Herrera, L. Kilian, and E. Pesavento (2021): "Impulse response analysis for structural dynamic models with nonlinear regressors," *Journal of Econometrics*, 225, 107–130, themed Issue: Vector Autoregressions.

——— (2023): "State-dependent local projections," *FRB of Dallas Working Paper*.

——— (2024): "State-dependent local projections," *Journal of Econometrics*, 105702.

Gourieroux, C. and Q. Lee (2023): "Nonlinear Impulse Response Functions and Local Projections," *arXiv preprint arXiv:2305.18145*.

Hamilton, J. D. (1983): "Oil and the Macroeconomy since World War II," *Journal of Political Economy*, 91, 228–248.

——— (2003): "What is an oil shock?" *Journal of Econometrics*, 113, 363–398.

——— (2011): "Nonlinearities and the macroeconomic effects of oil prices," *Macroeconomic Dynamics*, 15, 364–378.

Hamilton, J. D. and A. M. Herrera (2004): "Comment: Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy," *Journal of Money, Credit and Banking*, 36, 265–286.

Herrera, A. M., L. Gupta Lagalo, and T. Wada (2011): "Oil Price Shocks and Industrial Production: Is the Relationship Linear?" *Macroeconomic Dynamics*, 15, 472–497.

Herrera, A. M., L. G. Lagalo, and T. Wada (2015): "Asymmetries in the response of economic activity to oil price increases and decreases?" *Journal of International Money and Finance*, 50, 108–133.

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

Kilian, L. and R. J. Vigfusson (2011): "Are the responses of the U.S. economy asymmetric in energy price increases and decreases?" *Quantitative Economics*, 2, 419–453.

Kilian, L. and X. Zhou (2020): "The Econometrics of Oil Market VAR Models," Tech. rep., Federal Reserve Bank of Dallas.

Känzig, D. R. (2021): "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, 111, 1092–1125.

Lan, H. and A. Meyer-Gohde (2013): "Solving DSGE models with a nonlinear moving average," *Journal of Economic Dynamics and Control*, 37, 2643–2667.

Montiel Olea, J. L. and M. Plagborg-Møller (2019): "Simultaneous confidence bands: Theory, implementation, and an application to SVARs," *Journal of Applied Econometrics*, 34, 1–17.

——— (2021): "Local Projection Inference Is Simpler and More Robust Than You Think," *Econometrica*, 89, 1789–1823.

Mork, K. A. (1989): "Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results," *Journal of Political Economy*, 97, 740–744.

Nakamura, E. and J. Steinsson (2018): "Identification in Macroeconomics," *Journal of Economic Perspectives*, 32, 59–86.

Newey, W. K. and K. D. West (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Plagborg-Møller, M. and C. K. Wolf (2022): "Instrumental Variable Identification of Dynamic Variance Decompositions," *Journal of Political Economy*, 130, 2164–2202.

Plagborg-Møller, M. (2023): "Tradeoffs Between Impulse Response Inference Methods," Presentation in the 2023 ASSA Annual Meetings.

Plagborg-Møller, M. and C. K. Wolf (2021): "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 89, 955–980.

Rambachan, A. and N. Shephard (2021): "When do common time series estimands have nonparametric causal meaning?" *working paper*.

Ramey, V. A. and D. J. Vine (2011): "Oil, Automobiles, and the U.S. Economy: How Much Have Things Really Changed?" *NBER Macroeconomics Annual*, 25, 333–368.

Ramey, V. A. and S. Zubairy (2018): "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 126, 850–901.

Randles, R. H., M. A. Fligner, G. E. Policello, and D. A. Wolfe (1980): "An Asymptotically Distribution-Free Test for Symmetry Versus Asymmetry," *Journal of the American Statistical Association*, 75, 168–172.

Ravn, M. O. and M. Sola (2004): "Asymmetric Effects of Monetary Policy in the United States," *Federal Reserve Bank of St. Louis Review*, 86.

Tenreyro, S. and G. Thwaites (2016): "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 8, 43–74.

Yitzhaki, S. (1996): "On Using Linear Regressions in Welfare Economics," *Journal of Business and Economic Statistics*, 14, 478–486.

Appendix

A Proofs

A.1 Lemma 1

This follows the proof technique in Angrist et al. (2000) and Rambachan and Shephard (2021) closely. First, consider a regression of the form:

$$y_{t+h} = \alpha + \beta_h f(\varepsilon_t) + u_{t+h}. \tag{A.1}$$

The OLS estimand is $\beta_h = \dfrac{Cov(y_{t+h}, f(\varepsilon_t))}{Var(f(\varepsilon_t))}$. We now dissect each term by parts. First, the numerator:

$$\begin{aligned}
Cov(y_{t+h}, f(\varepsilon_t)) &= \mathbb{E}\left[g(\varepsilon_t, \bullet)(f(\varepsilon_t) - \mathbb{E}[f(\varepsilon_t)])\right] \\
&= \mathbb{E}\left[\left(g(\underline{\varepsilon}, \bullet) + \int_{\underline{\varepsilon}}^{\varepsilon_t} g_\varepsilon(a, \bullet)da\right)(f(\varepsilon_t) - \mathbb{E}[f(\varepsilon_t)])\right] \\
&= \mathbb{E}\left[\int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \mathbf{1}_{\{a \leqslant \varepsilon_t\}} g_\varepsilon(a, \bullet)da(f(\varepsilon_t) - \mathbb{E}[f(\varepsilon_t)])\right] \\
&= \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} Cov\left(\mathbf{1}_{\{a \leqslant \varepsilon_t\}}, f(\varepsilon_t)\right)\mathbb{E}[g_\varepsilon(a, \bullet)]da.
\end{aligned}$$

From first to second line, we use the fundamental theorem of calculus and the fact that $g(x, \bullet)$ is differentiable.[40] From second to third, we use the Independence assumption, so $g(\underline{\varepsilon}, \bullet)$ drops out.[41] We also rewrite the integral using indicators. From fourth to fifth we exchange integrals and expectations, which we can do under our assumptions, we use independence, and write the resulting expectation in covariance notation. An identical derivation (swapping $g(x, \bullet)$ by $f(x)$) yields $Var(f(\varepsilon_t)) = \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} Cov\left(\mathbf{1}_{\{a \leqslant \varepsilon_t\}}, f(\varepsilon_t)\right) f'(a)da$. Thus $\beta_{f,h} = \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \omega_f(a)\mathbb{E}[g_\varepsilon(a, \bullet)]da$ where $\omega_f(a) = \frac{Cov(\mathbf{1}_{\{a \leqslant \varepsilon\}}, f(\varepsilon_t))}{\int_{\underline{\varepsilon}}^{\bar{\varepsilon}}\left[Cov(\mathbf{1}_{\{a \leqslant \varepsilon_t\}}, f(\varepsilon_t))\right]f'(a)da}$.

