

Demand Composition and the Strength of Recoveries[†]

Martin Beraja

Christian K. Wolf

MIT & NBER

MIT & NBER

February 4, 2022

Abstract: Recoveries from demand-driven recessions with expenditure cuts concentrated in services or non-durables tend to be weaker than recoveries from recessions more biased towards durables. Intuitively, the smaller the bias towards more durable goods, the less the recovery is bolstered by pent-up demand. We first show that this prediction holds in many standard multi-sector business-cycle models. We then document two pieces of supporting evidence in U.S. data: pent-up demand effects appear present in aggregate time series, and the prediction on recovery strength is consistent with the cross-section of regional business cycles. Next we use (a) a semi-structural shift-share and (b) a structural model to quantify the effect of varying demand composition on recovery dynamics, and find it to be large. We thus also discuss implications for optimal stabilization policy.

Keywords: durables, services, demand recessions, pent-up demand, shift-share design, recovery dynamics, COVID-19. *JEL codes:* E32, E52

[†]Email: maberaja@mit.edu and ckwolf@mit.edu. We received helpful comments from George-Marios Angeletos, Gadi Barlevy, Florin Bilbiie, Ricardo Caballero, Lawrence Christiano, Martin Eichenbaum, Peter Ganong, François Gourio, Basile Grassi, Erik Hurst, Greg Kaplan, Andrea Lanteri, Jennifer La'O, Alisdair McKay, Simon Mongey, Ernesto Pasten, Matt Rognlie, Alp Simsek, Ludwig Straub, Silvana Tenreyro, Nicholas Trachter, Gianluca Violante, Iván Werning, Johannes Wieland (our discussant), Tom Winberry, Nathan Zorzi and seminar participants at various venues, and we thank Isabel Di Tella for outstanding research assistance.

1 Introduction

Basic consumer theory suggests that “pent-up demand” effects should be stronger for more durable goods (Mankiw, 1982; Caballero, 1993). When a consumer decides against a durable purchase in the midst of a recession — like a car or furniture — she simply postpones that spending for later. This reversal in spending is likely to be weaker or even absent for services and non-durables — the lost spending on haircuts, food away from home or gasoline may simply be foregone. What, if anything, does this logic imply for how differences in household expenditure composition across recessions affect subsequent recovery dynamics?

In this paper we argue that, because of the pent-up demand mechanism, recoveries from demand-driven recessions with spending cuts concentrated in services or non-durables will tend to be weaker than recoveries from recessions more biased towards durables. Our analysis begins with a standard multi-sector business-cycle model in which demand composition varies with (i) long-run sectoral expenditure shares and (ii) the sectoral incidence of shocks.¹ In this environment, the prediction holds if and only if, conditional on an aggregate demand shock to all categories of spending (e.g., a monetary policy shock), expenditures on more durable categories revert back faster. This testable condition on impulse responses receives ample support in U.S. data. Regional U.S. business cycles are similarly in line with the prediction: during the Great Recession and the COVID-19 pandemic, states where the local recession was less concentrated in durables saw weaker recoveries. To quantify the impact of varying demand composition on recovery strength, we then use (a) a semi-structural impulse response shift-share design and (b) a structural model estimated via impulse response matching. We find that empirically relevant differences in recession spending composition can have a sizable effect on recovery strength. In light of this, we conclude by discussing implications for optimal stabilization policy.

To transparently illustrate the pent-up demand mechanism, we first of all consider a stylized two-sector business-cycle model with transitory shocks, fully demand-determined output (e.g., due to perfectly rigid prices), and a representative household with separable utility over the two goods, called “durables” and “services.” The only difference between the two goods is that the durables stock depreciates at rate $\delta < 1$ while services depreciate instantly.² We then add three reduced-form demand shocks: one for each sector’s good, and one to aggre-

¹As we show in Section 4.1, long-run sectoral expenditure shares vary widely across countries and, in the U.S., past recessions have differed greatly in their sectoral spending composition.

²In the model, we make no distinction between services and non-durables, treating them interchangeably.

gate demand. In this environment, much previous research has established that — because of their higher intertemporal substitutability — durable goods *amplify* output declines in recessions (e.g. Barsky et al., 2007). We instead focus on how pent-up demand for durables affects the *shape* of dynamic responses to demand shocks. We ask: given a recession of a certain magnitude, how should we expect the recovery dynamics to differ depending on the recession’s sectoral spending composition?

In this simple model, the recovery is invariably weaker when the recession’s spending composition is more biased towards services. To show this, we first establish that, following an arbitrary combination of aggregate and sectoral demand shocks, the impulse response of durable expenditures is always Z-shaped — with a fraction $1 - \delta$ of the initial decline at time $t = 0$ reversed at time $t = 1$ — while that of services is V-shaped — spending declines at $t = 0$, and then returns to baseline at $t = 1$. The intuition is as follows. The household exits the recession with a durable stock below target. Her subsequent flow expenditure must thus overshoot to replenish this stock. For services, on the other hand, there are no such pent-up demand effects. Our claim on recovery strength is then immediate: when services account for a share ω of the output decline at $t = 0$, the $t = 1$ overshoot of output is equal to a fraction $(1 - \omega)(1 - \delta)$ of the initial drop. The cumulative impulse response of output normalized by its trough — a measure of the weakness of the recovery — is then equal to $1 - (1 - \omega)(1 - \delta)$, so recoveries are weaker for a larger ω . This result holds irrespective of whether ω is large due to (i) a high long-run expenditure share of services, or (ii) a particular realization of shocks that decreases the relative demand for services.

We then relax many of our stark simplifying assumptions and consider a richer class of business-cycle models featuring: persistent shocks; arbitrary adjustment costs on durables; partially sticky prices and wages; an arbitrary number of goods varying in their durability and adjustment costs; net substitutability between goods; incomplete markets and hand-to-mouth households; and informational frictions in household expectations. In this extended environment our conclusion on recovery dynamics does not go through automatically. However, we can prove that, if the monetary authority’s policy rule is neutral (in the sense of fixing the ex-ante expected real rate of interest), then the prediction holds *if and only if*, conditional on a contractionary aggregate demand shock, the cumulative impulse response of durables spending (normalized by the spending trough, nCIR for short) is smaller than the corresponding nCIR for services spending. In words, this condition says that durables expenditure reverts back to trend (or overshoots) faster than services expenditure.

Our key testable condition for pent-up demand effects to be strong enough — and so the

prediction to hold — receives ample support in U.S. data. As our main empirical test, we study sectoral expenditure dynamics following aggregate monetary policy shocks identified using conventional time-series methods. For our purposes, such monetary policy shocks are close to the ideal experiment: they are a particular example of the aggregate demand shocks required by our testable condition, and much previous work agrees on their identification and macroeconomic consequences. The estimated impulse responses are in line with the pent-up demand logic: durable expenditures decline and then feature a strong overshoot, while non-durables and in particular services instead return to baseline from below, with the services nCIR around 88 percent larger than that of durables.³

We complement this test of the pent-up demand mechanism with evidence on the relation between demand composition and subsequent recovery strength, leveraging data on regional business cycles in the U.S. Cross-regional analysis is particularly well-suited for our purposes because it addresses a key identification challenge. Unlike time series comparisons across recessions, the regional analysis implicitly controls for aggregate confounders that are likely to have varied across business-cycle episodes — e.g., the aggregate policy response or the persistence of the underlying shocks — and so could obscure our prediction on the role of demand composition. Our main result is that, for both the Great Recession as well as the COVID-19 pandemic, states where durables accounted for a smaller share of output changes during the recession also saw a relatively weaker subsequent recovery.

Given this supporting evidence, we proceed to *quantify* the effects of demand composition on the strength of the subsequent recoveries. We do so in two ways.

1. The first approach is a semi-structural shift-share: we take our estimated sectoral expenditure impulse responses to monetary policy shocks, and then simply re-weight them to give us an aggregate contraction of a given severity and with a given sectoral composition. For this exercise, we choose those sectoral weights in line with empirically observed (i) cross-country variation in sectoral expenditure shares and (ii) U.S. cross-recession variation in sectoral spending composition. We prove that, under the same conditions as those required for sufficiency of our empirical test, this seemingly naive summing of re-weighted impulse responses in fact gives valid counterfactual predictions for recovery dynamics when varying (i) long-run sectoral expenditure shares and (ii) sectoral shock incidence.

³For further empirical support, we document similar sectoral expenditure patterns for more disaggregated spending categories — e.g., semi-durables officially categorized as non-durables — as well as following (i) uncertainty shocks (Basu & Bundick, 2017), (ii) oil shocks (Hamilton, 2003), and (iii) reduced-form forecast errors of sectoral output.

2. Our second approach is fully structural, relying on an extended model that violates the conditions required by the shift-share. Similar to the classical analysis of Christiano et al. (2005), we estimate the model by matching the estimated sectoral impulse responses. We then use the model to compute recovery dynamics in economies with varying long-run expenditure shares and subject to different sectoral demand shock combinations.

The two approaches paint a consistent picture: the effects of recession spending composition on recovery strength are robustly large. For example, the nCIR of output in a recession that is largely concentrated on services is estimated to be about 70 to 90 per center larger than that of an average (largely durables-led) recession. Similarly, moving from the U.S. to a more durables-intensive economy (e.g., Canada), the output nCIR to a common aggregate demand shock decreases by about 15 per cent.

In light of this quantitative relevance, we conclude the paper with a brief discussion of optimal stabilization policy. We show that our two main sources of heterogeneity in demand composition — (i) differences in long-run sectoral expenditure shares and (ii) sectoral shock incidence — actually have rather different implications for optimal policy design. First, in an economy subject only to aggregate (i.e., not sectoral) demand disturbances, the optimal path of interest rates turns out to be completely independent of long-run expenditure shares. Intuitively, changes in shares affect not only the transmission of exogenous aggregate demand shocks, but also that of changes in the nominal interest rate. In our model, these two effects exactly offset, leaving optimal monetary policy unaffected. Second, in the face of contractionary *sector-specific* demand shocks, the monetary authority should optimally ease for longer the greater the shock’s bias towards the services or non-durables sector, exactly as expected given our discussion of the pent-up demand mechanism.

LITERATURE. This paper relates and contributes to several strands of literature.

First, we build on a long literature that studies the role of durable consumption in shaping business-cycle dynamics. One line of work (e.g. Erceg & Levin, 2006; Barsky et al., 2007) models durable adjustment at the intensive margin, embeds it into conventional business-cycle models with nominal rigidities (Christiano et al., 2005), and emphasizes the implications for recession severity. A separate strand argues that extensive-margin adjustments imply state-dependent shock elasticities (Berger & Vavra, 2015) and affect the transmission of monetary policy (McKay & Wieland, 2020). We study a different question — the impact of demand composition on recovery dynamics — while relying on a tractable intensive-margin set-up. In addition to distilling the economic forces at work, this analytical framework is

useful because it allows us to: (i) identify a key testable condition for the main prediction, i.e., the ranking of impulse responses; (ii) formally justify our measurement approaches, i.e., the semi-structural shift-share as well as impulse response matching à la Christiano et al.; and (iii) derive implications for optimal monetary policy.⁴ Mirroring the first part of our empirical analysis, Erceg & Levin (2006) and McKay & Wieland (2020) also show that durables spending tends to overshoot after monetary shocks. We offer additional insights by comparing impulse responses across different consumption categories; using them for our quantification exercises; and providing further supporting cross-sectional evidence.

Second, a large literature considers the business-cycle implications of sectoral heterogeneity on the production side. One branch highlights heterogeneity in nominal rigidities across sectors (Carvalho, 2006; Barsky et al., 2007; Nakamura & Steinsson, 2010); another one incorporates rich network structures (Carvalho & Grassi, 2019; Bigio & La’o, 2020), sometimes combined with nominal rigidities (Pasten et al., 2017; Farhi & Baqaee, 2020; Rubbo, 2020; La’O & Tahbaz-Salehi, 2020). We instead highlight the importance of heterogeneity on the *demand side*, sorting goods and sectors by their durability.

Third, many papers have sought to understand the determinants of strength and shape of recoveries. Popular mechanisms include: the nature of shocks (Galí et al., 2012; Beraja et al., 2019), structural forces (Fukui et al., 2018; Fernald et al., 2017), secular stagnation (Hall, 2016), social norms (Coibion et al., 2013), changes in beliefs (Kozlowski et al., 2020), and labor market frictions (Schmitt-Grohé & Uribe, 2017; Hall & Kudlyak, 2020). We contribute to this literature by emphasizing the importance of changes in demand composition.

Finally, our quantification exercise and U.S. regional evidence relate to recent work on the sectoral incidence of the COVID-19 pandemic (Chetty et al., 2020; Cox et al., 2020; Guerrieri et al., 2020), and shapes of the recovery (Gregory et al., 2020; Reis, 2020). While predicting the economic recovery from COVID-19 is a complex endeavor due to the many channels at play, our results shed light on the economics of one particular mechanism: pent-up demand.

OUTLINE. Section 2 provides analytical characterizations of business-cycle dynamics in a multi-sector general equilibrium model with demand-determined output. Section 3 presents our empirical evidence. Section 4 quantifies the effect of demand composition on recovery strength. Finally, Section 5 discusses implications for optimal stabilization policy. Section 6 concludes, with supplementary details and proofs relegated to several appendices.

⁴On the investment side, the same reversal effects are discussed in Appendix B.1 of Rognlie et al. (2018).

2 Pent-up demand and recovery dynamics

This section presents our main theoretical results on demand composition and the strength of recoveries. Section 2.1 outlines the model. Sections 2.2 and 2.3 illustrate the pent-up demand mechanism in a stripped-down variant and present implications for recovery dynamics. Finally, Sections 2.4 and 2.5 extend results to the full model and discuss generalizations.

2.1 Model

We consider a discrete-time, infinite-horizon economy populated by a representative household, monopolistically competitive retailers, and a government. Households consume durables and services (or non-durables).⁵ The only source of risk are shocks to household preferences over consumption bundles. We will throughout focus on the linearized model solution.

HOUSEHOLDS. Household preferences over services s_t , durables d_t and hours worked ℓ_t are represented by

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \{u(s_t, d_t; b_t) - v(\ell_t; b_t)\} \right],$$

where we assume

$$u(s, d; b) = \frac{e^{b^a + b^s} \tilde{\phi}^\gamma s^{1-\gamma} + e^{\alpha(b^a + b^d)} (1 - \tilde{\phi})^\gamma d^{1-\gamma} - 1}{1 - \gamma}, \quad v(\ell; b) = e^{b^a + \phi b^s + (1-\phi)b^d} \chi \frac{\ell^{1+\frac{1}{\varphi}}}{1 + \frac{1}{\varphi}},$$

and b_t^a is an aggregate demand shock to all categories of spending, while $\{b_t^s, b_t^d\}$ are sectoral services and durables demand shocks, respectively. We interpret these shocks as simple reduced-form stand-ins for more plausibly exogenous shocks to (sectoral) household spending — e.g., increased precautionary savings due to greater income risk ($b^a < 0$) or increased fear of consuming certain services during a pandemic due to greater infection risk ($b^s < 0$).⁶ The scaling factor α is chosen to ensure that, in the flex-price limit of our economy, the aggregate demand shock b_t^a has no real effects on equilibrium quantities (to first order), instead

⁵In this section we speak of services and non-durables interchangeably. We will distinguish them in our empirical measurement in Section 3 and the quantification in Section 4.

⁶(A.5)-(A.6) in Appendix A.1 reveal that our preference shocks may equivalently be interpreted as shocks to consumption taxes or shadow prices (and so as “wedges” in optimality conditions (Chari et al., 2007)).

only moving the path of real interest rates.⁷ The scalings $\{\phi, 1 - \phi\}$ in the disutility of hours worked ensure that a combined shock $b_t^d = b_t^s$ is isomorphic to an aggregate demand shock b_t^a of the same magnitude.⁸ Finally, we note that preferences are separable over services and durables, as in Barsky et al. (2007). We do so mainly for expository purposes; later, in Section 2.5, we show that all of our results are actually strengthened in the empirically relevant case of net substitutability (as in Berger & Vavra, 2015).