We now turn to the multivariate regression (4). Using the Frisch-Waugh-Lovell theorem, $\beta_{2,h} = \frac{Cov\left(y_{t+h}, f^\perp(\varepsilon_t)\right)}{Var(f^\perp(\varepsilon_t))}$ where $f^\perp(x) = f(x) - \frac{Cov(f(\varepsilon_t), \varepsilon_t)}{Var(\varepsilon_t)}x$ is the residual of linearly projecting

[40] If the function has kinks, then we abuse notation by writing $\int_{\underline{\varepsilon}}^{\varepsilon_t} g_\varepsilon(a, \bullet)da$ instead of $\sum_{k=1}^{K} \int_{\mathcal{I}_k \cap [\underline{\varepsilon}, \varepsilon_t]} g_\varepsilon(a, \bullet)da$ where $\mathcal{I}_k$ are the intervals where the function is differentiable.

[41] Since $\mathbb{E}\left[g(\underline{\varepsilon}, \bullet)\left(f(\varepsilon_t) - \mathbb{E}[f(\varepsilon_t)]\right)\right] = \mathbb{E}\left[g(\underline{\varepsilon}, \bullet)\right]\mathbb{E}\left[f(\varepsilon_t) - \mathbb{E}[f(\varepsilon_t)]\right] = 0$

f on a constant and ε_t. The weighted average representation follows from the discussion above.

Lastly, we show that the weights integrate to one and zero respectively. For $\omega_2(a)$, note that by Frisch-Waugh-Lovell theorem:

$$\omega_2(a) \propto Cov\left(\mathbf{1}_{\{a\leqslant\varepsilon_t\}}, f(\varepsilon_t) - \frac{Cov(f(\varepsilon_t),\varepsilon_t)}{Var(\varepsilon_t)}\varepsilon_t\right),$$

so

$$\int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \omega_2(a)da \propto \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} Cov(\mathbf{1}_{\{a\geqslant\varepsilon\}}, f(\varepsilon_t))da - \frac{Cov(f(\varepsilon_t),\varepsilon_t)}{Var(\varepsilon_t)} \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} Cov(\mathbf{1}_{\{a\geqslant\varepsilon_t\}}, \varepsilon_t)da.$$

Following the exact same derivation for the covariance term

$$\begin{aligned} Cov\left(f(\varepsilon_t),\varepsilon_t\right) &= \mathbb{E}\left[f(\varepsilon_t) - \mathbb{E}[f(\varepsilon_t)][\varepsilon_t - \underline{\varepsilon}]\right] \\ &= \mathbb{E}\left[(f(\varepsilon_t) - \mathbb{E}[f(\varepsilon_t)])\left[\int_{\underline{\varepsilon}}^{\varepsilon_t} da\right]\right] = \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} Cov(\mathbf{1}_{\{a\geqslant\varepsilon_t\}}, f(\varepsilon_t))da \end{aligned}$$

using the above result with $f(\varepsilon_t) = \varepsilon_t$ we get $Var(\varepsilon_t) = \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} Cov(\mathbf{1}_{\{a\geqslant\varepsilon_t\}}, \varepsilon_t)da$. Putting those two results together it is immediate that $\omega_2(a) \propto 0$. Equality follows from the constant of proportionality being bounded by the assumptions on finite moments for ε_t.

For the weights on β_1, note that:

$$\omega_1(a) = Cov\left(\mathbf{1}_{\{a\leqslant\varepsilon_t\}}, \varepsilon_t - \frac{Cov(f(\varepsilon_t),\varepsilon_t)}{Var(\varepsilon_t)} f(\varepsilon_t)\right) / Var\left(\varepsilon_t - \frac{Cov(f(\varepsilon_t),\varepsilon_t)}{Var(\varepsilon_t)} f(\varepsilon_t)\right)$$

But by using the same formulas as above:

$$\begin{aligned} \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \omega_1(a)da &= \left[Var\left(\varepsilon_t - \frac{Cov(f(\varepsilon_t),\varepsilon_t)}{Var(\varepsilon_t)} f(\varepsilon_t)\right)\right]^{-1} \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} Cov\left(\mathbf{1}_{\{a\leqslant\varepsilon_t\}}, \varepsilon_t - \frac{Cov(f(\varepsilon_t),\varepsilon_t)}{Var(\varepsilon_t)} f(\varepsilon_t)\right) da \\ &= \left[Var\left(\varepsilon_t - \frac{Cov(f(\varepsilon_t),\varepsilon_t)}{Var(\varepsilon_t)} f(\varepsilon_t)\right)\right]^{-1} \left[Var(\varepsilon_t) - \frac{Cov(f(\varepsilon_t),\varepsilon_t)}{Var(\varepsilon_t)} Cov(f(\varepsilon_t),\varepsilon_t)\right] = 1 \end{aligned}$$

A.2 Lemma 3

Consider first an even f. Under symmetry, $Cov(f(\varepsilon_t),\varepsilon_t) = 0$. Therefore, by FWL, the coefficient on β_2 is the same and the univariate regression coefficient. Thus $\omega_2(a) = Cov\left(\mathbf{1}_{\{a\leqslant\varepsilon_t\}}, f(\varepsilon_t)\right)/A$ where A is a positive constant. Define the auxiliar random variable

$X = -\varepsilon_t$. By symmetry, the marginal distribution of X is the same as the one of ε_t. Thus:

$$
\begin{aligned}
\omega_2(a) &= Cov\left(\mathbf{1}_{\{a \leqslant X\}}, f(X)\right)/A \\
&= Cov\left(\mathbf{1}_{\{a \leqslant -\varepsilon_t\}}, f(-\varepsilon_t)\right)/A \\
&= Cov\left(\mathbf{1}_{\{-a \geqslant \varepsilon_t\}}, f(\varepsilon_t)\right)/A \\
&= Cov\left(\mathbf{1}_{\{-a \geqslant \varepsilon_t\}} - 1, f(\varepsilon_t)\right)/A \\
&= -Cov\left(1 - \mathbf{1}_{\{-a \geqslant \varepsilon_t\}}, f(\varepsilon_t)\right)/A \\
&= -Cov\left(\mathbf{1}_{\{-a \leqslant \varepsilon_t\}}, f(\varepsilon_t)\right)/A = -\omega_2(-a)
\end{aligned}
$$

where from first to second line we used the definition of X, from second to third that f is even. From third to fourth, we added a constant to the covariance and flipped signs. Third to fourth, we used that $1 - \mathbf{1}_{\{-a \geqslant \varepsilon_t\}} = \mathbf{1}_{\{-a \leqslant \varepsilon_t\}}$, which follows from Assumption 4 since $\text{Prob}(\varepsilon_t = -a) = 0$.

Note that if the distribution of ε_t has a mass point at $-a$, then

$$
Cov\left(\mathbf{1}_{\{-a \leqslant \varepsilon_t\}}, f(\varepsilon_t)\right) = Cov\left(1 - \mathbf{1}_{\{-a > \varepsilon_t\}}, f(\varepsilon_t)\right) + Cov\left(\mathbf{1}_{\{-a = \varepsilon_t\}}, f(\varepsilon_t)\right)
$$

But since $Cov\left(\mathbf{1}_{\{-a = \varepsilon_t\}}, f(\varepsilon_t)\right) = Prob(\varepsilon_t = -a)\left(1 - Prob(\varepsilon_t = -a)\right) f(a)$, the mass point discontinuity is not a problem if $f(a) = 0$, which is the case for all non-linear transformations considered in the paper.