Households borrow and save in a single nominally risk-free asset a_t at nominal rate r_t^n , supply labor at wage rate w_t , and receive dividend payouts q_t . Letting p_t^s and p_t^d denote the real relative prices of services and durables, δ the depreciation rate of durables, and π_t the inflation rate in the numeraire, we can write the household budget constraint as

$$p_t^s s_t + p_t^d \underbrace{[d_t - (1 - \delta)d_{t-1}]}_{\equiv e_t} + \psi(d_t, d_{t-1}) + a_t = w_t \ell_t + \frac{1 + r_{t-1}^n}{1 + \pi_t} a_{t-1} + q_t$$

We consider a general adjustment cost function in Section 2.5, but for now — again for expository purposes — restrict attention to a standard quadratic stock specification:

$$\psi(d, d_{-1}) = \frac{\kappa}{2} \left(\frac{d}{d_{-1}} - 1 \right)^2 d \quad (1)$$

For convenience we normalize steady-state total consumption expenditure $p^s \bar{s} + p^d \delta \bar{d}$ to one, and let the steady-state expenditure shares of services and durables be⁹

$$\phi \equiv p^s \bar{s}, \quad 1 - \phi \equiv p^d \delta \bar{d}$$

Finally, we assume that household labor supply is intermediated by standard sticky-wage unions (Erceg et al., 2000); we relegate details of the union problem to Appendix A.1.

⁷The expression for α is provided in Appendix A.1. The neutrality property for the aggregate demand shock b_t^a is desirable because it holds in the textbook New Keynesian model with only non-durables. Our definition of b_t^a is the natural extension of this notion of an “aggregate demand shock” to a multi-sector economy with durables; in particular, it is isomorphic to a shock to the shadow price of the total household consumption bundle, and thus equivalent to standard monetary shocks (Proposition 3). See Appendix B.2 for a discussion and an alternative specification of b_t^a . We thank Johannes Wieland for raising this point.

⁸See Appendix A.1 for the relationship between ϕ and $\tilde{\phi}$. We note that the scalings $\{\phi, 1 - \phi\}$ also imply that, in an economy with symmetric sectors (i.e., $\delta = 1$ and $\kappa = 0$) and flexible prices, sectoral shocks will only re-shuffle production across sectors, without any effect on aggregate output.

⁹The household preference parameter $\tilde{\phi}$ is then pinned down to make these expenditure shares consistent with optimal behavior (see Appendix A.1 for details).

PRODUCTION. Both services and durable goods are produced by aggregating intermediate varieties sold by monopolistically competitive retailers. Consistent with the empirically documented absence of meaningful short-run relative price movements (House & Shapiro, 2008; McKay & Wieland, 2020), we assume that the aggregated intermediate good can be flexibly transformed into durable goods or services, implying fixed real relative prices. In Section 2.5 we consider an extension of our model in which sector-specific supply shocks lead to changes in those real relative prices. Finally, production of intermediates only uses labor, and price-setting is subject to nominal rigidities. Since the problem of retailers is entirely standard we relegate details to Appendix A.1.

In equilibrium, aggregate output y_t must equal total consumption expenditures.¹⁰ In log-deviations from the steady state (denoted by $\hat{\cdot}$) aggregate output then satisfies

$$\hat{y}_t = \phi \hat{s}_t + (1 - \phi) \hat{e}_t$$

POLICY. The monetary authority sets the nominal rate of interest on bonds, r_t^n . For much of this section we will consider a *neutral* monetary rule that fixes the (expected) real rate of interest, i.e., $\hat{r}_t^n = \mathbb{E}_t [\hat{\pi}_{t+1}]$. Whenever we impose this assumption, we will ensure equilibrium determinacy by imposing that output ultimately reverts back to steady-state:¹¹

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 [\hat{y}_t] = 0 \tag{2}$$

In our quantitative explorations in Section 4.3 we instead consider a more standard (Taylor-type) rule of the form

$$\hat{r}_t^n = \phi_\pi \hat{\pi}_t \tag{3}$$

SHOCKS. The three disturbances b_t^a , b_t^s and b_t^d follow exogenous AR(1) processes with common persistence ρ_b and innovation volatilities $\{\sigma_b^a, \sigma_b^s, \sigma_b^d\}$, respectively. The common *exogenous* persistence ρ_b ensures that all differences in spending dynamics across sectors are related to differences in *endogenous* propagation.

¹⁰For simplicity, we assume that durables adjustment costs are either perceived utility costs, or get rebated back lump-sum to households.

¹¹This equilibrium selection can be formally justified with the continuity argument of Lubik & Schorfheide (2004): For a rule $\hat{r}_t^n = \phi_\pi \times \mathbb{E}_t [\hat{\pi}_{t+1}]$ and $\phi_\pi \rightarrow 1^+$, our equilibrium selection delivers continuity in ϕ_π . Note that a neutral monetary policy rule of this sort is consistent with any degree of price stickiness except for the limit case of *perfect* price flexibility.

2.2 The pent-up demand mechanism

We first of all use a stripped-down version of our model to cleanly illustrate the pent-up demand mechanism. We assume that: (i) all shocks are transitory ($\rho_b = 0$), (ii) there are no adjustment costs ($\kappa = 0$), and (iii) monetary policy is neutral.

We then arrive at the following characterization of aggregate impulse response functions.

Lemma 1. *The impulse responses of services and durables consumption expenditures to a vector of time-0 shocks $\{b_0^a, b_0^s, b_0^d\}$ satisfy*

$$\widehat{s}_0 = \frac{1}{\gamma}(b_0^a + b_0^s), \quad \widehat{s}_t = 0 \quad \forall t \geq 1 \quad (4)$$

and

$$\widehat{e}_0 = \frac{1}{\gamma}(b_0^a + b_0^d) \frac{1}{\delta} \frac{1}{1 - \beta(1 - \delta)}, \quad \widehat{e}_1 = -(1 - \delta)\widehat{e}_0, \quad \widehat{e}_t = 0 \quad \forall t \geq 2 \quad (5)$$

The impulse response of aggregate output is thus

$$\widehat{y}_0 = \phi\widehat{s}_0 + (1 - \phi)\widehat{e}_0, \quad \widehat{y}_1 = -(1 - \delta)(1 - \phi)\widehat{e}_0, \quad \widehat{y}_t = 0 \quad \forall t \geq 2 \quad (6)$$

Figure 1 shows impulse responses to three possible sets of time-0 shock vectors $\{b_0^a, b_0^s, b_0^d\}$, each normalized to depress aggregate output by one per cent on impact, but heterogeneous in their sectoral incidence. This exercise reveals how the *shape* of impulse responses — the focus of our paper — is affected by sectoral incidence, while keeping *amplification* — the focus of much previous work (e.g. Barsky et al., 2007) — constant. The figure thus answers the following question: conditional on a recession of a given severity, how do we expect the recovery dynamics to vary with the sectoral spending composition of the recession?

First, the solid green lines depict impulse responses to a pure durables demand shock ($b_0^d < 0$). Durables spending and so equilibrium output initially decline at $t = 0$. Following this contraction in durables spending, the household enters period $t = 1$ with a stock of durables below its steady-state target level. This results in *pent-up demand* for durables: the household seeks to replenish their stock and return it to target, and thus flow spending must overshoot. Since output is demand-determined in equilibrium, it also shows precisely this pent-up-demand-induced overshoot at $t = 1$ — a Z-shaped cycle.

Second, the solid orange lines depict impulse responses to a services demand shock ($b_0^s < 0$). Services consumption falls while durables consumption does not. As a result, there is no pent-up demand, equilibrium consumption and output return to steady state at $t = 1$, and




Figure 1: Dynamics in the stripped-down model. Impulse responses for: a pure durables shock (green), a pure services shock (orange), and an aggregate demand shock in an economy with a low services share $\underline{\phi}$ (dashed green) and a high services share $\bar{\phi}$ (dark blue). For details on the model parameterization see Appendix A.1.

the cycle is V-shaped.

Third, the dashed green and solid blue lines show impulse responses to an aggregate demand shock ($b_0^a < 0$) in two economies: one with a low steady-state share of services expenditures $\underline{\phi}$, and one with a high share $\bar{\phi}$. The larger the services share, the weaker pent-up demand effects, and so the less pronounced the Z-shape in aggregate output.

2.3 Implications for recovery dynamics

Having discussed the basic pent-up demand logic, we now discuss its implications for the relation between sectoral spending composition and recovery dynamics. We begin by defining two objects. First, we denote the share of services expenditures in time-0 aggregate consumption expenditure changes by ω :

$$\omega \equiv \frac{\phi \hat{s}_0}{\phi \hat{s}_0 + (1 - \phi) \hat{e}_0} \quad (7)$$

We will say that demand composition is more biased towards services when ω is larger. Second, we denote the cumulative impulse response of output, normalized by its trough change (which here is the time-0 change), by \hat{y} :

$$\hat{y} \equiv \frac{\sum_{t=0}^{\infty} \hat{y}_t}{\hat{y}_0} \quad (8)$$

The normalized cumulative impulse response (nCIR) measures the weakness of the reversal of output in the recovery phase. In particular, the nCIR is smaller when output reverts to steady state faster (or overshoots). Therefore, we will say that *a recovery is stronger* whenever \hat{y} is smaller.¹²

With the definitions (7) and (8) in hand, we can now state a first version of our main result on demand composition and the strength of recoveries.

Proposition 1. *Consider an arbitrary vector of time-0 shocks $\{b_0^a, b_0^s, b_0^d\}$ with a services share ω . Then, the normalized cumulative impulse response of aggregate output satisfies*

$$\hat{y} = 1 - (1 - \omega)(1 - \delta). \quad (9)$$

Thus, if $\delta < 1$, recoveries from recessions concentrated in services are weaker.

Proposition 1 states that, at least in the stripped-down model of Section 2.2, recoveries from demand-driven recessions will invariably be weaker if the composition of expenditure changes during the recession is more biased towards services. The logic follows immediately from Figure 1 and the discussion surrounding it: the larger the services share ω , the weaker pent-up demand effects, and so the weaker the subsequent recovery. For example, in a pure services recession ($\omega = 1$) — as in the orange line in Figure 1 — the lost output during the recession is entirely foregone ($\hat{y} = 1$), whereas in a pure durables recession ($\omega = 0$) — as in the green line in Figure 1 — the recovery is boosted by pent-up demand and only a fraction δ of lost output is actually foregone ($\hat{y} = \delta$). Note that this result applies irrespective of *why* a given recession’s services share ω is high or low. For example, it applies for variation in ω generated by both (i) changes in long-run expenditure shares ϕ , fixing the shocks $\{b_0^a, b_0^s, b_0^d\}$, and (ii) changes in shocks $\{b_0^a, b_0^s, b_0^d\}$, fixing the long-run expenditure shares ϕ .

The objective for the remainder of this paper is to learn about this effect of demand composition on the strength of recoveries — i.e., $\frac{\partial \hat{y}}{\partial \omega}$ — in environments that are substantially richer than the one considered so far.

2.4 Recovery dynamics in the full model

We return to the full model of Section 2.1. While the pent-up demand mechanism is always present, we will now show that our conclusions on recovery strength in this extended setting

¹²An alternative but related measure of persistence is the half-life of output. However, since output dynamics may be non-monotone, the half-life is generally a less appropriate measure of persistence and recovery strength than the nCIR.

do not go through automatically — pent-up demand effects need to be “strong enough.” The key contribution of this section is to lay the groundwork for an empirical test: we will give an empirically measurable, theoretically necessary and sufficient condition for pent-up demand effects to indeed be consistent with our main prediction on recovery strength. Throughout, this section maintains the assumption of a neutral monetary policy rule that fixes the (expected) real rate of interest.

Proposition 1 holds in the simple model because following contractionary aggregate demand shocks — and so any mix of aggregate and sectoral demand shocks — the impulse response of durables spending relative to trough always lies above the corresponding normalized services impulse response, as illustrated by Figure 1. This result does not necessarily extend to the full model; in particular, adjustment costs on the durable stock can moderate the rebound in durables spending and so push the durables impulse response below that of services.¹³ Thus, if adjustment costs dominate pent-up demand effects, then recoveries may instead be *stronger* after services-led recessions.

Our key insight is that, even in the full model, the relative ranking of impulse responses on the one hand and the effect of demand composition on recovery dynamics on the other are always two sides of the same coin. Proposition 2 formalizes this claim.

Proposition 2. *Suppose that monetary policy is neutral. Let $\widehat{\mathbf{s}}^a$ and $\widehat{\mathbf{e}}^a$ denote the normalized cumulative impulse responses of services and durables expenditure to a recessionary aggregate demand shock $b_0^a < 0$, defined as in (8). Then, the normalized cumulative impulse response of aggregate output $\widehat{\mathbf{y}}$ is increasing in the services share ω if and only if*

$$\widehat{\mathbf{s}}^a > \widehat{\mathbf{e}}^a \tag{10}$$

Proposition 2 states that information on sectoral nCIRs following an aggregate demand shock b_0^a is enough to sign the effect of sectoral demand composition on the strength of recoveries. Loosely speaking, the condition (10) guarantees that pent-up demand effects are “strong enough,” and in particular dominate the slowdown in durables spending reversal associated with strong adjustment costs.

In our view, the main advantage of stating the condition in terms of the ranking of sectoral impulse responses — two endogenous objects — as opposed to primitive, structural parameters is that these impulse responses can in principle be directly measured in the data. A second advantage is that (10) continues to apply even in environments substantially

¹³In Appendix B.1 we use the full model’s closed-form solution to make this point formally.

richer than the one considered so far. We will first of all briefly discuss these extensions in Section 2.5, and then turn to empirical measurement of (10) in Section 3.

2.5 Further generalizations

We conclude our theoretical analysis by investigating the robustness of Proposition 2 to several model generalizations. We begin with the case of an active monetary rule (together with partially sticky prices), and then consider further departures from our baseline model of Section 2.1. We summarize the main conclusions here, and relegate formal statements to Appendices A.2, B.1 and B.2.

ACTIVE MONETARY RULES AND PARTIALLY STICKY PRICES. In Appendix B.1, we relax the assumption of a neutral monetary rule, and provide a generalized version of (10).

With an active monetary rule and partially sticky prices, we find that the condition (10) is generally only necessary, not sufficient. Inflation overshoots by more in the durables-led recovery, so the monetary authority increases nominal (and so real) rates by more, thereby moderating the subsequent overshoot in aggregate output. This limitation of our analytical results motivates the model simulations in Section 4.3 and Figure B.3. Those simulations show that — when structural models with active monetary rules are calibrated to satisfy (10) — it remains robustly true that \hat{y} is increasing in ω .

MANY SECTORS. Consider an extension of the baseline model with N sectors, with each good heterogeneous in its depreciation rate δ_i , adjustment cost parameter κ_i , and output share ϕ_i . Following the same steps as in the proof of Proposition 2, we can show that the output nCIR \hat{y} for an arbitrary shock mix $\{b_0^a, \{b_0^i\}_{i=1}^N\}$ that results in shares $\{\omega_i = \frac{\phi_i \hat{e}_0^i}{\hat{y}_0^a}\}_{i=1}^N$ is given by

$$\hat{y} = \sum_{i=1}^N \omega_i \times \hat{e}_i^a \quad (11)$$

Equation (11) is a natural extension of our previous two-sector results: recoveries are stronger whenever the recession spending shares ω_i are larger for sectors with a smaller spending nCIR to an aggregate demand shock \hat{e}_i^a .

The multi-sector perspective (11) also makes particularly clear that, at least conceptually, our statements about recovery dynamics are not at all tied to differences in durability: heterogeneity in spending nCIRs across sectors *for whatever reason* will translate into heterogeneous recovery dynamics after sectorally biased recessions. We simply choose to focus

our analysis on durability and pent-up demand because it is one natural — and, as we will see, empirically supported — candidate for nCIR variation across sectors.¹⁴

GENERAL ADJUSTMENT COSTS. Our baseline model considered a very particular form of adjustment costs. Consider instead a general adjustment cost function of the form

$$\psi(\{d_{t-\ell}\}_{\ell=0}^{\infty}) \tag{12}$$

The function (12) is general enough to nest arbitrary forms of adjustment costs, including in particular adjustment costs on expenditure flows (rather than stocks). Given this, we lose the ability to characterize impulse response functions in closed form. Nevertheless, as long as monetary policy is neutral, it is still true that

$$\hat{\mathbf{y}} = \omega \hat{\mathbf{s}}^a + (1 - \omega) \hat{\mathbf{e}}^a, \tag{13}$$

for any vector of shocks $\{b_0^a, b_0^s, b_0^d\}$ resulting in a services share ω .¹⁵ Thus (10) still applies. Intuitively, the crucial restriction is that the system of first-order conditions for consumer demand remains separable in services s_t and durables d_t .

NON-SEPARABILITY. Consider an alternative preference specification where services and durables are aggregated in a CES fashion with an inverse elasticity of substitution of ζ .¹⁶ We show that (10) remains sufficient in the empirically relevant case of net substitutability ($\zeta < \gamma$). The condition is sufficient but not necessary because pent-up demand effects are now in fact reinforced by substitution patterns: a low durables stock pushes up services demand, and so further increases the strength of the recovery.

STICKY INFORMATION. Recent work on durables spending dynamics has highlighted the importance of informational frictions on the household side (McKay & Wieland, 2020), notably in the form of sticky information constraints (Mankiw & Reis, 2003). Such frictions

¹⁴By the same logic, investment expenditures could also be subject to pent-up demand effects. Empirically, however, the reversal in investment does not seem to be as strong as for consumer durables spending, suggesting that pent-up demand effects for investment might be more muted.

¹⁵Note that, with arbitrary adjustment costs, the expenditure troughs for different sectors may in principle occur at different time periods. We thus construct $\hat{\mathbf{s}}^a$ and $\hat{\mathbf{e}}^a$ by fixing a common normalization period t^* (e.g., the trough of $\hat{\mathbf{y}}^a$), and then note that (13) applies for *any* choice of t^* .

¹⁶Our baseline separable preferences are a special case with $\zeta = \gamma$.

again generally rule out simple characterizations of impulse response functions, but do not affect the system's separability in s_t and d_t , so (10) yet again applies.

We already note at this point that (i) more general forms of adjustment costs, as in (12), and (ii) the sticky information friction will play an important role in our quantitative impulse response matching exercise in Section 4.3. Importantly, by the discussion provided here, these two frictions are entirely consistent with our main empirical test based on Proposition 2.

INCOMPLETE MARKETS. Proposition B.1 and Proposition 2 continue to apply without change in a model extension with liquidity-constrained households. Formally, we consider an extension of the baseline framework of Section 2.1 in which a fraction μ of households cannot save or borrow in liquid bonds, and so is hand-to-mouth in each period. In this environment, depending on the cyclical nature of income for hand-to-mouth households, the impulse responses in Lemma B.1 are scaled up or down. Impulse response *shapes*, however, are unaffected by this scaling, and so our conclusions on recovery dynamics are entirely unaffected.

SUPPLY SHOCKS. As our final extension, we allow for the production of durables and services out of the common intermediate good to be subject to productivity shocks. By perfect competition in final goods aggregation, it follows that these productivity shocks transmit directly into real relative prices. Thus, at least in our baseline case of a neutral monetary policy rule, supply shocks are isomorphic to our demand shocks (which are effectively shocks to shadow prices), and so all results extend without any change.¹⁷ This extension motivates our alternative empirical test using oil shocks, discussed in Section 3.2.

3 Empirical Evidence

The previous sections gave conditions under which recoveries from demand-driven recessions concentrated in services or non-durables will tend to be weaker than recoveries from recessions biased towards more durable goods. In particular, we have seen that, in a family of relatively standard business-cycle models, the prediction holds if and only if durable expenditures exhibit a stronger reversal than services (and non-durables) expenditures *conditional on an aggregate demand shock*. In this section, we begin by testing this condition and, there-

¹⁷Of course, by the economy's production technology, supply and demand shocks necessarily have different effects on hours worked. With a fixed real rate of interest, however, these differences in hours worked do not affect any other equilibrium aggregates.

fore, the pent-up demand mechanism. In particular, Section 3.1 studies sectoral expenditure dynamics conditional on standard monetary policy shocks, and Section 3.2 extends the analysis for several other shocks. Finally, in Section 3.3, we complement this indirect test of the mechanism with suggestive *direct* evidence in support of its main prediction, exploiting variation across U.S. regional business cycles during the Great Recession and COVID-19.

3.1 Monetary policy shocks

As the main empirical test of the pent-up demand mechanism, we study the response of different consumption categories to identified monetary policy shocks.

We focus on monetary shocks for two reasons. First, among all of the macroeconomic shocks studied in applied work, monetary shocks are arguably the most prominent, and much previous work is in agreement on their effects on the macro-economy (Ramey, 2016; Wolf, 2020). Our contribution thus need not lie in shock identification; instead, we can focus on the impulse responses themselves and their connections to our theory. Second, when viewed through the lens of the model in Section 2.1, monetary shocks are equivalent to our notion of an aggregate demand shock b_t^a , and so directly map into the empirical test of Proposition 2. To establish this claim, we extend the model to allow for AR(1) shocks m_t to the monetary rule. We then arrive at the following equivalence result.¹⁸

Proposition 3. *Consider the model of Section 2.1, extended to feature exogenous AR(1) deviations m_t from the central bank’s rule. The impulse responses of $x \in \{s, e, d, y\}$ to (i) to a recessionary aggregate demand shock $b_0^a < 0$ with persistence ρ_b , and (ii) a contractionary monetary shock $m_0 = -(1 - \rho_b)b_0^a$ with persistence $\rho_m = \rho_b$ are identical:*

$$\hat{x}_t^a = \hat{x}_t^m$$

Intuitively, equivalence obtains because both our aggregate demand shock as well as monetary shocks move the shadow price of the total household consumption bundle. We can thus test the key condition (10) using sectoral impulse responses to monetary policy shocks.

EMPIRICAL FRAMEWORK. Our analysis of monetary policy transmission follows Christiano et al. (1999). We estimate a Vector Autoregression (VAR) in measures of consumption, output, prices and the federal funds rate, and identify monetary policy shocks as the innovation

¹⁸Note that the proof covers both a simple Taylor rule as in (3) as well as our baseline case of a neutral monetary policy rule.




Figure 2: Quarterly impulse responses to a recursively identified monetary policy shock (as in Christiano et al. (1999)) by consumption spending category, all normalized to drop by -1% at the trough. The solid blue line is the posterior mean, while the shaded areas indicate 16th and 84th percentiles of the posterior distribution, respectively.

to the federal funds rate under a recursive ordering, with the policy rate ordered last. We consider quarterly data, with the sample period ranging from 1960:Q1 to 2007:Q4. To keep the dimensionality of the system manageable, we fix aggregate consumption, output, prices and the policy rate as a common set of observables, and then estimate three separate VARs for three broad categories of household spending — durables, non-durables, and services.¹⁹ We include four lags, and estimate the models using Bayesian techniques. Further details of the set-up as well as results for finer sectoral expenditure series are provided in Appendix C.1.

RESULTS. Consistent with previous work, we find that a contractionary monetary policy shock lowers output and consumption. Figure 2 then decomposes the response of aggregate consumption into its three components: durables, non-durables, and services. We are mostly interested in the comparison of services and durables spending impulse responses; however, since non-durables as measured by the BEA also contain semi-durables, a comparison with the non-durables spending impulse response provides a useful additional test. We find that the reversal in durables spending is much weaker than that of services spending, with the former in fact overshooting.

¹⁹Note that the baseline NIPA measure of total services contains housing and utilities. Imputed rent arguably has a durable component, so we show in Appendix C.1 that our results change little if housing and utility services are stripped out.

To formally test the condition in Proposition 2, we compute the posterior distribution of

$$\frac{\widehat{\mathbf{s}}^a}{\widehat{\mathbf{e}}^a} - 1 \quad (14)$$

At the posterior mode, the services nCIR is 88 per cent larger than the durables nCIR, with this difference statistically significant at conventional levels.²⁰ The non-durables spending nCIR is between the two, at around 22 per cent larger than the durables nCIR.

We emphasize that our results very much agree with previous work; notably, both Erceg & Levin (2006) and McKay & Wieland (2020) report an overshoot in durables expenditure following monetary shocks. Relative to those papers, our incremental contribution is to (i) look in more detail across consumption categories and (ii) show that those results constitute a highly informative test on the effects of demand composition on recovery dynamics. We can overall conclude that both our as well as previous empirical evidence are very much consistent with the key testable condition in Proposition 2.

FINER SECTORAL SPENDING DYNAMICS. As already emphasized in Section 2.5, the logic of Proposition 2 is not at all tied to whether a given type of spending is “officially” categorized as being a durable or a service: recoveries are stronger after recessions biased towards sectors with comparatively small demand shock nCIRs, irrespective of the underlying mechanism. Viewed in this light, Figure 2 merely confirms that, as predicted by standard consumer theory, coarse measures of durability are a powerful predictor of nCIR differences. As a final step, we now repeat our analysis at finer sectoral levels, asking whether the responses in Figure 2 actually mask meaningful within-category nCIR heterogeneity.

Table 3.1 summarizes our results by reporting a large number of sectoral monetary policy nCIRs.²¹ We emphasize two key take-aways. First, going more granular does not materially alter our main conclusions: durables and (to a lesser extent) non-durables nCIRs are quite consistently smaller than those of services, across sectors. The variation within coarse

²⁰In computing the nCIRs, we truncate at a maximal horizon $T^* = 20$, consistent with our focus on short-run business-cycle fluctuations. Our results are even stronger for longer horizons. To construct the posterior credible set, we estimate a single VAR containing all consumption measures, compute the nCIR ratio for each draw from the posterior, and then report percentiles.

²¹We include all the major sub-categories of durables, non-durables and services (excluding housing and utilities, see Footnote 19) as reported in the NIPA tables. The slight quantitative differences relative to the discussion surrounding (14) are related to the fact that our construction of posterior credible intervals for the coarse spending series relies on a single VAR with all three sectoral consumption series, while for Table 3.1 we estimate separate VARs for each consumption series.

Durables		Non-Durables		Services	
<i>All</i>	1.00	<i>All</i>	1.28	<i>All</i>	1.97
<i>Motor Vehicles</i>	1.14	<i>Food</i>	1.01	<i>Health</i>	1.98
<i>Furniture</i>	1.31	<i>Clothes</i>	0.97	<i>Transport</i>	1.65
<i>Recreation Goods</i>	0.86	<i>Gas</i>	1.53	<i>Recreation</i>	1.43
<i>Other</i>	1.19	<i>Other</i>	0.76	<i>Food</i>	2.19
				<i>Financial</i>	1.24
				<i>Other</i>	1.55

Table 3.1: Monetary policy nCIRs for sectoral sub-categories, with all nCIRs are expressed relative to the durables nCIR. For further details see Appendix C.1.

sectoral categories is intuitive; for example, clothes (a “semi-durable”) have a small nCIR. Second, ex-ante some services may be expected to behave more like durables — for example, health expenditures, recreation services, and food away from home may have a “memory” component (Hai et al., 2013). However, we do not observe much evidence of strong pent-up demand effects for such services, at least in response to past U.S. monetary shocks.²²

3.2 Other shocks

While impulse responses to monetary policy innovations are, for the reasons discussed in Section 3.1, a close-to-ideal test of the pent-up demand mechanism, they are of course not the only possible one. In this section we summarize the results of several other empirical exercises, with details for all relegated to Appendices C.2 to C.4.

UNCERTAINTY. Uncertainty shocks are a natural candidate for the reduced-form aggregate demand shocks b_t^a , and as such a promising alternative to the baseline monetary policy experiment. Following Basu & Bundick (2017), we identify uncertainty shocks as an innovation in the VXO, a well-known measure of aggregate uncertainty.

Our results are very similar to the monetary policy experiment. All components of consumption drop on impact, but durables expenditure recovers quickly and then overshoots, while the recoveries in non-durable and in particular service expenditure are more sluggish.

²²This is actually consistent with the findings in Hai et al. (2013). They estimate that a typical memory good stock depreciates by about two-thirds within a year — a much higher rate of depreciation than for the average durable good.

However, given the relatively short sample, our estimates are somewhat less precise than for our main experiment on monetary policy shock transmission.

OIL. As a third test, we study oil price shocks, identified as in Hamilton (2003) and then ordered first in a recursive VAR. While such shocks can generate broad-based recessions, they are special in that they directly affect the *relative prices* of consumption goods. As discussed in Section 2.5, such relative supply shocks will generate pent-up demand effects exactly like the demand shocks presented in Section 2.1. In particular, a sudden increase in oil prices will increase the effective relative price of all transport-related consumption, allowing us to test the ranking of nCIRs at a finer sectoral level, as in (11).

Again, the results support the pent-up demand mechanism. Since transport-related expenditures are an important component of durables expenditure (e.g., motor parts and vehicles), total durable consumption is strongly affected by the shock and follows the predicted Z-shaped pattern. Food and finance expenditures instead dip in the initial recession, but then simply return to baseline, without any further overshoot.

REDUCED-FORM FORECASTS. So far, we have focussed on dynamics *conditional* on particular structural shocks, thus allowing us to directly connect empirics and the theory in Proposition 2. As our final check, we here instead look at *unconditional* sectoral expenditure dynamics. Implicitly, in looking at such reduced-form forecasts, we are assuming that sectoral dynamics are largely driven by aggregate demand shocks; in that case, it is straightforward to see that unconditional forecasts can also be used for the test in (10).

To implement the forecasting exercise, we estimate a high-order reduced-form VAR representation in our three baseline sectoral output categories, and then separately trace out the implied impulse response of aggregate consumption to reduced-form innovations in each equation, with the impact impulse response in all cases normalized to one. Consistent with both theory and our previous empirical results, we find that innovations to the durables expenditure equation move aggregate consumption much less persistently than equally large innovations to non-durables and services expenditures. In particular, we find that the total consumption nCIR for an innovation to services spending is around 120 per cent larger than the nCIR corresponding to a durables innovation. These unconditional results are at least qualitatively consistent with the *conditional* results for monetary policy shocks.

3.3 Evidence from U.S. regional business cycles

The analysis in the previous subsections provided a test of the pent-up demand mechanism and so (indirectly) of our prediction regarding the effects of demand composition on recovery strength. In this section we provide some suggestive *direct* evidence on this prediction.

THE IDENTIFICATION CHALLENGE. Anecdotally, evidence from U.S. business cycles on the relation between spending composition and recovery strength is mixed. For example, as the economy emerged from the 1945 recession, the recovery was buoyed by consumer purchases of previously rationed durable goods (e.g., cars, furniture, and appliances). Similarly, the 1960 and 1973 recessions were associated with strong recoveries and a particularly large share of durables in overall expenditure declines.²³ On the other hand, the COVID-19 recession-recovery episode appears to contradict our main prediction: despite the unusually large share of services in expenditure declines, the economy has bounced-back relatively quickly when compared, for instance, to the recovery from the Great Recession.

Such comparisons across historical U.S. episodes, however, are unlikely to allow a credible test of the effects of demand composition on the strength of recoveries. There are two reasons; one statistical, the other conceptual. Statistically, the relatively small number of observations — twelve recessions in the U.S. since 1948 — precludes sharp inference. Conceptually, and more importantly, several other likely determinants of a recovery’s strength — including in particular the persistence of the underlying shock and the response of policy — are likely to have differed across episodes. For example, the ongoing recovery from the COVID-19 recession may well have been comparatively fast only because the shock was rather transitory (given the quick discovery of vaccines) and the policy response was unusually strong.

REGIONAL EVIDENCE. To address these two challenges, we propose to use variation across U.S. *regional* business cycles during the COVID-19 and Great Recession episodes. As we show in Appendix B.3, our results on spending composition and recovery strength in Proposition 1 extend straightforwardly to differences between regions in a model of a monetary union where regions produce tradable durable goods and non-tradable services. Furthermore, and most importantly, cross-regional variation is a natural candidate to address our two challenges: statistically, our tests can now leverage substantial cross-regional variation;

²³In the average recession, durable spending accounts for around 65 per cent of total spending contractions (see Figure C.7). In contrast, between 1960M04-1961M02 and 1973M11-1975M03, durable spending accounted for 105 per cent and 165 per cent of the total spending declines, respectively.