For the case of an odd f function, the weights depend on $f^{\perp}(\varepsilon_t) = f(\varepsilon_t) - \frac{Cov(f(\varepsilon_t), \varepsilon_t)}{Var(\varepsilon_t)}\varepsilon_t$. But since ε_t is an odd function, $f^{\perp}(\varepsilon_t)$ is a linear combination of odd functions, so it is odd. The proof concludes by using similar argument to the one made above for even function case.

Supplemental Appendix

This appendix contains supplemental material for the article "Disentangling sign and size non-linearities". Here we provide: (i) details and an extension of the illustrative example in B;(ii) Additional results on weighting schemes for different non-linear transformations for gaussian regressors in C; (iii) extensions on further topics of practical relevance, such as imperfect measurement of the shock, inclusion of controls, different estimation methods and inference in D; and (iv) further details and robustness exercises on the empirical application in section E.

B Details of Example 1

B.1 Microfoundations

We consider the textbook three equation New Keynesian model. Derivation of the IS curve is standard with discount factor shock. The monetary policy rule is $i_t = E_t[\pi_{t+1}] + \sigma^{-1} E_t[y_{t+1}]$, and ensures determinacy as long $\sigma^{-2} < 2$. The derivation of the Phillips curve is as follows. First, standard Calvo price-setting with a mark-up shock implies:

$$\pi_t = \lambda mc_t + \beta E_t[\pi_{t+1}] + \varepsilon_t^s$$

The non-linearity comes from the mapping between mc_t and y_t. In an economy in which the only input is labor, we are assuming that the only non-linearity comes from the labor market. Any kind of downward wage rigidity will generate non-linearities in this mapping. In general, we assume a general (possibly) non-linear mapping $mc_t = \tilde{c}(y_t)$. However, this nests several examples of interest.

Linear. if $\tilde{c}(y_t) = \lambda^{-1} \kappa y_t$, we arrive to the standard new Keynesian Phillips curve.

Downward wage rigidity. Assume TFP is constant. Then, following the standard derivation, we obtain that marginal cost (in log-deviations) is linear in wages, which are proportional to output gap $mc_t = w_t = (\sigma + \varphi) y_t$ where φ is the inverse Frisch elasticity. In that context, we add rigidity on real wages: they cannot fall by more than $\gamma(\sigma + \varphi)/\kappa$.

Then, we have that if $y_t \geqslant -\gamma/\kappa$, the usual derivation works. Otherwise, if output gap was required to be below $-\gamma/\kappa$, the downward rigidity binds. We assume that $\sigma c_t + \varphi n_t = w_t$ is replaced by $y_t = -\gamma/\kappa$: the intratemporal optimality condition does not hold anymore and there is involuntary unemployment. Note that market clearing and the production function still hold. The economy thus gets stuck in a regime in which demand shocks now adjust purely though quantities. Therefore, this generates a c function of the form used in the main text.

B.2 Weights for uniform distribution

Using the definition of the weights when the distribution is $U[-b, b]$ yields the following implicit weighting schemes for the non-linear regressor:

(i) **Linear only.** In this case the weights are:

$$\omega_1(a) = \frac{3}{4}(1 - (x/b)^2)\mathbf{1}_{\{-b \leqslant a \leqslant b\}} \tag{B.1}$$

(ii) **Linear and $f(\varepsilon_t) = (\varepsilon_t)^2$.** In this case, the weights on the linear term are given by (B.1), whereas the weights on the non-linear term are given by

$$\omega_2(a) = \frac{15}{8b^5}\left(ab^2 - a^3\right) \tag{B.2}$$

(iii) **Linear and $f(\varepsilon_t) = (\varepsilon_t)^3$.** The weights on the non-linear term are given by:

$$\omega_3(a) = \frac{35}{32b^7}\left(6b^2a^2 - 5a^4 - b^4\right) \tag{B.3}$$

All estimands follow from using the appropriate weights and marginal effect functions, and computing the corresponding integrals. Exact values for the regressors can be computed directly from the weights:

$$\alpha_2 = \int_{-b}^{b} \omega_2(a)\frac{\partial g_\pi}{\partial \varepsilon_t^d}(a)da = \frac{15}{32b}\kappa$$
$$\theta_2 = \int_{-b}^{b} \omega_3(a)\frac{\partial g_\pi}{\partial \varepsilon_t^d}(a)da = 0$$

B.3 Estimands when $\gamma > 0$

Consider the setting of Example 1, but now let $\gamma > 0$. The marginal effect of a demand shock on inflation is now given by:

$$\frac{\partial g_\pi}{\partial \varepsilon_t^d}(x) = \begin{cases} \kappa & \text{if } x > -\gamma \\ 0 & \text{if } x < -\gamma \end{cases} \tag{B.4}$$

Intuitively, now the downward wage rigidity only binds when the demand shock is negative, but large in magnitude. In that case, the f function now has a kink, but not at zero. How does this kind of non-linearity show up in our set-up? In this case, there is both a sign and a size non-linearity: negative shocks are different from positive shocks, because only negative shocks can make the economy hit the downward wage rigidity. On the other hand, this is true only for large negative shock, so there is a size asymmetry as well: small shocks never make the inequality binding, large shocks might.

In order to understand how this would manifest in the kind of regressions we run, consider the even-odd decomposition of the $\frac{\partial g_\pi}{\partial \varepsilon_t^d}$ function in this context. This is depicted in Figure B.1. As we can see, both the odd and the even part are non-constant. The odd part is similar to the top panel 1, but now there is a part of the function that evaluates to zero, for $a \in [-\gamma, \gamma]$. This is capturing the fact that there is asymmetry by sign, as discussed: negative shocks can hit the kink, in which the marginal response of inflation to demand is lower. Thus, positive shocks have higher marginal effect than negative shocks. On the other hand, the even part is not constant anymore: it now takes a different value for $a \in [-\gamma, \gamma]$ than for the rest of the domain. In particular, the value is higher for small shocks. This captures the fact that, if the shock is small, the economy does not hit the threshold at which the wage rigidity starts to matter. In contrast, large shocks may induce the economy to hit the threshold at which the downward rigidity is binding, and the impact on inflation is lower.

In this case, we would expects coefficients of U-shaped even regressors to have positive sign: positive shocks have higher marginal impact on inflation that negative shocks. On the other hand, for size related regressors, if we use an increasing odd regressor as x^3, we would expect the sign of the coefficient for that regressor to be negative: small shocks have a higher marginal effect on inflation than large shocks (because, in this world, small shocks are always partially accommodated via inflation unlike large shocks).

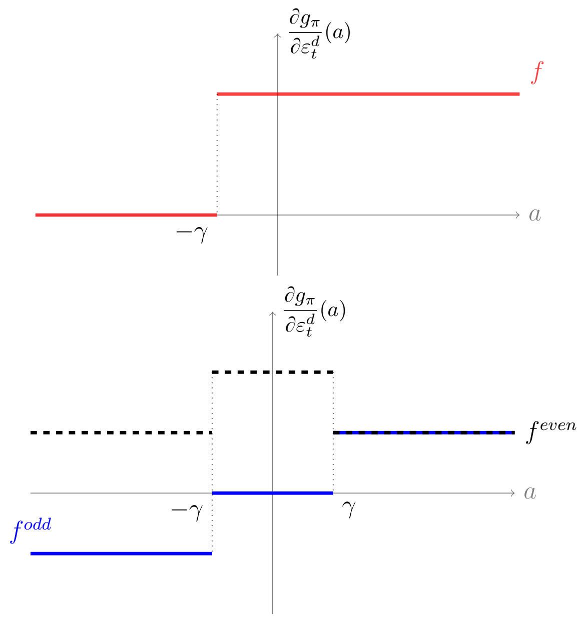

Figure B.1: The first panel depicts the derivative of the kinked function (red). The second panel depicts the even-odd decomposition of the function.