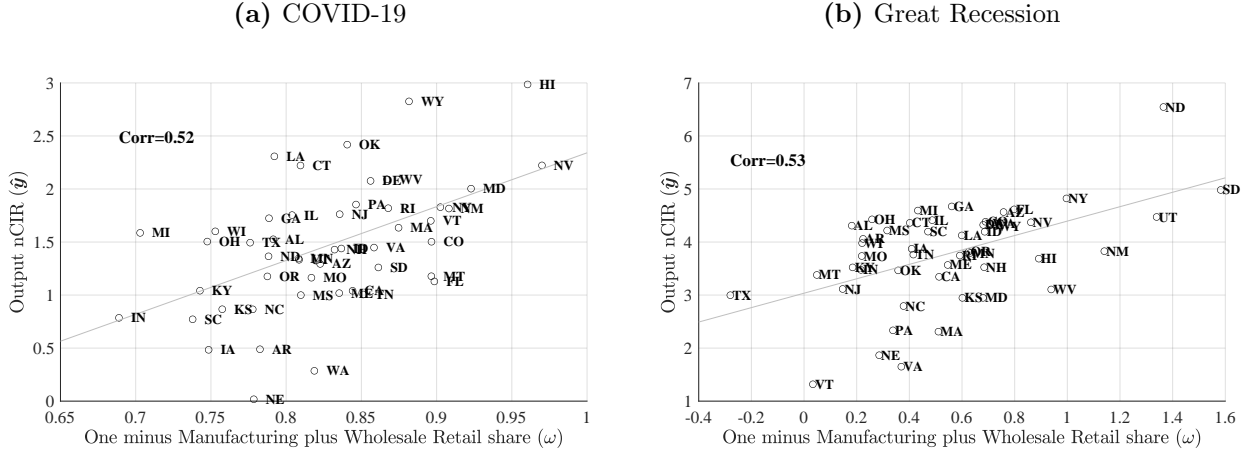


Figure 3: Recovery strength and recession sectoral composition. The nCIR of output (\hat{y}) is the cumulative percent change in private industry real GDP from peak to four quarters after the trough, normalized by the peak-to-trough change. The manufacturing plus wholesale retail share ($1-\omega$) is the ratio of the percent change in manufacturing plus wholesale retail output over the percent change in private industry output from peak to trough. Data is from the BEA, Table SQGDP11.

and conceptually, looking across regions in a given recession allows us to control for the — at least largely common — aggregate shock persistence and policy response.

Our empirical analysis proceeds as follows. In our extension of Proposition 1 to cross-regional differences (see Appendix B.3), the share ω is interpreted as the services plus non-durables share in local *output* (rather than consumption). Consistent with this, we use state-level, sectoral real GDP series from the BEA (quarterly frequency). For each episode, we identify the state-specific peak and trough of private industry output. We then compute the services plus non-durables share of the recession — i.e., the theoretical ω — as one minus the ratio of the percent change in manufacturing plus wholesale retail output over the percent change in private industry output from peak to trough. From here, we compute the nCIR of output — the theoretical \hat{y} — as the cumulative percent change in private industry output from peak to four quarters after the trough, normalized by the peak-to-trough change.²⁴ Appendix C.5 provides more details on the construction of all variables, and repeats the analysis for other plausible definitions of the share ω .

Results are reported in Figure 3, which plots (ω, \hat{y}) pairs for the COVID-19 (left panel) and Great Recession (right panel) episodes. We see that sectoral recession composition is a strong predictor of recovery strength across U.S. states, for both recessions. On average, the

²⁴We choose four quarters because this is the longest horizon for which data is available after the second quarter of 2020 (the trough in most states during the COVID-19 recession).

recovery was stronger (lower \hat{y}) in states where manufacturing and wholesale retail accounted for more of the output decline of the recession (lower ω).²⁵ Our analysis of cross-regional data is central to this conclusion: looking across time, we see that average state-level output nCIR \hat{y} and share ω are instead *negatively* related — i.e., lower average \hat{y} and higher ω in the COVID-19 recession relative to the Great Recession — reflecting the confounding factors of differential shock persistence and policy response. Further supplementary evidence that looks at a longer history of regional business cycles is provided in Appendix C.5.

DISCUSSION. Our analysis in this subsection has demonstrated that past regional business cycles in the U.S. have been broadly consistent with our prediction on sectoral composition and recovery strength. Of course, the documented reduced-form patterns are unlikely to be driven *exclusively* by the pent-up demand dynamics that we emphasize; for example, in the aftermath of the COVID-19 recession, stimulus checks together with substitution towards durables are very likely to have also materially contributed to the particularly fast rebound of durables spending. We thus in the next section proceed to *quantify* the importance of the pent-up demand mechanism using explicitly (semi-)structural methods.

4 Quantifying the effects of demand composition

Are pent-up demand effects likely to be strong enough to drive meaningful variation in recovery strength across recessions with different sectoral composition? Section 4.1 begins by asking why and how much spending composition varies across typical recession episodes. We then consider two approaches to translating this observed variation into predicted differences in recovery strength.

First, in Section 4.2, we show that — even in relatively general variants of our baseline model from Section 2.1 — the desired counterfactual can be consistently estimated through a simple shift-share design that leverages estimated impulse responses to an aggregate demand shock. We implement this semi-structural approach with our monetary policy shock estimates from Section 3.1. Second, in Section 4.3, we follow a fully structural approach: we consider a model variant that violates the conditions required for the shift-share, discipline it through impulse response matching, and then use it to recover the desired counterfactual.

²⁵Note that this is not mechanically driven by such states experiencing milder recessions, as the nCIR is normalized by the peak-to-trough change in output, precisely to account for differences in the magnitude of recessions.

4.1 Sources of variation in demand composition

We want to investigate whether empirically plausible variations in recession spending composition — the ω 's in Section 2 — can have meaningful effects on recovery dynamics. As a first step, we thus need to get a sense of what constitutes plausible variation in ω . Figure 3 gives an indication for *regional* business cycles. To speak to aggregate cyclical fluctuations, we will look at the following two counterfactual exercises.

1. In the first exercise we ask: fixing an aggregate demand shock b_t^a , how different would the recovery dynamics be in economies with different long-run expenditure shares of durables, nondurables, and services? This exercise is motivated by the observed cross-country differences in expenditure shares, illustrated in the left panel of Figure C.7 in Appendix C.6. For example, the 2017 U.S. durables and services shares are 0.09 and 0.68, respectively; for comparison, the durables share can be as large as 0.13 in an economy like Canada and the services share can be as low as 0.37 in an economy like Russia.
2. In the second exercise we ask: fixing a given parameterization of the economy, how different would the recovery dynamics be following different combinations of shocks $\{b_t^a, b_t^s, b_t^d\}$ that induced different recession spending compositions? This exercise is motivated by the wide range of sectoral patterns observed in past U.S. recessions, as illustrated in the right panel of Figure C.7 in Appendix C.6. As is well known, a typical U.S. recession tends to be more biased towards durables: the contribution of durable expenditure changes to real personal consumption expenditure changes is 62 percent (average of peak-to-trough changes from 1960 to 2019). At one extreme of the spectrum was the recession following the 1973 oil crisis: as gas prices increased, consumers cut car purchases much more than in a typical recession, and so durables spending overall accounted for far more than 100 percent of the total decline. At the opposite extreme, the COVID-19 pandemic triggered a recession in which services spending cuts accounted for almost all of the total decline — fearing infection, consumers mostly cut down on food away-from-home as well as travel- and health-related services.

The remainder of this section will translate the range of observed sectoral shares into implied variation in recovery strength, relying first on a semi-structural shift-share (in Section 4.2) and then on a structural model (in Section 4.3). We emphasize that these counterfactuals are explicitly designed to vary demand composition *alone*, keeping everything else fixed. Our counterfactuals should thus *not* be understood as forecasting recoveries in any

given country or from any particular U.S. recession — for example, those could also differ in the persistence of the underlying shocks or in the responses of monetary and fiscal policy.²⁶

4.2 Semi-structural shift-share

In Section 3.1, we estimated the impulse responses of all components of consumer expenditures to a change in the monetary policy stance and so, under the conditions of Proposition 3, to an aggregate demand shock b_t^a . We now propose a shift-share approach that arrives at our desired counterfactuals simply by suitably re-weighting those sectoral impulse responses.

Our key result supporting this approach is Proposition 4.

Proposition 4. *Consider the model of Section 2.1, and suppose that monetary policy is neutral up to shocks m_t . Now let \widehat{s}_t^m and \widehat{e}_t^m denote the impulse responses of services and durables expenditures, respectively, to a monetary policy shock. Then:*

1. *In an alternative economy with services share ϕ' , the impulse response of aggregate output to an aggregate demand shock b_0^a with persistence $\rho_b = \rho_m$ and $\widehat{y}_0 = -1$ is*

$$\widehat{y}_t = - \left[\frac{\phi'}{\phi' \widehat{s}_0^m + (1 - \phi') \widehat{e}_0^m} \widehat{s}_t^m + \frac{1 - \phi'}{\phi' \widehat{s}_0^m + (1 - \phi') \widehat{e}_0^m} \widehat{e}_t^m \right]$$

2. *The impulse response of aggregate output to an arbitrary combination of aggregate and sectoral demand shocks $\{b_0^a, b_0^s, b_0^d\}$ with persistence ρ_m and such that $\{\widehat{y}_0 = -1, \phi \widehat{s}_0 = -\omega, (1 - \phi) \widehat{e}_0 = -(1 - \omega)\}$ is*

$$\widehat{y}_t = - \left[\omega \frac{\widehat{s}_t^m}{\widehat{s}_0^m} + (1 - \omega) \frac{\widehat{e}_t^m}{\widehat{e}_0^m} \right]$$

Proposition 4 gives conditions under which our two desired counterfactuals — recovery dynamics in economies (i) with different long-run expenditure shares and (ii) subject to different shock combinations — can be estimated semi-structurally, simply by re-weighting the estimated sectoral impulse responses to monetary shocks.

RESULTS. We now leverage Proposition 4 together with the sectoral monetary policy impulse responses from Figure 2 to construct our two counterfactuals. Given that our estimated

²⁶The COVID-19 recession is a particularly instructive example: there, the shock is likely to be relatively transitory, and the policy response was particularly strong.

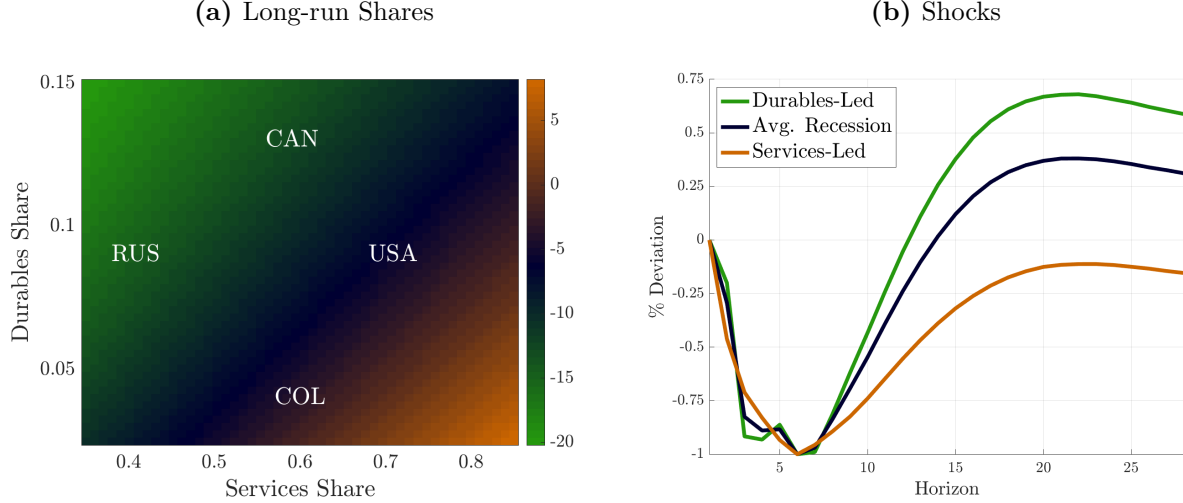


Figure 4: Left panel: nCIR to an aggregate demand shock b_0^a as a function of long-run expenditure shares and relative to the U.S. Right panel: Impulse response of total consumption to shock combinations giving a durables/non-durables/services trough spending composition of: (i) [100%, 50%, -50%] for a durables-led recession, (ii) [62%, 45%, -7%] for the average U.S. recession, and (iii) [12.5%, 12.5%, 75%] for a services-led recession. All are normalized to lead to a peak-to-trough consumption contraction of -1 per cent. For both panels we use the posterior mode point estimates from Figure 2.

impulse responses are hump-shaped (and in particular mechanically equal to zero on impact), we compute nCIRs by normalizing the *trough* of the impulse response (rather than the impact response, as in (8) and Proposition 4).²⁷

We begin with our first counterfactual: nCIRs for an aggregate demand shock b_0^a as a function of the long-run expenditure shares. Results are reported in the left panel of Figure 4. In the figure, we have scaled the nCIR of an economy with the sectoral composition of the U.S. to one. The color shadings reveal that, as sectoral shares are adjusted, the strength of recoveries as measured by the nCIR changes substantially. On the one hand, in an economy as durables-intensive as Canada or with a services share as low as in Russia, the nCIR is around 15 per cent smaller; on the other hand, for economies with a durables share as low as that in Colombia, the nCIR can be around 5 per cent larger.

For our second counterfactual, we compare impulse response paths for a vector of sectoral demand shocks with a peak effect on consumption of -1 per cent and sectoral composition of

²⁷This normalization is consistent with our discussion in Footnote 15. Also note that, since our empirical estimates in Section 3.1 are split spending into three categories, the quantitative results below rely on the natural three-sector extension of Proposition 4, derived easily from our general multi-sector characterizations in Section 2.5.

expenditure changes from peak-to-trough as in (i) a strongly durables-led recession, (ii) an average U.S. recession, and (iii) a strongly services-led recession. We choose the compositions corresponding to the strongly durables-led and services-led recessions to approximately capture the variation observed across past U.S. recessions (e.g., from the 1973 oil crisis to the COVID-19 recessions). Results are reported in the right panel of Figure 4. As expected, the durables-led contraction shows a strong reversal and overshoot, while the recovery from an ordinary recession is more gradual, and the recovery from a heavily services-biased recession is even weaker. In nCIR terms, the effects are large; for example, at the point estimates displayed in Figure 4, the nCIR of output in a heavily services-biased recession is 67.8 per cent larger compared to an average recession, with the difference strongly statistically significant.

DISCUSSION. In addition to being simple and transparent to implement, the semi-structural shift-share approach is attractive because it is not tied to any particular parametric model. In particular, it frees the researcher from needing to take a stance on the nature of adjustment costs and informational frictions — two structural model features that matter for the importance of the pent-up demand mechanism (see McKay & Wieland (2020) and our results in Section 4.3), yet are not easy to pin down credibly.²⁸

Its key drawback is that it relies on some restrictive assumptions, notably that of a neutral monetary rule, and applies only to demand shocks that are exactly as persistent as the identified monetary shock.

4.3 Structural counterfactuals

We now instead compute our two counterfactuals in explicit structural models, estimated through impulse response matching (as in Christiano et al., 2005). We then rely on numerical simulations from the model, both at our preferred parameter point estimates as well as across a wider range of plausible parameterizations.

MODEL. Our baseline model in Section 2.1 was designed to illustrate the pent-up demand mechanism as transparently as possible. Unfortunately, for the purposes of impulse response matching, this baseline model is not quite rich enough.

²⁸Even though Proposition 4 is stated in the context of the model of Section 2.1, it is straightforward to show that the result applies unchanged in our extended economies with incomplete markets, arbitrary adjustment costs, and informational frictions, and extends naturally to an N-sector environment.

As it turns out, however, two simple changes — both already discussed in Section 2.5 — are enough to get the desired agreement between data and model: (i) quadratic adjustment costs on the durables spending *flow* instead of the stock, similar to the treatment of investment in many quantitative business-cycle models (Smets & Wouters, 2007; Justiniano et al., 2010); and (ii) households that only learn gradually about the shocks hitting the economy, as in the sticky information model of Mankiw & Reis (2002) (and as applied to the durable consumption setting by McKay & Wieland (2020)). Details for both extensions are provided in Appendix A.2. Table 4.1 reports our preferred parameterization of this extended model.

CALIBRATED PARAMETERS. We first of all calibrate a subset of parameters. The two preference parameters (β, γ) are standard.²⁹ We consider a broad notion of durables, and thus set the depreciation rate δ as quarterly durable depreciation divided by the total durable stock in the BEA Fixed Asset tables. Given δ , we set the preference share ϕ to fix durables expenditure as 10 per cent of total steady-state consumption expenditure. We set wages to be moderately flexible, roughly consistent with the estimates in Beraja et al. (2019) and Grigsby et al. (2019), and consider a relatively flat New Keynesian Phillips Curve, as in Ajello et al. (2020) or Hazell et al. (2020). For monetary policy, we consider the Taylor rule (3). Finally, as our baseline, we set the persistence of demand shocks to that of our empirically identified monetary policy shock, recovered from the estimated impulse response of nominal interest rates to the shock.

ESTIMATED PARAMETERS. The remaining parameters are chosen to ensure a tight match between empirically estimated and model-implied impulse responses, as displayed in Figure 5. We consider a model without any level adjustment costs ($\kappa = 0$). The estimation then calls for moderate adjustment costs on expenditures (κ_e), and strong informational frictions ($1 - \theta$ is the fraction of households receiving the latest information each period).

The implied impulse response fit is excellent: both sets of impulse responses follow hump shapes, and expenditure on durables shows a strong overshoot, while expenditure on services does not, exactly as in the data.³⁰ Both sticky information and the flow adjustment costs are essential to the quality of the fit. First, sticky information allows us to match the hump-

²⁹Given our results in Section 2.5, we chose to keep the assumption of non-separable preferences because it is conservative. If durables and services were net substitutes, then the nCIR differences reported in Figure 7 would be even larger.

³⁰Note that, even though Figure 5 only reports *normalized* impulse responses, our impulse response matching in fact targeted *levels*.

Parameter	Description	Value	Source/Target
<i>Households</i>			
β	Discount Rate	0.99	Annual Real FFR
γ	Inverse EIS	1	Standard
ϕ	Durables Consumption Share	0.1	NIPA
θ	Sticky Information Friction	0.95	IRF matching
<i>Firms & Unions</i>			
ζ_p	Slope of the NKPC	0.02	Ajello et al. (2020)
ε_w	Labor Substitutability	10	Standard
δ	Depreciation Rate	0.021	BEA Fixed Asset
ϕ_w	Wage Re-Set Probability	0.2	Beraja et al. (2019)
κ	Level Adjustment Cost	0	IRF matching
κ_e	Flow Adjustment Cost	0.2	IRF matching
<i>Policy</i>			
ϕ_π	Inflation Response	1.5	Literature
<i>Shocks</i>			
$\rho_b = \rho_m$	Shock Persistence	0.83	Section 3.1

Table 4.1: Baseline parameterization of the quantitative structural model.




Figure 5: Impulse response matching: recursively identified monetary policy shock impulse responses from Figure 2 (grey) vs. model-implied impulse responses at the parameterization of Table 4.1 (purple).

shaped response of services spending. Second, while flow adjustment costs alone are enough to give a hump shape in durables expenditure, sticky information ensures that (i) this hump is delayed enough and (ii) adjustment costs can be sufficiently weak so as to not excessively moderate the subsequent, pent-up demand-related spending overshoot.

COMPUTING COUNTERFACTUALS. Note that, at our preferred parameterization, the model violates the assumptions required for the shift-share: prices and wages are only partially sticky, and the monetary rule is not neutral, so Proposition 4 does not apply. We are thus forced to rely on simulations from the model to compute our desired counterfactuals.

To compute those counterfactuals we proceed as follows. First, for variation in long-run expenditure shares ϕ , we compare cumulative impulse responses to an aggregate demand shock b_0^a , normalized so that at the trough total consumer expenditure declines by one per cent. Second, for variation in shock incidence, we choose combinations of sectoral and aggregate demand shocks $\{b_t^a, b_t^s, b_t^d\}$ that again have a trough effect on total expenditure of -1 per cent, with the shares at the trough corresponding to some given target.³¹

³¹To map our two-sector model to our empirical estimates, we assign non-durables to services. We emphasize that this choice is conservative: Since the share of non-durables in total expenditure declines is larger in ordinary and durables-led recessions than in services-led recessions, the estimated IRF differences and nCIR ratios in the right panels of Figures 6 and 7 would be even larger if non-durable spending was split between

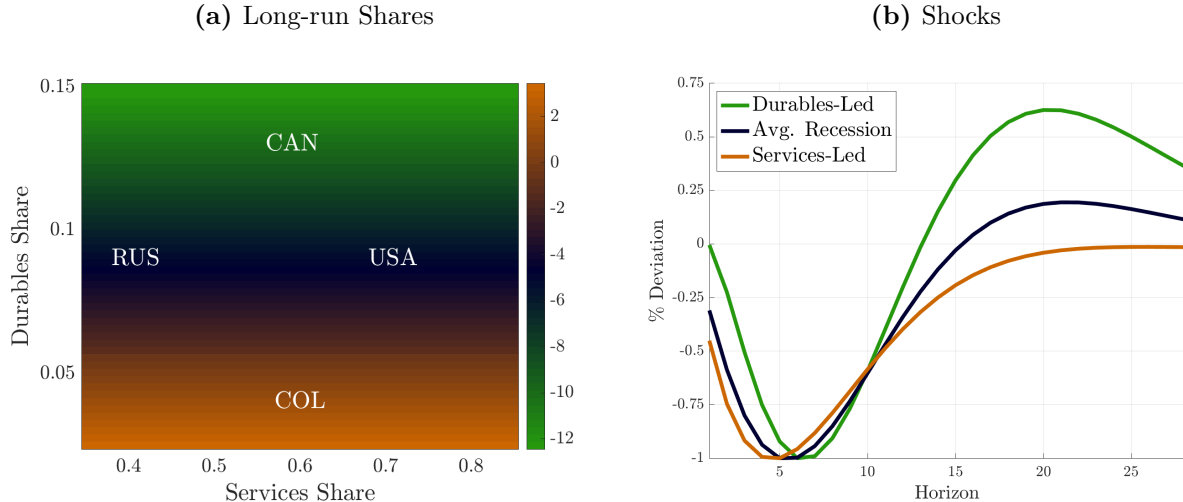


Figure 6: Left panel: nCIR to an aggregate demand shock b_0^a as a function of long-run expenditure shares and relative to the U.S. Right panel: Impulse response of total consumption to shock combinations giving a durables/non-durables/services trough spending composition of: (i) [100%, 50%, -50%] for a durables-led recession, (ii) [62%, 45%, -7%] for the average U.S. recession, and (iii) [12.5%, 12.5%, 75%] for a services-led recession. All are normalized to lead to a peak-to-trough consumption contraction of -1 per cent. For both panels we solve the estimated model at the posterior mode (Table 4.1).

RESULTS. We first compute our two counterfactuals at the posterior mode of the model, and then explore the robustness of our conclusions to changes in the strength of adjustment costs, the stickiness of prices, the monetary policy rule, and the persistence of shocks.

Figure 6 presents the analogue of the shift-share counterfactuals in Figure 4. The results are in line with those of the shift-share. First, changes in long-run shares result in meaningful differences in recovery dynamics following an aggregate demand shock b_0^a . Second, the shape of dynamic impulse responses following different menus of sectoral shocks exactly mimics that in the right panel of Figure 4: pent-up demand-driven overshoot for a durables-led recession, and a weak recovery after a services-led contraction. For both, the slightly muted effects relative to the results displayed in Figure 4 largely reflect our (conservative) choice to include non-durables in services. We conclude that, even though the simple shift-share approach is not exactly valid in our extended structural model, it remains a reliable guide to our two desired counterfactuals.

Having confirmed close agreement between shift-share and structural model at the mode, we now investigate the robustness of these conclusions to plausible variations in the model's

durables and services.




Figure 7: Left panel: Percentage gap between the nCIR to an aggregate demand shock b_0^a in an economy with the U.S. vs. Canada long-run expenditure shares, as a function of adjustment costs (x -axis) and the NKPC slope (y -axis). Right panel: Percentage gap between the nCIR to demand shocks (b_0^a, b_0^d, b_0^s) inducing a composition of expenditure changes on impact as in a services-led vs. an average U.S. recession, again as a function of adjustment costs (x -axis) and the NKPC slope (y -axis). The red cross in both figures indicates our preferred parameterization. All other parameters are set as in Table 4.1.

parameterization. Since our analysis has revealed the strength of adjustment costs as well as the degree of price stickiness to play an important role in shaping the strength of pent-up demand effects, we begin in Figure 7 by examining sensitivity to the adjustment cost parameter κ_e and the NKPC slope ζ_p . To do so, we consider many different pairs of $\{\kappa_e, \zeta_p\}$, and for each compute nCIRs for (i) aggregate demand shocks in the U.S. and Canada — two economies with very different durables shares — and (ii) shock combinations that lead to a recession with an ordinary spending composition vs. that of a services-led recession. Our main finding is that recessions more biased towards services are invariably and robustly followed by weaker recoveries. Around our preferred parameter estimates of $\zeta_p = 0.02$ and $\kappa_e = 0.2$ (marked with the red cross), the results align quantitatively with those of the shift-share, as discussed before. As expected, the nCIR difference declines as adjustment costs become stronger (offsetting the pent-up demand effect) and prices become more flexible (leading to a more aggressive policy response). Even with counterfactually large price flexibility and adjustment costs, however, the implied nCIR differences robustly remain sizable.

POLICY RESPONSE AND SHOCK PERSISTENCE. Appendix B.4 presents results for an analogous exercise that varies the shock persistence ρ_b and the monetary policy rule coefficient

ϕ_π . We again find that, across a wide range of plausible parameter values, our conclusions go through largely unchanged. We emphasize that this exercise is particularly informative about the recovery from the COVID-19 recession. Here, the relatively fast arrival of vaccines in the U.S. can be thought of as the underlying shock having a low persistence. For ρ_b equal to 0.2 instead of our benchmark 0.9, output quickly returns to baseline five quarters after the trough, about the time it took for U.S. consumption expenditures to reach their pre-crisis level. Importantly, however, the pent-up demand logic suggests that we would have expected the recovery to be even stronger if not for the COVID-19 recession’s unusual demand composition. In other words, the recovery has been “quick” compared to some recent episodes like the Great Recession but rather “weak” compared to the counterfactual, given how transitory the shock was (as well as how unusually strong the monetary and fiscal policy responses were). Our empirical cross-regional evidence in Section 3.3 is very much consistent with this interpretation.

4.4 Summary

We view our two approaches in Sections 4.2 and 4.3 as complementary: the shift-share is appealingly semi-structural but relies on some restrictive assumptions, while the structural model relaxes those assumptions yet commits to particular frictions in durables and services demand. Encouragingly, both of them paint a consistent picture: viewed through the lens of standard multi-sector business-cycle models, the empirical evidence in Figure 2 robustly implies that recoveries from demand-driven recessions concentrated in services would tend to be weaker than after recessions more biased toward durables.

5 Policy implications

We have argued that the sectoral expenditure composition during demand-driven recessions can have sizable effects on recovery dynamics. Given this quantitative significance, we conclude our analysis by studying the implications of the pent-up demand mechanism for optimal (monetary) stabilization policy. Mirroring our two counterfactuals in Section 4, we proceed in two steps: first we study optimal policy following shocks to all categories of spending, and then we turn to explicitly sectoral demand shocks.

5.1 Optimal policy under aggregate demand shocks

We return to our analytically tractable baseline model of Section 2.1. For our first experiment we assume that the economy is subject only to *aggregate* demand shocks b_t^a , and rule out any sectoral shocks b_t^s or b_t^d . In this setting, we are interested in studying whether monetary policy should behave differently (or not) in more durables- or services-intensive economies. The following proposition characterizes the flexible-price allocation and so the first-best policy.

Proposition 5. *Consider the model of Section 2.1, simplified to feature only shocks to aggregate demand b_t^a . Then the first-best monetary policy sets*

$$\widehat{r}_t^n = (1 - \rho_b)b_t^a \tag{15}$$

In particular, it follows that the optimal monetary policy is independent of the long-run durables expenditure share $1 - \phi$.

The intuition is as follows. Changes in the durables share $1 - \phi$ affect the transmission of *both* aggregate demand shocks b_t^a and conventional interest rate policy. With our definition of an aggregate demand shock b_t^a , these two effects exactly offset, leaving optimal monetary policy as a function of b_t^a completely unchanged. It follows in particular that the Wicksellian rate of interest — defined in Woodford (2011) as the equilibrium rate of return with fully flexible prices — is independent of the durables share, and so behaves exactly as in conventional business-cycle models with only non-durable consumption.

Proposition 5 connects the findings of McKay & Wieland (2020) to questions of optimal policymaking. McKay & Wieland study the transmission of monetary policy *shocks* in an environment with durable consumption, and argue that monetary authorities face an intertemporal trade-off: interest rate cuts today pull demand forward in time, pushing output below its natural level in the future. The results here reveal that this does not necessarily imply any trade-off for *optimal* policymaking: while interest rate cuts today do indeed lead to deficient demand tomorrow, negative exogenous demand shocks today at the same time lead to excess demand tomorrow. Putting the two together, optimal policy is left unaffected.

While clearly knife-edge, this independence result illustrates a more general principle: when studying the consequences of a particular mechanism for the conduct of stabilization policy, one needs to assess not only how that mechanism affects the transmission of policy itself, but also whether it alters the transmission of the underlying exogenous shocks that the policy is trying to offset.

5.2 Optimal policy under sectoral demand shocks

We now study how optimal policy depends on the sectoral incidence of shocks. We are able to state a particularly clean result in the special case of transitory shocks ($\rho_b = 0$) and no adjustment costs ($\kappa = 0$).

Proposition 6. *Consider the model of Section 2.1 with $\kappa = 0$ and $\rho_b = 0$, and let $\hat{r}_t(b_0^i)$ with $i \in \{s, d\}$ denote the flexible-price equilibrium real interest rate at t given a time-0 shock b_0^i . Then, for shocks b_0^s and b_0^d such that $\hat{r}_0(b_0^s) = \hat{r}_0(b_0^d) < 0$, we have*

$$\hat{r}_t(b_0^s) < \hat{r}_t(b_0^d), \quad \forall t \geq 2 \quad (16)$$

Thus, the optimal monetary policy eases strictly longer following a services demand shock compared to a durables demand shock.

We have already seen that, under a neutral monetary policy rule, a services shock leads to a persistent contraction in aggregate output, while an otherwise identical durables shock leads to a more short-lived one. If the monetary authority cuts rates in the face of such sectoral shocks, it invariably stimulates the initially unaffected sector. Proposition 6 reveals what this stimulus — written in terms of real rates — should look like: persistent in the case of a recession biased towards services, and more short-lived after a durables-led contraction.

The intuition for the ordering of the real rate paths in (16) is simple. Given a negative services shock, the monetary authority cuts real rates, stimulating durables expenditures. In the following periods, the durables stock is gradually run down, so services consumption can remain relatively elevated. This high level of services consumption is supported through persistently low real interest rates. Conversely, given a negative durables shock, future real interest rates are relatively high to depress services expenditures and allow the durables stock to be re-built gradually. It follows that, relative to the baseline equilibrium rate of interest for aggregate demand shocks, the Wicksellian rate paths for pure services and durables demand shocks are tilted down and up, respectively.

Finally we emphasize that, while Proposition 6 is stated only for a very special case of the model, the same forces remain present in richer environments. In the more general model with persistent shocks ($\rho_b > 0$) and adjustment costs ($\kappa > 0$), we have that $\hat{r}_t(b_0^a) = -\rho_b^t$ (consistent with Proposition 5) and

$$\hat{r}_t(b_0^s) = -\rho_b^t - \zeta_s \sum_{q=0}^{t-1} \rho_b^{t-q} \vartheta^q \quad (17)$$

$$\widehat{r}_t(b_0^d) = -\rho_b^t + \zeta_d \sum_{q=0}^{t-1} \rho_b^{t-q} \vartheta^q \quad (18)$$

where the parameters $\{\zeta_s, \zeta_d, \vartheta\}$ are functions of primitive model parameters. Without adjustment costs and with transitory shocks, we can prove that $\{\zeta_s, \zeta_d, \vartheta\}$ are all strictly positive, establishing the desired result. Numerically, we have verified that they remain positive in the more general model for a wide range of values of ρ_b and κ , ensuring that an extended version of Proposition 6 continues to hold (see Appendix B.4).