C Weights for different non-linear transformations and a Gaussian regressor

This section shows weights under a Gaussian distribution for different non-linear transformations of the shock. Assume $\varepsilon \sim N(0,1)$. Figure C.1 shows the weights generated by different non-linear transformations. Throughout this subsection, note that by symmetry $Cov(x, f(x)) = 0$ for any even function $f(x)$.

Linear. First, if we do not add any extra non-linear regressor, we can compute $w(a)$. It turns out that it equals the pdf of the standard normal. The top left panel in Figure C.1

Adding even as regressors. We turn now to the case in which either $f(x) = |x|$, $f(x) = x^2$ or $f(x) = |x|x^2$ are added as regressors. The top panel in Figure C.2 plots these functions. Note that because the shock distribution is symmetric, $Cov(\varepsilon, \varepsilon^2) = Cov(\varepsilon, |\varepsilon|) = 0$. Thus, the linear term and the extra regressor are orthogonal, which implies that the weights on β_1 are the same as depicted in the top left panel. Regarding the weights on the second regressor, these are depicted in the bottom left panel. As we showed, both functions are odd. Thus, if marginal effects were symmetric (i.e the derivatives were the same for a

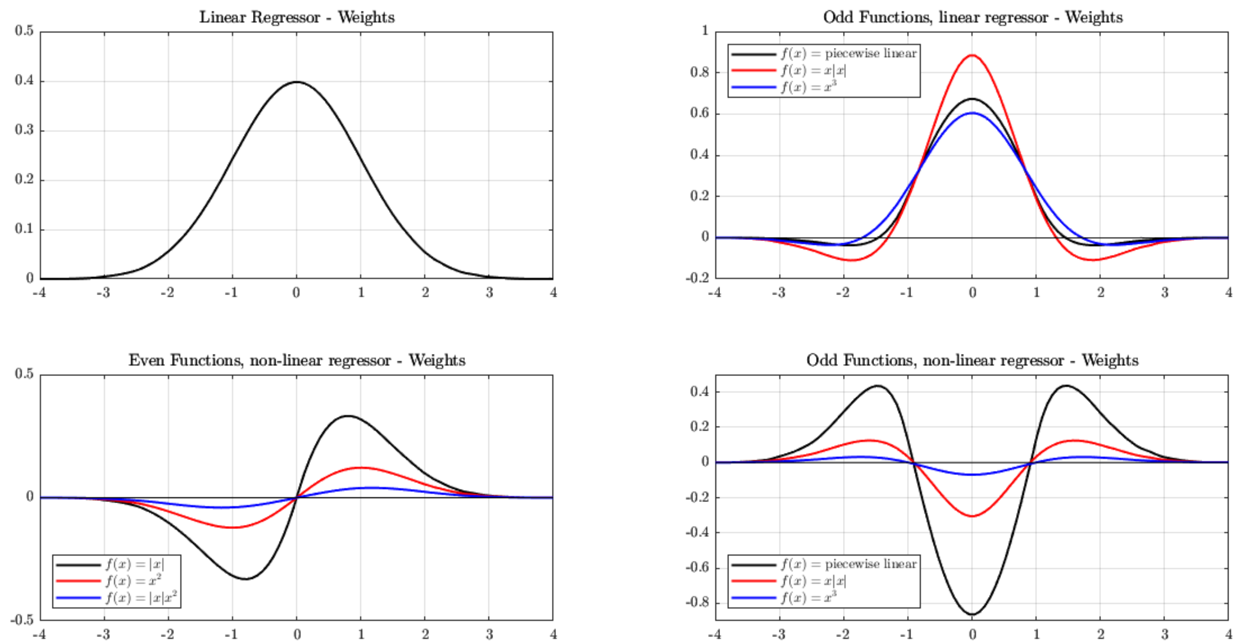

Figure C.1: The figure depicts the weights if the the regressor is standard normal. All panels depict weights implicit in a regression of the form $y_{j,t+h} = \beta_0 + \beta_1\varepsilon_t + \beta_2 f(\varepsilon_t) + u_{j,t+h}$. Top left panel: weights on β_1 with an even regressor. Since the standard normal is a symmetric distribution, $Cov(x, f(x)) = 0$ for any even f, so these weights are the same as in a regression that only included a linear term. Top right: weights on β_1 for different odd regressors. Bottom left: weights on β_2 for different even regressors. Bottom right: weights on β_2 for different odd regressors.

shock of size x than for a shock of size $-x$), this coefficient would be zero. On the other hand, if the marginal effect for positive shocks is, on average, higher than the one for negative shocks, we would have a positive β_2.

We see another pattern: the size of the weights in $|x|$ is higher than in x^2, which in turn are higher than in $f(x) = |x|x^2$. This is because the function $|x|$ is more sensitive (i.e its derivative is larger) for the values in which a normal distribution has more mass: values in $[-1, 1]$. The picture is even more steep if we use $|x|^\alpha$ is $\alpha < 1$.

Adding odd regressors. We now study what happens with the weights if we used a linear term and different odd regressors: $f(x) = 1_{x\leqslant -1}(x+1) + 1_{x\geqslant 1}(x-1)$, $f(x) = |x|x$ and $f(x) = x^3$. The first two regressors warrant some explanaition. First, $f(x) = |x|x$ is what one obtains if we take the values of x^2 for negative x, and reflect them along the x axis. That is, it has the same shape as in x^3, but with degree 2 instead of 3. For $f(x) = 1_{x\leqslant -1}(x+1) + 1_{x\geqslant 1}(x-1)$, this is a piecewise linear function that equals zero if