6 Conclusions

We have argued that recoveries from demand-driven recessions with expenditure cuts concentrated in services or non-durables will tend to be weaker than recoveries from recessions biased towards durables. This prediction follows from standard consumer theory together with output being demand-determined. We have documented empirical support for the underlying pent-up demand mechanism in U.S. time series data, and also shown evidence consistent with the prediction in the cross-section of U.S. regions. Our quantitative analysis furthermore suggests that the effect of demand composition on recovery strength can be meaningful. Finally, moving from positive to normative analysis, we have shown that the pent-up demand mechanism matters for optimal policy design: if the monetary authority were to ignore the sectoral incidence of shocks and instead apply a simple one-size-fits-all policy to all recessions, then interest rates would be hiked *too fast* in services recessions, and aggregate output would remain depressed for longer.

We leave several important avenues for future research. First, time-varying household adjustment hazards for durable spending imply that the strength of pent-up demand effects may well vary over the business cycle. Second, we have only studied the optimal policy implications of the pent-up demand mechanism for aggregate and sectoral *demand* shocks. It would be interesting to see how it interacts with other kinds of prominent business-cycle disturbances. Third, we have emphasized that our results on sectoral spending impulse responses and predicted recovery dynamics are not intrinsically tied to the pent-up demand mechanism and so the durables-services dichotomy. It would be interesting to use the semi-structural shift-share developed here to see whether other mechanisms predict similarly large differences in recovery dynamics.

References

- Ajello, A., Cairó, I., Cúrdia, V., Lubik, T., & Queralto, A. (2020). Monetary policy tradeoffs and the federal reserve's dual mandate. FEDS Working Paper.
- Arias, J. E., Rubio-Ramírez, J. F., & Waggoner, D. F. (2018). Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications. *Econometrica*, *86*(2), 685–720.
- Auclert, A., Rognlie, M., & Straub, L. (2020). Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model. Technical report, National Bureau of Economic Research.
- Barsky, R. B., House, C. L., & Kimball, M. S. (2007). Sticky-price models and durable goods. *American Economic Review*, *97*(3), 984–998.
- Basu, S. & Bundick, B. (2017). Uncertainty shocks in a model of effective demand. *Econometrica*, *85*(3), 937–958.
- Beraja, M., Hurst, E., & Ospina, J. (2019). The aggregate implications of regional business cycles. *Econometrica*, *87*(6), 1789–1833.
- Berger, D. & Vavra, J. (2015). Consumption dynamics during recessions. *Econometrica*, *83*(1), 101–154.
- Bigio, S. & La'ò, J. (2020). Distortions in production networks. *The Quarterly Journal of Economics*, *135*(4), 2187–2253.
- Bilbiie, F. O. (2018). Monetary policy and heterogeneity: An analytical framework. CEPR Discussion Paper No. DP12601.
- Bilbiie, F. O. (2020). The new keynesian cross. *Journal of Monetary Economics*, *114*, 90–108.
- Boppart, T., Krusell, P., & Mitman, K. (2018). Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control*, *89*, 68–92.
- Caballero, R. J. (1993). Durable goods: An explanation for their slow adjustment. *Journal of Political Economy*, *101*(2), 351–384.

- Carvalho, C. (2006). Heterogeneity in price stickiness and the real effects of monetary shocks. *The BE Journal of Macroeconomics*, 6(3).
- Carvalho, V. M. & Grassi, B. (2019). Large firm dynamics and the business cycle. *American Economic Review*, 109(4), 1375–1425.
- Chari, V. V., Kehoe, P. J., & McGrattan, E. R. (2007). Business cycle accounting. *Econometrica*, 75(3), 781–836.
- Chetty, R., Friedman, J. N., Hendren, N., & Stepner, M. (2020). Real-time economics: A new platform to track the impacts of covid-19 on people, businesses, and communities using private sector data. Technical report, Mimeo.
- Christiano, L. J., Eichenbaum, M., & Evans, C. L. (1999). Monetary policy shocks: What have we learned and to what end? *Handbook of macroeconomics*, 1, 65–148.
- Christiano, L. J., Eichenbaum, M., & Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of political Economy*, 113(1), 1–45.
- Coibion, O., Gorodnichenko, Y., & Koustas, D. (2013). Amerisclerosis? the puzzle of rising us unemployment persistence. Technical report, National Bureau of Economic Research.
- Cox, N., Ganong, P., Noel, P., Vavra, J., Wong, A., Farrell, D., & Greig, F. (2020). Initial impacts of the pandemic on consumer behavior: Evidence from linked income, spending, and savings data. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2020-82).
- Erceg, C. & Levin, A. (2006). Optimal monetary policy with durable consumption goods. *Journal of monetary Economics*, 53(7), 1341–1359.
- Erceg, C. J., Henderson, D. W., & Levin, A. T. (2000). Optimal Monetary Policy with Staggered Wage and Price Contract. *Journal of Monetary Economics*, 46(2), 281–313.
- Farhi, E. & Baqaee, D. R. (2020). Supply and demand in disaggregated keynesian economies with an application to the covid-19 crisis. Working Paper.
- Fernald, J. G., Hall, R. E., Stock, J. H., & Watson, M. W. (2017). The disappointing recovery of output after 2009. Technical report, National Bureau of Economic Research.

- Fukui, M., Nakamura, E., & Steinsson, J. (2018). Women, wealth effects, and slow recoveries. Technical report, National Bureau of Economic Research.
- Galí, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.
- Galí, J., Smets, F., & Wouters, R. (2012). Slow recoveries: A structural interpretation. *Journal of Money, Credit and Banking*, 44, 9–30.
- Gregory, V., Menzio, G., & Wiczer, D. G. (2020). Pandemic recession: L or v-shaped? Technical report, National Bureau of Economic Research.
- Grigsby, J., Hurst, E., & Yildirmaz, A. (2019). Aggregate nominal wage adjustments: New evidence from administrative payroll data. Technical report, National Bureau of Economic Research.
- Guerrieri, V., Lorenzoni, G., Straub, L., & Werning, I. (2020). Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? Technical report, National Bureau of Economic Research.
- Hai, R., Krueger, D., & Postlewaite, A. (2013). On the welfare cost of consumption fluctuations in the presence of memorable goods. Technical report, National Bureau of Economic Research.
- Hall, R. E. (2016). Macroeconomics of persistent slumps. In *Handbook of Macroeconomics*, volume 2 (pp. 2131–2181). Elsevier.
- Hall, R. E. & Kudlyak, M. (2020). Why has the us economy recovered so consistently from every recession in the past 70 years? *NBER Working Paper*, (w27234).
- Hamilton, J. D. (2003). What is an oil shock? *Journal of econometrics*, 113(2), 363–398.
- Hazell, J., Herreño, J., Nakamura, E., & Steinsson, J. (2020). The slope of the phillips curve: evidence from us states. Technical report, National Bureau of Economic Research.
- House, C. L. & Shapiro, M. D. (2008). Temporary investment tax incentives: Theory with evidence from bonus depreciation. *American Economic Review*, 98(3), 737–68.
- Justiniano, A., Primiceri, G. E., & Tambalotti, A. (2010). Investment shocks and business cycles. *Journal of Monetary Economics*, 57(2), 132–145.

- Kozlowski, J., Veldkamp, L., & Venkateswaran, V. (2020). The tail that wags the economy: Beliefs and persistent stagnation. *Journal of Political Economy*, 128(8), 2839–2879.
- La’O, J. & Tahbaz-Salehi, A. (2020). Optimal monetary policy in production networks. Technical report, National Bureau of Economic Research.
- Lubik, T. A. & Schorfheide, F. (2004). Testing for indeterminacy: An application to us monetary policy. *American Economic Review*, 94(1), 190–217.
- Mankiw, N. G. (1982). Hall’s consumption hypothesis and durable goods. *Journal of Monetary Economics*, 10(3), 417–425.
- Mankiw, N. G. & Reis, R. (2002). Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics*, 117(4), 1295–1328.
- Mankiw, N. G. & Reis, R. (2003). What measure of inflation should a central bank target? *Journal of the European Economic Association*, 1(5), 1058–1086.
- McKay, A. & Wieland, J. F. (2020). Lumpy durable consumption demand and the limited ammunition of monetary policy. Technical report, National Bureau of Economic Research.
- Nakamura, E. & Steinsson, J. (2010). Monetary non-neutrality in a multisector menu cost model. *The Quarterly journal of economics*, 125(3), 961–1013.
- Olney, M. L. & Pacitti, A. (2017). The rise of services, deindustrialization, and the length of economic recovery. *Economic Inquiry*, 55(4), 1625–1647.
- Pasten, E., Schoenle, R., & Weber, M. (2017). Price rigidity and the origins of aggregate fluctuations. *NBER Working Paper*, (w23750).
- Plagborg-Møller, M. & Wolf, C. K. (2021). Local projections and vars estimate the same impulse responses. *Econometrica*, 89(2), 955–980.
- Ramey, V. A. (2016). Macroeconomic shocks and their propagation. In *Handbook of macroeconomics*, volume 2 (pp. 71–162). Elsevier.
- Reis, R. (2020). An ABC recovery. Technical report, <https://threadreaderapp.com/thread/1253988696749150208.html>.

- Rognlie, M., Shleifer, A., & Simsek, A. (2018). Investment hangover and the great recession. *American Economic Journal: Macroeconomics*, 10(2), 113–53.
- Rubbo, E. (2020). Networks, phillips curves and monetary policy. Working Paper.
- Schmitt-Grohé, S. & Uribe, M. (2017). Liquidity traps and jobless recoveries. *American Economic Journal: Macroeconomics*, 9(1), 165–204.
- Smets, F. & Wouters, R. (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American economic review*, 97(3), 586–606.
- Wolf, C. K. (2020). Svar (mis) identification and the real effects of monetary policy shocks. *American Economic Journal: Macroeconomics*, 12(4), 1–32.
- Woodford, M. (2011). *Interest and prices: Foundations of a theory of monetary policy*. princeton university press.

Online Appendix for: Demand Composition and the Strength of Recoveries

This online appendix contains supplemental material for the article “Demand Composition and the Strength of Recoveries.” We provide (i) further details for the various structural models used in the paper, (ii) several supplementary theoretical results, and (iii) implementation details for all empirical exercises. The end of this appendix contains further proofs and auxiliary lemmas.

Any references to equations, figures, tables, assumptions, propositions, lemmas, or sections that are not preceded “A.” – “D.” refer to the main article.

A Model appendix

In this appendix we provide further details on the structural models of Section 2. First, in Appendix A.1, we elaborate on the baseline model of Section 2.1. Then, in Appendix A.2, we present the various model extensions introduced in Section 2.5.

A.1 Detailed model outline

HOUSEHOLDS. The household consumption-savings problem is described fully in Section 2.1, up to the link between the preference parameter $\tilde{\phi}$ and the spending share ϕ , as well as the scaling factor α in our specification of household preferences.

For $\tilde{\phi}$, we from the steady-state first-order conditions get

$$\left(\frac{\tilde{\phi}}{1-\tilde{\phi}}\right)^\gamma = \frac{1}{1-\beta(1-\delta)} \left(\frac{\phi}{\frac{1}{\delta}(1-\phi)}\right)^\gamma \quad (\text{A.1})$$

Next, we set the scaling factor $\{\alpha\}$ to ensure that b_t^a has no first order effects on any real quantities in a flexible price equilibrium. The required factor can be shown to be:

$$\alpha \equiv 1 + \frac{\beta(1-\delta)(1-\rho_b)}{1-\beta(1-\delta)} \quad (\text{A.2})$$

For future reference, it will be useful to let

$$c_t \equiv \left[e^{b_t^a + b_t^s} \tilde{\phi}^\gamma s_t^{1-\gamma} + e^{\alpha(b_t^a + b_t^d)} (1-\tilde{\phi})^\gamma d_t^{1-\gamma} \right]^{\frac{1}{1-\gamma}}$$

denote the total household consumption bundle. Note that, to first order, this bundle satisfies

$$\begin{aligned} \hat{c}_t = & \frac{\phi}{\phi + [1-\beta(1-\delta)]\frac{1}{\delta}(1-\phi)} \left(\hat{s}_t + \frac{1}{1-\gamma} (b_t^a + b_t^s) \right) \\ & + \frac{[1-\beta(1-\delta)]\frac{1}{\delta}(1-\phi)}{\phi + [1-\beta(1-\delta)]\frac{1}{\delta}(1-\phi)} \left(\hat{d}_t + \frac{\alpha}{1-\gamma} (b_t^a + b_t^d) \right) \end{aligned} \quad (\text{A.3})$$

We now state the first-order conditions characterizing optimal household behavior. The marginal utility of wealth λ_t satisfies

$$\hat{\lambda}_t = \hat{r}_t^n - \mathbb{E}_t [\hat{\pi}_{t+1}] + \mathbb{E}_t [\hat{\lambda}_{t+1}] \quad (\text{A.4})$$

The first-order conditions for services and durables are

$$-\gamma \widehat{s}_t + (b_t^a + b_t^s) = \widehat{\lambda}_t \quad (\text{A.5})$$

$$\begin{aligned} -\gamma \widehat{d}_t + \alpha(b_t^a + b_t^d) &= \frac{1}{1 - \beta(1 - \delta)} \left[\widehat{\lambda}_t + \kappa(\widehat{d}_t - \widehat{d}_{t-1}) \right] \\ &\quad - \frac{\beta(1 - \delta)}{1 - \beta(1 - \delta)} \mathbb{E}_t \left[\widehat{\lambda}_{t+1} + \frac{\kappa}{1 - \delta} (\widehat{d}_{t+1} - \widehat{d}_t) \right] \end{aligned} \quad (\text{A.6})$$

Given the definition of the scaling factor α , we can re-write those conditions as

$$-\gamma \widehat{s}_t = \widehat{\lambda}_t - (b_t^a + b_t^s), \quad (\text{A.7})$$

$$\begin{aligned} -\gamma \widehat{d}_t &= \frac{1}{1 - \beta(1 - \delta)} \left[\widehat{\lambda}_t - (b_t^a + b_t^d) + \kappa(\widehat{d}_t - \widehat{d}_{t-1}) \right] \\ &\quad - \frac{\beta(1 - \delta)}{1 - \beta(1 - \delta)} \mathbb{E}_t \left[\widehat{\lambda}_{t+1} - (b_{t+1}^a + b_{t+1}^d) + \frac{\kappa}{1 - \delta} (\widehat{d}_{t+1} - \widehat{d}_t) \right] \end{aligned} \quad (\text{A.8})$$

where we have used that $\mathbb{E}_t(b_{t+1}^a) = \rho_b b_t^a$ and $\mathbb{E}_t(b_{t+1}^d) = \rho_b b_t^d$. This alternative way of writing the first-order conditions reveals cleanly that our aggregate and sectoral demand shocks are constructed to be isomorphic to shocks to the shadow prices of the total household consumption bundle and the two sectoral goods, respectively. In particular, this ensures that the aggregate demand shock can be perfectly offset by movements in real interest rates (via (A.4)), ensuring the desirable neutrality property discussed in Footnote 7.

Finally, optimal household labor supply relates real wages \widehat{w}_t , inflation $\widehat{\pi}_t$, hours worked $\widehat{\ell}_t$, the marginal utility of wealth $\widehat{\lambda}_t$, and shocks $\{b_t^a, b_t^s, b_t^d\}$:

$$\widehat{\pi}_t^w = \frac{(1 - \beta\phi_w)(1 - \phi_w)}{\phi_w(\frac{\varepsilon_w}{\varphi} + 1)} \left[\frac{1}{\varphi} \widehat{\ell}_t - \left(\widehat{w}_t + \widehat{\lambda}_t - (b_t^a + \phi b_t^s + (1 - \phi)b_t^d) \right) \right] + \beta \mathbb{E}_t [\widehat{\pi}_{t+1}^w] \quad (\text{A.9})$$

PRODUCTION. We assume that both durables and services are produced by aggregating a common set of varieties sold by monopolistically competitive retailers, modeled exactly as in Galí (2015, Chapter 3). This set-up implies that real relative prices are always equal to 1 (i.e., $\widehat{p}_t^s = \widehat{p}_t^d = 0$).