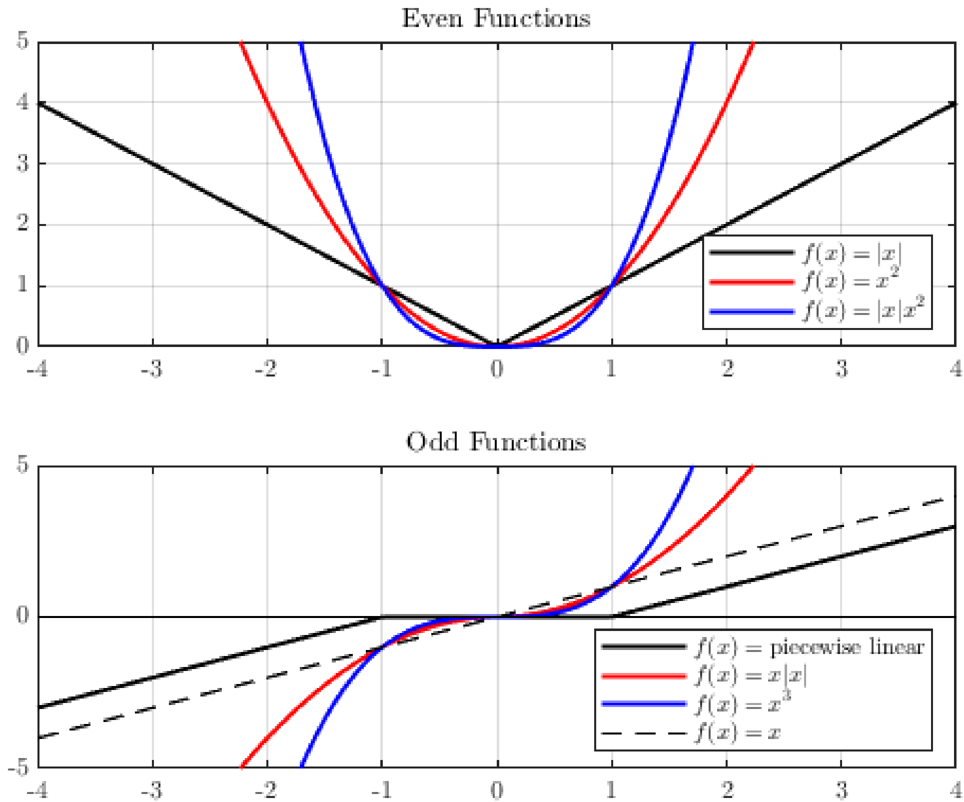

Figure C.2: Non-linear functions used as regressors. The top panel depicts different even functions, the bottom panel depicts different odd functions ($f(x) = x$ is also depicted for comparison).

$x \in [-1, 1]$. From both -1 and 1, the function continuous linearly, with a slope of 1. The bottom panel in Figure C.2 plots these functions.

Finally, the bottom right panel plots the weights on the coefficient on the non-linear term. As expected, $\omega_2(a)$ is an even function. This will capture size asymmetries by giving positive weight to large realizations in absolute value, and negative weights to small realizations in absolute value. Any asymmetry unrelated to size will be wiped out. Once again, the lower the order of the function, the more pronounced the differences in weighting.

D Implementation in practice: measurement, estimation and inference

Measurement. In the empirical application, we assume that there are many oil supply news shocks, but we are just measuring (perfectly) as subset of them. As long as there is no clear evidence of predictability of these shocks with past information, one can interpret measured series of proxies as the shock. This assumption is related to the one in Plagborg-Møller and Wolf (2022) for Forecast Variance Decompositions.[42]

One might be worried about measurement error contaminating the separation result. However, if the distribution of the measurement error is also assumed to be symmetric, this is not the case.[43] In order to see this, assume that we observe $z_t = \varepsilon_t + v_t$ where v_t is iid, independent of $\{\varepsilon_t\}$ and with a symmetric marginal distribution. A similar derivation as the one in Appendix A.1 shows that the weights on the non-linear regressor are:

$$\omega_2(a) = \frac{Cov(\mathbf{1}_{\{a \leqslant \varepsilon\}}, f^{\perp}(z_t))}{Var(f^{\perp}(z_t))} \tag{D.1}$$

Consider first the case in which $f(\cdot)$ is even. In that case, $f^{\perp}(z_t) = f(z_t)$. Then:

$$\begin{aligned}
\omega_2(a) = Cov(\mathbf{1}_{\{a \leqslant \varepsilon_t\}}, f(z_t))/Var(f(z_t)) &= Cov(\mathbf{1}_{\{a \leqslant -\varepsilon_t\}}, f(-\varepsilon_t + v_t))/Var(f(z_t)) \\
&= Cov(\mathbf{1}_{\{a \leqslant -\varepsilon_t\}}, f(-z_t))/Var(f(z_t)) = -\omega_2(-a).
\end{aligned}$$

The first equality comes from symmetry of ε_t. From first to second line, we use symmetry of v_t and replace to get $-z_t$. The last equality follows from the same steps as in the proof in Appendix A.2. Therefore, the presence of a symmetric measurement error does not affect the separation result. Of course, the estimand is now different, since it has weights (D.1). This the manifestation of classical attenuation bias in our context. Note also that these weights are not directly estimable, since they would require direct knowledge of ε_t.

Adding controls. It is straightforward to generalize the weights formulas adding controls. It is instructive to consider how to modify the weight formulas if we added controls,

[42] As argued in that paper, one of the bounds corresponds to the case where the shock is assumed to be measured perfectly. Since measurement error would produce attenuation bias, intuitively it only reduces the power to test for non-linearities.

[43] We specially thank Fernando Alvarez for suggesting this extension.

which is indeed the case in the empirical application. Generically, by controls we mean lags of y_t or other variables. Crucially, we still maintain the assumption that ε_t is independent of these controls. We are considering the case in which controls are added to reduce variance, not because of identification purposes. This is often the case in applied work, where external measurements of the shock are usually not predictable by past information.

In order to obtain formulas for the weights, we again invoke FWL. Define $x_t^* = x_t - \hat{x}_t$ where $\hat{x}_t = E^*[x_t|\text{controls}]$. We only consider linear projections on the controls because this is what is done in practice, any non-linearities can be controlled for by adding non-linear lagged terms are additional regressors. Then, by FWL:

$$\begin{aligned}\beta_1 = \frac{Cov(y_{t+h}^*, \varepsilon_t^*)}{Var(\varepsilon_t^*)} &= \frac{Cov(y_{t+h}^*, \varepsilon_t - \beta^* f(\varepsilon_t) - L(\text{ controls }))}{Var(\varepsilon_t^*)} \\ &= \frac{Cov(y_{t+h}^*, \varepsilon_t - \beta^* f(\varepsilon_t))}{Var(\varepsilon_t^*)}\end{aligned}$$

where β^* is the multiple regression coefficient of $f(\varepsilon_t)$ when residualizing ε_t, and $L(\cdot)$ is a linear function of controls. Using the same steps as before, we can get weights as:

$$\omega(a) = \frac{Cov(\mathbf{1}_{\{a \leqslant \varepsilon_t\}}, \varepsilon_t - \beta^* f(\varepsilon_t))}{Var(\varepsilon_t^*)}$$

Thus, it is essentially the same formula as above, but using the multiple regression coefficient for $f(\varepsilon_t)$. This the formula we use to compute the weights depicted in the main text. Note also that the interpretation changes slightly: we are now estimating the causal effect of ε_t on y_{t+h}^*, not on y_{t+h} directly.

Estimation. Estimation is carried by running local projections augmented with a non-linear term, and additional controls (lags of the outcome variable and other variables). Building on Plagborg-Møller and Wolf (2021), the estimand of this procedure could be alternatively implementing by ordering the linear and non-linear transformations first and second in a VAR, and the variable of interest third. Although the population estimand is the same with an inifinite lag length, a word of caution is in order: if we use a finite order VAR(p), then responses for horizons $h > p$ will be extrapolated from the estimated (linear) dynamics of the VAR. Now, if we are suspicious that there are underlying non-linearities in the data, interpolating from linear dynamics is less defensible. While it is true that under some particular forms of non-linearity this is indeed correct (see Gonçalves et al. (2021), Forni et al. (2023)), in general it does not need to hold. For example, it is well known that menu cost models feature size non-linearities in the response to monetary shocks. Furthermore, the propagation over time is differs depending on the size: large shocks are much less persistent than small shocks. Thus, this kinds of models feature different dynamic patterns for shocks of different sizes, which is hard to accomodate under linear dynamics.