We can thus summarize the production side of the economy with a single aggregate New Keynesian Phillips curve, relating inflation $\widehat{\pi}_t$ to the real wage \widehat{w}_t and hours $\widehat{\ell}_t$:

$$\widehat{\pi}_t = \zeta_p \left(\widehat{w}_t - \frac{y''(\ell)\ell}{y'(\ell)} \widehat{\ell}_t \right) + \beta \mathbb{E}_t [\widehat{\pi}_{t+1}] \quad (\text{A.10})$$

where ζ_p is a function of the discount factor β , the production function of retailers $y(l)$, and the degree of price stickiness. For much of our analysis we need to merely assume that prices are not perfectly flexible, so $\zeta_p < \infty$; if so, the central bank can fix the expected real interest rate, and — under our assumptions on equilibrium selection — the NKPC (A.10) as well as the details of the production function $y = y(\ell)$ are irrelevant for all aggregate quantities. Throughout this paper, we restrict attention to the simple case of a linear production technology, so $y''(\ell) = 0$.

Firms discount at the stochastic discount factor of their owners (the representative household), and pay out dividends q_t . The dynamics of dividends are irrelevant for our purposes, so we do not discuss them further.

EXAMPLE PARAMETERIZATION. For our simple graphical illustration in Figure 1 we set $\gamma = 1$, $\beta = 0.99$, $\delta = 0.021$, $\rho_b = 0$, $\kappa = 0$ and $\underline{\phi} = 0.9$. We then choose $\bar{\phi}$ to construct a recession with equal shares for durables and services spending.

A.2 Further extensions

MANY SECTORS. Household preferences over consumption bundles are now given as

$$u(d; b) = \frac{\sum_{i=1}^N e^{\alpha_i(b^a + b^i)} \tilde{\phi}_i d_{it}^{1-\gamma} - 1}{1 - \gamma}$$

where the scaling coefficients α_i are defined as in (A.2). We normalize the expenditure share of good i to ϕ_i ; the preference parameters $\tilde{\phi}_i$ are then defined implicitly via optimal household behavior, as discussed in Appendix A.1. The budget constraint becomes

$$\sum_{i=1}^N \left\{ p_t^i \underbrace{[d_{it} - (1 - \delta_i)d_{it-1}]}_{e_{it}} + \psi_i(d_{it}, d_{it-1}) \right\} + a_t = w_t \ell_t + \frac{1 + r_{t-1}^n}{1 + \pi_t} a_{t-1} + q_t$$

and finally the linearized output market-clearing condition is

$$\hat{y}_t = \sum_{i=1}^N \phi_i \hat{e}_{it}.$$

All other model equations are unchanged.

GENERAL ADJUSTMENT COSTS. The adjustment cost function is

$$\psi(\{d_{t-\ell}\}_{\ell=0}^{\infty}).$$

In particular, this specification encompasses the quadratic adjustment costs on durable expenditure flows that we consider in Section 4.3:

$$\psi(e, e_{-1}) = \frac{\kappa_e}{2} \left(\frac{e}{e_{-1}} - 1 \right)^2 e. \quad (\text{A.11})$$

NON-SEPARABILITY. The preferences are now given by

$$u(s, d; b) = \frac{\left[e^{b^a + b^s} \tilde{\phi}^{\zeta} s^{1-\zeta} + e^{\alpha(b^a + b^d)} (1 - \tilde{\phi})^{\zeta} d^{1-\zeta} \right]^{\frac{1-\gamma}{1-\zeta}} - 1}{1 - \gamma}, \quad v(\ell; b) = e^{\varsigma_a b^a + \varsigma_s b^s + \varsigma_d b^d} \chi \frac{\ell^{1 + \frac{1}{\varphi}}}{1 + \frac{1}{\varphi}}$$

where ζ is the inverse elasticity of substitution between services and durables, and preference scaling coefficients are now given by

$$\begin{aligned} \left(\frac{\tilde{\phi}}{1 - \tilde{\phi}} \right)^{\zeta} &= \frac{1}{1 - \beta(1 - \delta)} \left(\frac{\phi}{\frac{1}{\delta}(1 - \phi)} \right)^{\zeta} \\ \alpha &\equiv 1 + \varsigma_a \frac{\beta(1 - \delta)(1 - \rho_b)}{1 - \beta(1 - \delta)} \\ \varsigma_a &\equiv \frac{1 + \frac{\zeta - \gamma}{1 - \zeta}}{1 - \frac{\zeta - \gamma}{1 - \zeta} \frac{\frac{1}{\delta}(1 - \phi)\beta(1 - \delta)(1 - \rho_b)}{\phi + (1 - \beta(1 - \delta))^{\frac{1}{\delta}}(1 - \phi)}} \\ \varsigma_s &\equiv \varsigma_a \phi \\ \varsigma_d &\equiv \varsigma_a (1 - \phi) \end{aligned}$$

The above coefficients are obtained from the modified household first-order conditions

$$(\zeta - \gamma)\widehat{c}_t - \zeta\widehat{s}_t + (b_t^a + b_t^s) = \widehat{\lambda}_t \quad (\text{A.12})$$

$$\begin{aligned} (\zeta - \gamma)\widehat{c}_t - \zeta\widehat{d}_t + \alpha(b_t^a + b_t^d) &= \frac{1}{1 - \beta(1 - \delta)} \left[\widehat{\lambda}_t + \kappa(\widehat{d}_t - \widehat{d}_{t-1}) \right] \\ &\quad - \frac{\beta(1 - \delta)}{1 - \beta(1 - \delta)} \mathbb{E}_t \left[\widehat{\lambda}_{t+1} + \frac{\kappa}{1 - \delta} (\widehat{d}_{t+1} - \widehat{d}_t) \right], \quad (\text{A.13}) \end{aligned}$$

with

$$\begin{aligned}\widehat{c}_t &= \frac{\phi}{\phi + [1 - \beta(1 - \delta)]^{\frac{1}{\delta}}(1 - \phi)} \left(\widehat{s}_t + \frac{1}{1 - \zeta} (b_t^a + b_t^s) \right) \\ &+ \frac{[1 - \beta(1 - \delta)]^{\frac{1}{\delta}}(1 - \phi)}{\phi + [1 - \beta(1 - \delta)]^{\frac{1}{\delta}}(1 - \phi)} \left(\widehat{d}_t + \frac{\alpha}{1 - \zeta} (b_t^a + b_t^d) \right)\end{aligned}\tag{A.14}$$

and the wage setting equation

$$\begin{aligned}\widehat{\pi}_t^w &= \frac{(1 - \beta\phi_w)(1 - \phi_w)}{\phi_w(\frac{\varepsilon_w}{\varphi} + 1)} \left[\frac{1}{\varphi} \widehat{\ell}_t - \left(\widehat{w}_t + \widehat{\lambda}_t - (\varsigma_a b_t^a + \varsigma_s b_t^s + \varsigma_d b_t^d) \right) \right] \\ &+ \beta \mathbb{E}_t [\widehat{\pi}_{t+1}^w]\end{aligned}$$

STICKY INFORMATION. To study sticky information we cast our model in sequence-space, as for example in Auclert et al. (2020) and McKay & Wieland (2020).³² In sequence-space, we can for each consumption variable $x \in \{c, s, d, e\}$ summarize consumer behavior with aggregate consumption functions of the form

$$\mathbf{x} = \mathcal{X}(\mathbf{r}^n, \boldsymbol{\pi}, \mathbf{b}^a, \mathbf{b}^s, \mathbf{b}^d)$$

where boldface notation denotes time paths. Since we are restricting attention to small shocks, we can in fact focus on the consumption derivative matrices

$$\mathcal{X}_p \equiv \frac{\partial \mathcal{X}(\bullet)}{\partial \mathbf{p}}\tag{A.15}$$

where $p \in \{r^n, \pi, b^a, b^s, b^d\}$ denotes a generic input to the consumption-savings problem.

We use R superscripts to indicate the derivative matrices corresponding to the baseline full-information, rational expectations consumption-savings problem summarized by (A.3) - (A.6). Sticky information then simply modifies these baseline matrices as follows:

$$\mathcal{X}_p = \begin{pmatrix} (1 - \theta)\mathcal{X}_{p,0,0}^R & (1 - \theta)\mathcal{X}_{p,0,1}^R & (1 - \theta)\mathcal{X}_{p,0,2}^R & \dots \\ (1 - \theta)\mathcal{X}_{p,1,0}^R & (1 - \theta)\mathcal{X}_{p,1,1}^R + (\theta - \theta^2)\mathcal{X}_{p,0,0}^R & (1 - \theta)\mathcal{X}_{p,1,2}^R + (\theta - \theta^2)\mathcal{X}_{p,0,1}^R & \dots \\ (1 - \theta)\mathcal{X}_{p,2,0}^R & (1 - \theta)\mathcal{X}_{p,2,1}^R + (\theta - \theta^2)\mathcal{X}_{p,1,0}^R & (1 - \theta)\mathcal{X}_{p,2,2}^R + (\theta - \theta^2)\mathcal{X}_{p,1,1}^R + (\theta^2 - \theta^3)\mathcal{X}_{p,0,0}^R & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

³²All results presented here continue to apply unchanged to the first-order state-space representation derived from the familiar perturbation solution, by the equivalence discussed in Boppart et al. (2018).

or more succinctly

$$\mathcal{X}_{p,i,j} = \sum_{s=0}^{\min\{i,j\}} [\theta^s - \theta^{s+1}] \mathcal{X}_{p,i,j}^R$$

All other model equations are unaffected by the presence of sticky information.

INCOMPLETE MARKETS. The model is populated by a mass $1 - \mu$ of households identical to the representative household of Section 2.1, and a residual fringe $\mu \in (0, 1)$ of hand-to-mouth households. Following Bilbiie (2018), we simply impose the reduced-form assumption that total income (and so total consumption) of every hand-to-mouth household H satisfies

$$\phi \widehat{s}_t^H + (1 - \phi) \widehat{e}_t^H = \eta \widehat{y}_t$$

Hand-to-mouth households have the same preferences as unconstrained households. Their consumption problem is thus to optimally allocate their exogenous income stream between durable and non-durable consumption, subject to the constraint that their bond holdings have to be zero at all points in time. We present the equations characterizing optimal behavior of hand-to-mouth households in the proof of Proposition B.6. All other model blocks are unaffected by the presence of hand-to-mouth households.³³

SUPPLY SHOCKS. We consider a simple model of (sectoral) productivity shocks in which innovations in productivity are completely passed through to goods prices. Analogously to our baseline model, we consider three shocks $\{z_t^a, z_t^s, z_t^d\}$ with common persistence ρ_z ; their relative volatilities are irrelevant for all results discussed here. Assuming that monetary policy fixes the real rate in terms of intermediate goods prices, real relative prices satisfy

$$\widehat{p}_t^s = -(z_t^a + z_t^s) \tag{A.16}$$

$$\widehat{p}_t^d = -(z_t^a + z_t^d) \tag{A.17}$$

The output market-clearing condition then becomes

$$\widehat{y}_t = [z_t^a + \phi z_t^s + (1 - \phi) z_t^d] + \widehat{\ell}_t = \phi \widehat{s}_t + (1 - \phi) \widehat{e}_t \tag{A.18}$$

All other model equations are unchanged.

³³In particular this includes the labor supply relation (A.9): unions bargain for both types of households, so the aggregate labor supply relation is unchanged.

ALTERNATIVE SPECIFICATION FOR b_t^a . A natural alternative specification for household consumption preferences is

$$u(s, d; b) = \frac{e^{b^a} \left[e^{b^s} \tilde{\phi}^\gamma s^{1-\gamma} + e^{b^d} (1 - \tilde{\phi})^\gamma d^{1-\gamma} \right] - 1}{1 - \gamma} \quad (\text{A.19})$$

Here, the aggregate demand shock b^a affects the valuation of the total consumption bundle in one period relative to the next. With this specification, (A.4) still applies without change, but the first-order conditions for services and durables become

$$-\gamma \widehat{s}_t + (b_t^a + b_t^s) = \widehat{\lambda}_t \quad (\text{A.20})$$

$$\begin{aligned} -\gamma \widehat{d}_t + (b_t^a + b_t^d) &= \frac{1}{1 - \beta(1 - \delta)} \left[\widehat{\lambda}_t + \kappa(\widehat{d}_t - \widehat{d}_{t-1}) \right] \\ &\quad - \frac{\beta(1 - \delta)}{1 - \beta(1 - \delta)} \mathbb{E}_t \left[\widehat{\lambda}_{t+1} + \frac{\kappa}{1 - \delta} (\widehat{d}_{t+1} - \widehat{d}_t) \right] \end{aligned} \quad (\text{A.21})$$

Note that the two sectoral shocks $\{b_t^s, b_t^d\}$ enter exactly as in our baseline system (up to scale). The aggregate demand shock b_t^a is now however scaled down in the durables first-order condition, breaking the desired real neutrality property as well as the equivalence to ordinary monetary shocks (which enter like shocks to the path of λ_t).

B Supplementary theoretical results

This section offers various supplementary theoretical results. First, in Appendix B.1, we offer closed-form solutions for impulse responses in the general model, supplementing the analysis in Section 2.4. In Appendix B.2 we then state formal results for the various model extensions discussed in Section 2.5, and in Appendix B.3 we derive an open-economy version of Proposition 1. Finally, in Appendix B.4, we provide several further robustness checks for the model-based exercises in Section 4.3 and for the optimal policy analysis in Section 5.

B.1 Baseline model

We characterize equilibria in our baseline model with and without the simplifying assumption of a neutral monetary policy. Consistent with the discussion in Section 2.4, we pay particular attention to the role of adjustment costs in muting pent-up demand effects.

NEUTRAL MONETARY RULE. We proceed exactly as in Section 2.2: first characterizing impulse response paths for arbitrary shock mixtures $\{b_0^a, b_0^s, b_0^d\}$, and then discussing implications for recovery dynamics.

Lemma B.1. *Suppose that monetary policy is neutral. Then the impulse responses of services and durables consumption expenditures to a vector of time-0 shocks $\{b_0^a, b_0^s, b_0^d\}$ satisfy*

$$\widehat{s}_t = \frac{1}{\gamma}(b_0^a + b_0^s)\rho_b^t \quad (\text{B.1})$$

and

$$\widehat{e}_t = \frac{1}{\gamma}(b_0^a + b_0^d)\frac{\theta_b}{\delta} \left(\rho_b^t - (1 - \delta - \theta_d)\frac{\theta_d^t - \rho_b^t}{\theta_d - \rho_b} \right) \quad (\text{B.2})$$

where $\{\theta_d, \theta_b\}$ are functions of model primitives with $\theta_d \in [0, 1)$ and $\theta_b > 0$. The impulse response of aggregate output is thus

$$\widehat{y}_t = \phi\widehat{s}_0\rho_b^t + (1 - \phi)\widehat{e}_0 \left(\rho_b^t - (1 - \delta - \theta_d)\frac{\theta_d^t - \rho_b^t}{\theta_d - \rho_b} \right) \quad (\text{B.3})$$

Lemma B.1 implies that the pent-up demand logic at the heart of the discussion in Section 2.2 remains present in a richer environment with persistent shocks and adjustment costs. To see this, consider first the case of $\rho_b > 0$ but $\kappa = 0$. In that case $\theta_d = 0$, and so the pent-up demand logic is entirely unaffected: the impulse response of services expenditures

decays at a constant rate ρ_b , while the impulse response of durables expenditures is scaled by $\rho_b^t - (1 - \delta)\rho_b^{t-1}$. Thus, while durables expenditures may not literally *overshoot* following sufficiently persistent negative shocks, durables expenditures will still be pushed up relative to expenditures on services. Visually, for negative shocks, the normalized durables spending impulse response always lies above the services impulse response. Second, for $\kappa > 0$, adjustments in the durables stock are slowed down, adding endogenous persistence (i.e., $\theta_d > 0$) that at least partially offsets pent-up demand effects.