INFERENCE. We use standard asymptotic inference for testing $\beta_{2,h} = 0$. Nothing in the derivation of the (robust) standard errors in the typical OLS setting relies on linearity of the true DGP: all what matters is heterogeneity in variances and autocovariances of terms like $\varepsilon_t u_{t-j}$, where ε_t is the regressor and u_{t-j} is the error in the projection.

To see this more clearly, simply define the estimands (and hence the estimation problem) in the GMM framework. Assume we are interested in regressions of the form:

$$y_{t+h} = \beta_0 + \beta_1 \varepsilon_t + \beta_2 f(\varepsilon_t) + u_{t+h} \tag{D.2}$$

Collect all regressors in a column vector $x_t = [1, \varepsilon_t, f(\varepsilon_t)]$, and the estimands as $\beta = [\beta_0, \beta_1, \beta_2]'$ such that:

$$\mathbb{E}\left[x_t(y_{j,t+h} - x_t'\beta)\right] = 0 \tag{D.3}$$

Under the assumption that $\mathbb{E}[x_t x_t']$ is invertible, β is identified. Furthermore, the in-sample estimate is given by the typical OLS formula. In order to compute the standard errors, we can use simple time series GMM asymptotic theory. Define $p(x, y, \beta) = x(y - x'\beta)$. Then, under regularity assumptions we can use the CLT to get:

$$\sqrt{T}(\hat{\beta} - \beta) \to N\left(0, (\mathbb{E}[x_t x_t'])^{-1} S_\infty (\mathbb{E}[x_t x_t'])^{-1}\right) \tag{D.4}$$

where $S_\infty = \Gamma_0 + \sum_{l=0}^{l=\infty} \Gamma_l + \Gamma_{-l}$, $\Gamma_l = \mathbb{E}[u_{t+h} u_{t+h+l} x_t x_{t+l}']$. Thus, any non-linearity in y_{t+h} matters as long as it introduces heteroskcedasticy and/or autocorrelation in $u_{t+h} x_t$. But standard inferencial procedures for local projections are already adapted to that. Thus, inference can be conducted as usual.

The fact that we use standard inference makes it easy to adapt to the procedure to multiple testing. In order to account for this, in our empirical application we present the bands both for the typical point-wise standard errors, and using the sup-t confidence intervals introduced in Montiel Olea and Plagborg-Møller (2019).

Confidence interval for weights. Given a measure of the shock, the weights can be computed in sample in the following way:

(i) Create a grid of values of $a \in \mathcal{A}$. Ideally the grid $\mathcal{A}$ should contain in the support of the sample distribution of the shock.

(ii) For each $a \in \mathcal{A}$, compute:

$$\hat{w}(a) = \widehat{Cov}(\mathbf{1}_{\{a \leqslant \varepsilon_t\}}, f(\varepsilon_t)) = T^{-1} \sum_{t=1}^{T} \mathbf{1}_{\{a \leqslant \varepsilon_t\}} \left(f(\varepsilon_t) - \hat{E}[f(\varepsilon_t)] \right) \tag{D.5}$$

with $\hat{E}[x_t] = T^{-1} \sum_{t=1}^{T} x_t$.

Note that we can alternatively express the weights as:

$$\hat{w}(a) = \hat{E} \left[\mathbf{1}_{\{a \leqslant \varepsilon_t\}} \left(f(\varepsilon_t) - \hat{E}[f(\varepsilon_t)] \right) \right] \tag{D.6}$$

As the weights are nothing more than a sample average, under the usual regularity conditions we can apply the CLT to obtain the asymptotic distribution:

$$T^{1/2}(\hat{w}(a) - w(a)) \to N \left(0, Var(\mathbf{1}_{\{a \leqslant \varepsilon_t\}} \left(f(\varepsilon_t) - \mathbb{E}[f(\varepsilon_t)] \right)) \right) \tag{D.7}$$

note that in principle, all autocovariances of $\mathbf{1}_{\{a \leqslant \varepsilon_t\}} \left(f(\varepsilon_t) - \mathbb{E}[f(\varepsilon_t)] \right)$ should appear. However, by our assumptions that random variable is uncorrelated over time.
The asymptotic normality of the estimator justifies a simple wild bootstrap approach to compute confidence bands for the errors:

(i) Fix a number of repetitions M. In each repetition, draw a sample $\{\varepsilon_t^*\}$ with replacement from the observed distribution of ε_t.

(ii) For each bootstrapped sample, compute the weights $w^*(a)$ as outlined above.

1. Repeat M times. For a two sided, α significance bands, use $[q^*_{\alpha/2}, q^*_{1-\alpha/2}]$ where q^*_γ is the γ-th quantile of the bootstrapped distribution of weights.

Variable	Large Shocks	Small Shocks	No Shock	Total
Real oil price	-8.82	3.26	-1.05	-1.16
	(38.03)	(28.21)	(31.36)	(31.73)
Industrial Production	0.40	2.13	2.36	2.11
	(4.98)	(3.32)	(3.90)	(3.98)
Consumer Price Index	2.70	2.66	2.64	2.65
	(1.44)	(1.14)	(1.34)	(1.31)
Fed Funds Rate	3.65	3.85	3.93	3.88
	(2.76)	(2.93)	(3.08)	(3.01)
N	45	70	292	420

Table E.1: Balance Table for large and small shocks. A large magnitude shock is defined as being below $20-th$ (these are negative, large magnitude shocks) or above the $80-th$ percentile of the conditional distribution of the shocks, conditional on being non-zero. Small shocks are defined as values of the shock between the $20-th$ and the $80-th$ percentile of the shock conditional on non-zero shocks. The row "No Shock" contains averages for observations where the shock is not observed (it is zero). The table shows the average on the year on year percent variation of the variable listed at $t-1$ for a shock that hits at t. Numbers in parenthesis are standard errors. Fed Funds rate is levels.

E Empirical application

E.1 Balance between large and small shocks

Table E.1 shows balance between "small" (i.e shocks less than the bound in absolute value) and "large" shocks (shocks above the bound). The values for each variable indicate the interanual variation in the variable for the period $t-1$ (for a shock hitting at time t).[44] As we can see, the state of the economy when a small, large or no shock hits is comparable. This is to be expected, since the shock is unpredictable by using information at time $t-1$.

E.2 Asymmetry by size: robustness

We conducted the following robustness checks. We present results with the most relevant differences from the baseline specification here. Results not reported here are available under request.

- Figure E.1 shows the results when using only real oil price as a control instead of the full set of oil market variables. The response of CPI does not display significant

[44]Except for the Fed Funds Rate, which is in levels.

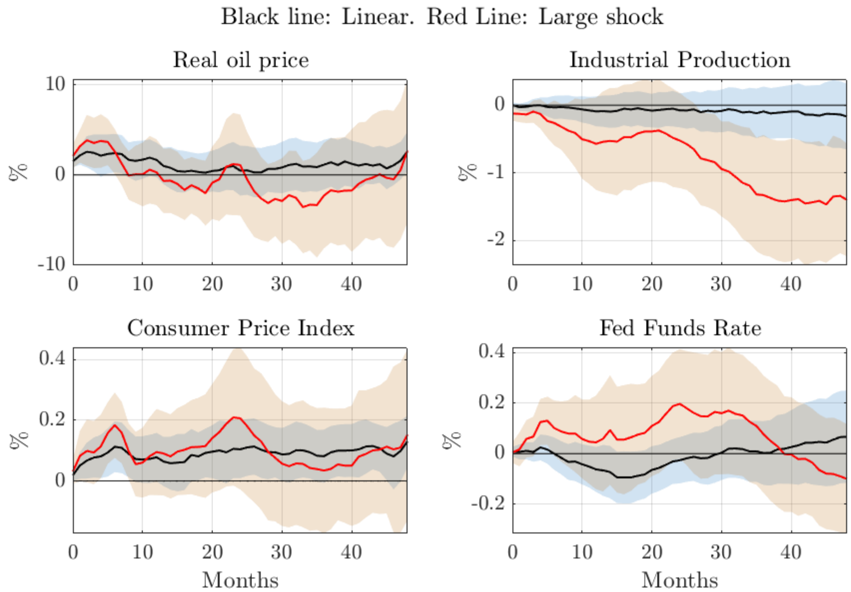

Figure E.1: IRFs estimated for the linear case (black) and for a large shock (red). Shaded areas indicate 95 percent confidence bands. Only lags of the outcome variable and real oil price are used as controls.

non-linearity, as explained in the main text.

- Figure E.2 shows testing results when $f(x) = x^3$ is used as a the non-linear transformation. They are the same as in the case of are the same regardless of whether we use the piecewise linear function.

- Figure E.3 shows testing results are robust when values adjusted for multiple testing (Montiel Olea and Plagborg-Møller (2019)). Results are similar to the case on the main text. In particular, the non-linearity in the initial response of industrial production is significant at 95%, whereas the non-linearity in the response of the Fed Funds rate is significant at 90%.

- Results are robust to using 12,18,24 lags (using the same variables as in the main text as controls).

- Standard errors are the virtually the same if we use Heteroskedasticy (but not AR) robust SE instead of Newey-West with lag-length set to the horizon of the local projection.

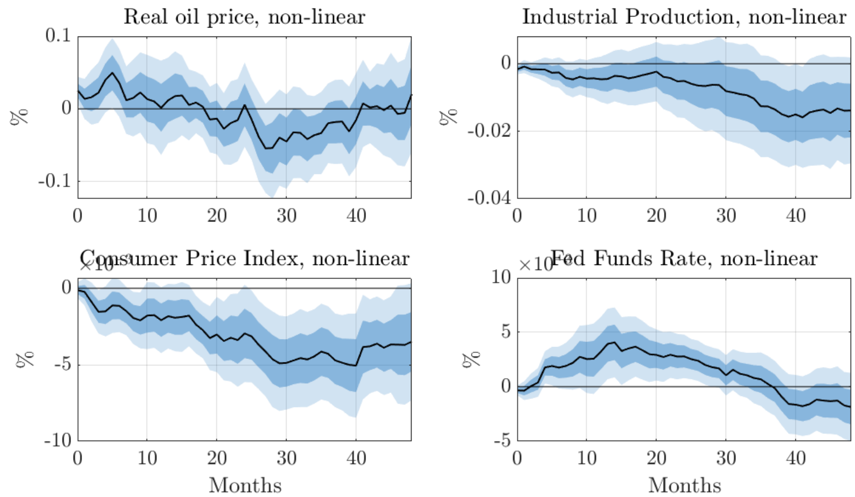

Figure E.2: Estimates of the coefficients on the non-linear terms obtained via Local Projections, i.e each panel depicts the sequence $\hat{\beta}_{j,h,2}$ that arises from estimating: $y_{j,t+h} = \alpha_{j,h} + \hat{\beta}_{j,h,1}\varepsilon_t^{oil} + \hat{\beta}_{j,h,2}(\varepsilon_t^{oil})^3 +$ controls $+u_{j,t+h}$ for each different outcome variable j, and for horizons $h = 0, \ldots, 48$. Shaded areas indicate 68 and 95 percent confidence bands respectively.

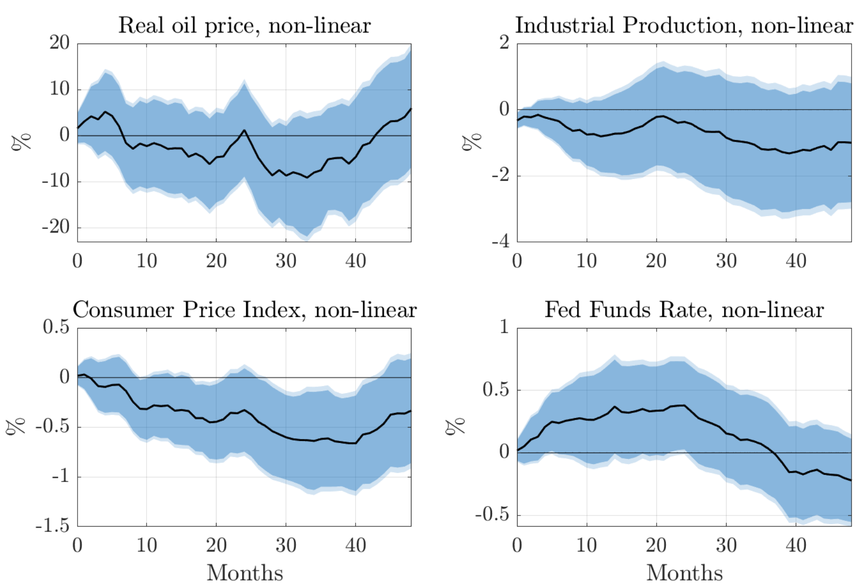

Figure E.3: The figure depicts the coefficients on the non-linear term (odd piecewise linear function) under the baseline specification. Shaded areas indicate joint confidence bands using the sup-t statistic of Montiel Olea and Plagborg-Møller (2019). Dark (light) shade indicates 90% (95%) confidence bands.

E.3 Asymmetry by sign

Figure E.4 shows the coefficients on the non-linear term. As explained in the main text, some significant coefficients appear for CPI and real oil price in long horizons. There are also some significant differences in Fed Funds rate in short horizons, but this is not robust to alternative selection of controls. The response of industrial production does not feature any significant non-linearity. Figure E.5 reports the results when Core CPI is used instead of CPI. As explained in the main text, the difference is now quite small and non-significant. Figure E.6 shows the implicit weights on the non-linear coefficient. As expected, positive and negative shocks get weights of different sign. Finally, Figure E.7 shows the weights on the transformed coefficients. The coefficient on positive (negative) shocks has almost all the weight on positive (negative) shocks. Thus, the interpretation of these coefficients as some average of positive (negative) effects is supported by the implicit weighting scheme being used.

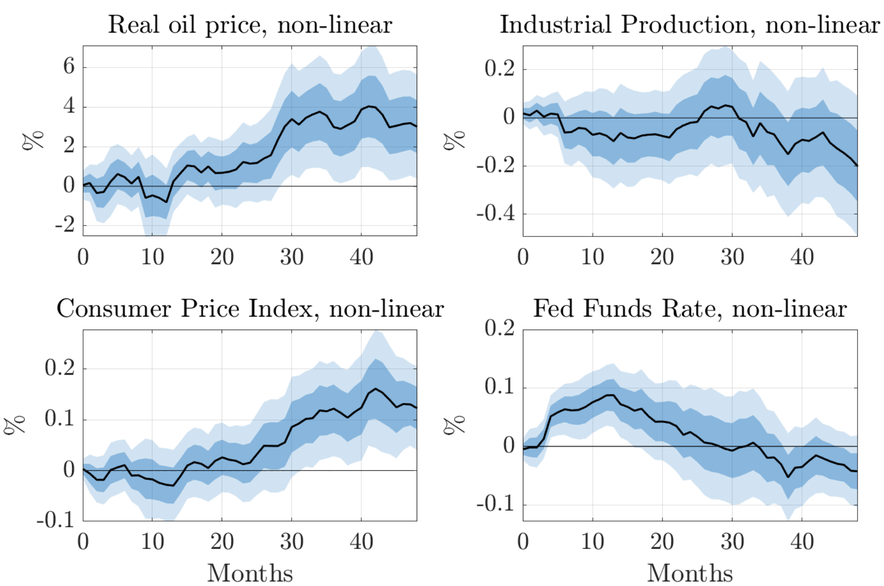

Figure E.4: Estimates of the coefficients on the non-linear terms. These are obtained by as $\hat{\beta}_{j,h,2}$ in the regression $y_{j,t+h} = \alpha_{j,h} + \hat{\beta}_{j,h,1}\varepsilon_t^{oil} + \hat{\beta}_{j,h,2}|\varepsilon_t^{oil}| +$ controls $+ u_{j,t+h}$ for each different outcome variable j, and for horizons $h = 0, \ldots, 48$. Shaded areas indicate 68 and 95 percent confidence bands respectively.

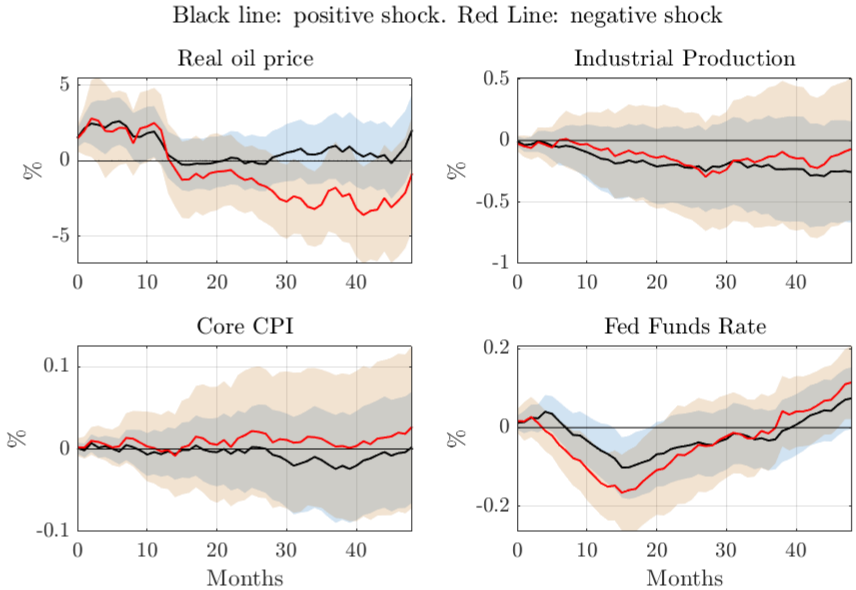

Figure E.5: IRFs estimated for the positive (black) and negative (red) shocks. Core CPI is used instead of CPI. Shaded areas indicate 95 percent confidence bands.

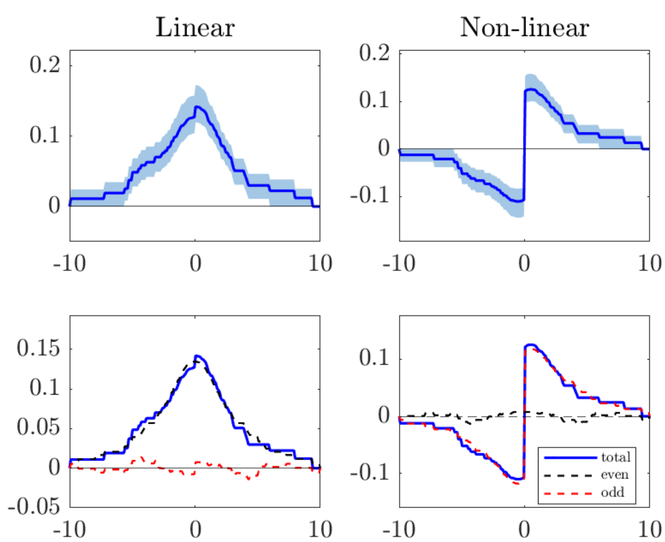

Figure E.6: Weights implicit in the Local Projection $y_{j,t+h} = \alpha_{j,h} + \hat{\beta}_{j,h,1}\varepsilon_t^{oil} + \hat{\beta}_{j,h,2}|\varepsilon_t^{oil}| +$ controls $+ u_{j,t+h}$. The shaded area indicates a 90 percent confidence band compute via bootstrap.

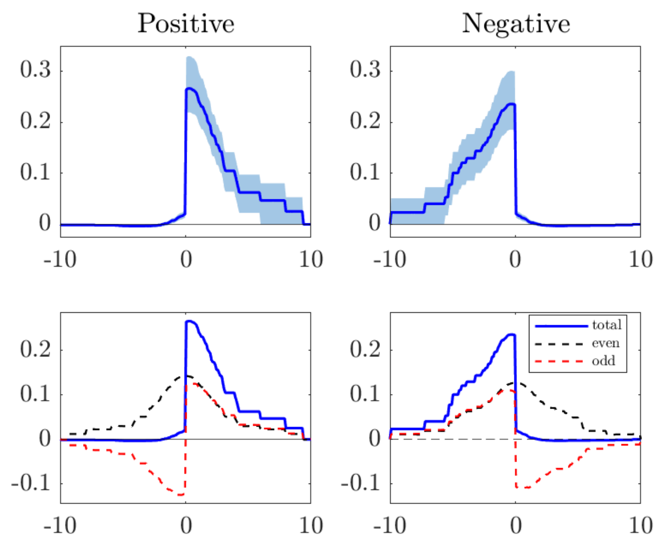

Figure E.7: Weights implicit in estimates for negative and positive shocks. These are obtained by as $\hat{\beta}_{j,h,1} - \hat{\beta}_{j,h,2}$ (negative) $\hat{\beta}_{j,h,1} + \hat{\beta}_{j,h,2}$ and (positive) from running $y_{j,t+h} = \alpha_{j,h} + \hat{\beta}_{j,h,1}\varepsilon_t^{oil} + \hat{\beta}_{j,h,2}|\varepsilon_t^{oil}| +$ controls $+ u_{j,t+h}$ for each different outcome variable j, and for horizons $h = 0, \ldots, 48$. The shaded area indicates a 90 percent confidence band compute via bootstrap.