We can now as before translate Lemma B.1 into a result relating demand composition and the strength of the recovery.

Proposition B.1. *Suppose that monetary policy is neutral, and consider a vector of time-0 shocks $\{b_0^a, b_0^s, b_0^d\}$ with a services share ω . Then, the normalized cumulative impulse response of aggregate output satisfies*

$$\hat{\mathbf{y}} = \frac{1}{1 - \rho_b} \left[1 - (1 - \omega) \left(1 - \frac{\delta}{1 - \theta_d} \right) \right] \quad (\text{B.4})$$

Proposition B.1 reveals that, in the presence of adjustment costs ($\theta_d > 0$), our conclusions on the effect of demand composition on the strength of the subsequent recovery do not go through automatically — they hold if and only if pent-up demand effects are strong enough, i.e. when $\theta_d < 1 - \delta$. Proposition 2 is the key necessary and sufficient condition ensuring that this is indeed the case.

θ_d vs. $1 - \delta$. Our quantitative explorations in Appendix B.4 and Section 4.3 reveal that, for a wide range of model calibrations, adjustment costs dampen — but do not come close to offsetting — pent-up demand effects. We here provide an analytical argument to rationalize this finding. The key result is the following:

Proposition B.2. *Consider the model of Section 2.4. If $\theta_d = 1 - \delta$, then*

$$\theta_b = \frac{\delta}{\gamma}$$

The durables share after an aggregate demand shock, $\frac{(1-\phi)\tilde{e}_0^a}{\phi\tilde{s}_0^a + (1-\phi)\tilde{e}_0^a}$, is thus equal to $1 - \phi$.

If the adjustment cost κ is large enough to completely offset the pent-up demand mechanism, then durables spending is also not more volatile than services spending, sharply at odds

with empirical evidence.³⁴ Thus, as long as relative spending volatilities are data-consistent, pent-up demand effects invariably dominate the adjustment cost-related slowdown.

ACTIVE MONETARY RULE AND PARTIALLY STICKY PRICES. First, we consider the generalization of Proposition B.1 to an economy with partially sticky prices and a monetary policy rule which does not fix the real rate. For simplicity we here assume that wages are perfectly flexible.

Proposition B.3. *Consider the model of Section 2.1 with arbitrary degrees of price stickiness (governed by ζ_p) but flexible nominal wages ($\phi_w = 0$). Let $\{b_0^a, b_0^s, b_0^d\}$ be a vector of time-0 shocks resulting in services share ω . Then the normalized cumulative impulse response of aggregate output satisfies*

$$\hat{y} = \frac{1}{1 - \rho_b} \left[1 - (1 - \omega) \left(1 - \frac{\delta}{1 - \theta_d} \left(1 + \frac{\phi}{1 - \phi} \theta_s \right) \right) \right]$$

where $\{\theta_d, \theta_s\}$ are functions of model primitives.

As we show in the proof of Proposition B.3, the coefficients $\{\theta_d, \theta_s\}$ govern the response of durables and services consumption $\{\hat{d}_t, \hat{s}_t\}$ to changes in the past durable stock \hat{d}_{t-1} . We then have the following generalization of Proposition 2.

Proposition B.4. *Let \hat{s}^a and \hat{e}^a denote the normalized cumulative impulse responses of services and durables expenditure to a recessionary aggregate demand shock $b_0^a < 0$, defined as in (8), resulting in services share ω^a .*

Then, the normalized cumulative impulse response of aggregate output \hat{y} in Proposition B.3 is increasing in the services share ω if and only if

$$\hat{s}^a > \hat{e}^a \left(1 + \frac{\phi}{1 - \phi} \frac{1}{\omega^a} \theta_s \right)$$

We expect that $\theta_s > 0$, so (10) becomes only necessary and not sufficient. Our numerical simulations in Appendix B.4 (for quadratic stock adjustment costs) and Section 4.3 (for

³⁴The relative unconditional volatility documented in the bottom panel of Figure C.7 suffices as evidence under the assumption that business cycles are largely driven by aggregate demand disturbances. A stronger test looks at relative volatilities *conditional* on a particular aggregate demand shock, e.g. to monetary policy. It is well-known that, even conditional on such shocks, durables spending is much more volatile than non-durables spending (e.g. Christiano et al., 1999).

the impulse response matching exercise) confirm this intuition, but reveal the policy-related offset to be small for standard parameterizations of monetary Taylor-type rules.

B.2 Extensions

This section discusses how our results generalize to the various model extensions discussed in Section 2.5 and for an alternative specification of the aggregate demand shock b_t^a .

MANY SECTORS. In an economy with N sectors and fixed real rates, our results in Section 2.4 and Proposition B.1 apply sector-by-sector.

Proposition B.5. *Consider the extended model with N sectors, and suppose that the monetary rule is neutral. Consider an arbitrary shock mix $\{b_0^a, \{b_0^i\}_{i=1}^N\}$ with sectoral spending shares ω_i . Then*

$$\hat{\mathbf{y}} = \frac{1}{1 - \rho_b} \times \sum_{i=1}^N \omega_i \frac{\delta_i}{1 - \theta_d^i} = \sum_{i=1}^N \omega_i \hat{\mathbf{e}}_i^a \quad (\text{B.5})$$

where θ_d^s and θ_d^d are functions of model primitives.

GENERAL ADJUSTMENT COSTS & STICKY INFORMATION. Our results on the equivalence between nCIR rankings and the effects of spending composition on recovery dynamics go through without change in models with arbitrary adjustment costs — including in particular the flow adjustment costs in Appendix B.4 — as well as sticky information. Formally, in either case, the normalized cumulative impulse response of aggregate output to an arbitrary shock vector $\{b_0^a, b_0^s, b_0^d\}$ still satisfies

$$\hat{\mathbf{y}} = \omega \hat{\mathbf{s}}^a + (1 - \omega) \hat{\mathbf{e}}^a$$

and so $\hat{\mathbf{y}}$ is increasing in ω if and only if

$$\hat{\mathbf{s}}^a > \hat{\mathbf{e}}^a$$

The argument is immediate: neither friction changes the separability of the consumer demand system in (s, d) that lies at the heart of Proposition 2. Thus, relative to the results in Proposition B.1, we only lose the ability to characterize $\hat{\mathbf{s}}^a$ and $\hat{\mathbf{e}}^a$.

NON-SEPARABILITY. Propositions B.3 and B.4 extend without change to the case where preferences are non-separable as in Appendix A.2 (with inverse elasticity of substitution ζ). The proofs of both propositions in this extended case are identical to the ones in our baseline model and thus we omit them.

We pay particular attention to the special case of net substitutability ($\zeta < \gamma$) together with the assumption of a neutral monetary policy. In that case, we expect the coefficient θ_s in the characterizations of Propositions B.3 and B.4 to be strictly negative — intuitively, an elevated lagged durables stock should depress current services spending. Formally, combining equations (A.12) - (A.14), and solving for θ_s, θ_d , we obtain

$$\theta_s = -\theta_d \frac{[1 - \beta(1 - \delta)]^{\frac{1}{\delta}}(1 - \phi)}{\gamma\phi + \zeta[1 - \beta(1 - \delta)]^{\frac{1}{\delta}}(1 - \phi)}(\gamma - \zeta)$$

$$0 = \left((\gamma - \zeta) \frac{[1 - \beta(1 - \delta)]^{\frac{1}{\delta}}(1 - \phi)}{\phi\gamma + [1 - \beta(1 - \delta)]^{\frac{1}{\delta}}(1 - \phi)\zeta} + 1 \right) \zeta\theta_d + \frac{\kappa}{1 - \beta(1 - \delta)}(\theta_d - 1)(1 - \beta\theta_d).$$

One can readily see that, if $\zeta < \gamma$, then $1 > \theta_d > 0$ and θ_s is indeed negative. Therefore, our conclusions are in fact strengthened: (10) now is not necessary, but it remains sufficient.

INCOMPLETE MARKETS. In the simple spender-saver extension of our baseline model, our main results go through unchanged, as impulse responses are merely scaled up or down at all horizons.

Proposition B.6. *Consider the extended model with hand-to-mouth consumers, and suppose that monetary policy is neutral. Then:*

1. *If $\eta = 1$, all impulse responses are exactly as in the model with $\mu = 0$.*
2. *For arbitrary η , all normalized impulse responses are exactly as in the model with $\mu = 0$.*

Thus, in both cases, Proposition B.1 and Proposition 2 apply without change.

SUPPLY SHOCKS. In the model of Section 2.4, our results on demand recessions apply without change to our particular notion of supply shock-induced recessions.

Proposition B.7. *Consider the extended model with supply shocks. Then Propositions 2 and B.1 apply without change to a vector of time-0 supply shocks $\{z_0^a, z_0^s, z_0^d\}$.*

ALTERNATIVE SPECIFICATION FOR b_t^a . We note that many — but not all — of our main results extend to the alternative preference specification (A.19).

Consider first the equilibrium characterizations and recession recovery results in Sections 2.2 to 2.4. Comparing (A.20)-(A.21) to (A.7)-(A.8), and using the fact that all shocks follow AR(1) processes with common persistence ρ_b , we can conclude that the sectoral expenditure impulse responses \hat{s}_t and \hat{e}_t to any shock tuple $\{b_0^a, b_0^d, b_0^s\}$ are unaffected *up to scale*. Intuitively, the new preference specification is isomorphic to a rescaling of the shocks b_t^a and b_t^d in (A.8); since Propositions 1 and B.1 condition on sectoral spending declines, they (as well as Proposition 2) are entirely unaffected by the change in preferences.³⁵ In results available upon request, we also show that the conclusions from our quantitative analysis in Section 4.3 are barely affected by a change of the preference specification to (A.19). By the preceding discussion, the quantitative results are *exactly* unaffected at the boundary of the parameter range with fixed prices; they then only change slightly as we increase the degree of price flexibility.

Second, consider our empirical analysis in Section 3, and in particular the equivalence to monetary shocks. Given our new preference specification, Proposition 3 now fails: assuming for simplicity that monetary policy is such that the shock m_t equals the equilibrium response of (expected) real interest rates (e.g., because of perfectly fixed prices), straightforward algebra reveals that, with preferences as in (A.19), a monetary shock m_t with persistence ρ_m is now equivalent (up to scale) to a mixture of sectoral demand shocks

$$b_t^s + \frac{1 - \beta(1 - \delta)\rho_b}{1 - \beta(1 - \delta)} b_t^d \tag{B.6}$$

The weight on b_t^d is simply equal to our preference scaling parameter α , exactly as expected.

Third, the discussion of optimal policy in Section 5 is necessarily sensitive to the preference specification. If the economy is only buffeted by the mixture of sectoral demand shocks in (B.6), then our result on the independence of optimal monetary policy with respect to long-run expenditure shares continues to apply: such shock combinations are simply shocks to the equilibrium real rate of interest, and optimal policy tracks that interest rate. Thus, even though the transmission of monetary policy is affected, the monetary authority again faces no intertemporal trade-off in optimal stabilization.

³⁵The analogous expressions for Lemmas 1 and B.1 change, however, as an aggregate demand shock b_t^a is now a different weighted average of pure sectoral shocks.

B.3 Model with many regions

We consider an economy consisting of many regions, all part of a common monetary union. Each region produces non-tradable services and tradable durables. Production uses labor, assumed to be immobile and fully sector-specific. Total output in region i is given as $y_{it} = s_{it} + e_{it}$, where s_{it} is services output and e_{it} is durables output. We show here that a version of Proposition 1 applies to regional production y_{it} in this multi-region economy.

As in our stripped-down model of Section 2.2, we assume that the representative household in region i has separable preferences over local services as well as durables produced by all other regions. In the symmetric steady state, all regions produce the same amounts of services ($\bar{s} = \phi$) and durables ($\delta\bar{d} = 1 - \phi$). We assume that the preferences of households in each region i are subject to demand shocks for all goods (b_{it}^a), for local services (b_{it}^s), and for the durables of each region j (b_{ijt}^d). Now let p_{it}^s denote the real relative price of local services, and similarly let p_{jt}^d denote the real relative price of durables from region j . The first-order conditions for region i 's demand of services and durables produced by any region j are given by

$$\begin{aligned} -\gamma\widehat{s}_{it} &= \widehat{\lambda}_{it} + \widehat{p}_{it}^s - (b_{it}^a + b_{it}^s) \\ -\gamma\widehat{d}_{ijt} &= \frac{1}{1 - \beta(1 - \delta)}(\widehat{\lambda}_{it} + \widehat{p}_{jt}^d - (b_{it}^a + b_{ijt}^d)) - \frac{\beta(1 - \delta)}{1 - \beta(1 - \delta)}(\widehat{\lambda}_{it+1} + \widehat{p}_{jt+1}^d - (b_{it+1}^a + b_{ijt+1}^d)) \end{aligned}$$

where λ_{it} is the local marginal utility of income.

In keeping with our stripped-down model of Section 2.2, we now furthermore assume that asset markets are complete and the prices of all goods are fixed, implying that we have $\widehat{\lambda}_{it} = 0$. Considering again the case of perfectly transitory shocks, and aggregating over all regions, we obtain the total demand for durables produced in region i , \widehat{d}_{it} , as

$$-\gamma\widehat{d}_{it} = \frac{1}{1 - \beta(1 - \delta)}(-(b_{it}^a + b_{it}^d)) - \frac{\beta(1 - \delta)}{1 - \beta(1 - \delta)}(-(b_{it+1}^a + b_{it+1}^d))$$

where $b_t^a \equiv \int b_{jt}^a dj$ and $b_{it}^d \equiv \int b_{ijt}^d dj$ are aggregated demand shocks. Proceeding exactly as in the proof of Lemma 1, it follows that durables output from region i in response to time-0 shocks satisfies

$$\begin{aligned} -\gamma\delta\widehat{e}_{i0} &= -\frac{1}{1 - \beta(1 - \delta)}(b_0^a + b_{i0}^d) \\ \widehat{e}_{i1} &= -(1 - \delta)\widehat{e}_{i0} \end{aligned}$$

and $\widehat{e}_{it} = 0$ otherwise. Analogously, market-clearing in the (non-tradable) services sector implies that services output in region i satisfies

$$-\gamma \widehat{s}_{i0} = -(b_{i0}^a + b_{i0}^s)$$

and $\widehat{s}_{it} = 0$ otherwise.

Computing total output from region i , we obtain

$$\begin{aligned} \widehat{y}_{i0} &= \phi \frac{1}{\gamma} (b_{i0}^a + b_{i0}^s) + (1 - \phi) \frac{1}{1 - \beta(1 - \delta)} \frac{1}{\gamma \delta} (b_0^a + b_{i0}^d) \\ \widehat{y}_{i1} &= -(1 - \delta)(1 - \phi) \frac{1}{1 - \beta(1 - \delta)} \frac{1}{\gamma \delta} (b_0^a + b_{i0}^d) \\ \widehat{y}_{it} &= 0 \quad \forall t > 1 \end{aligned}$$

Regional output nCIRs as well as the recession's output share of services are thus given as

$$\begin{aligned} \widehat{\mathbf{y}}_i &= 1 - (1 - \delta) \frac{(1 - \phi) \frac{1}{1 - \beta(1 - \delta)} \frac{1}{\delta} (b_0^a + b_{i0}^d)}{\phi (b_{i0}^a + b_{i0}^s) + (1 - \phi) \frac{1}{1 - \beta(1 - \delta)} \frac{1}{\delta} (b_0^a + b_{i0}^d)} \\ \omega_i &= \frac{\phi (b_{i0}^a + b_{i0}^s)}{\phi (b_{i0}^a + b_{i0}^s) + (1 - \phi) \frac{1}{1 - \beta(1 - \delta)} \frac{1}{\delta} (b_0^a + b_{i0}^d)} \end{aligned}$$

It follows that $\widehat{\mathbf{y}}_i = 1 - (1 - \omega_i)(1 - \delta)$ — the same as in Proposition 1, but where ω_i is now interpreted as the share of services in the change of total output (rather than consumption).

B.4 Quantitative analysis

In Section 4.3 we study the effect of pent-up demand on expected recovery dynamics in a structural model estimated via impulse response matching, and in particular focus on the role of adjustment costs and price stickiness. We here provide two additional exercises. First, departing again from our estimated model, we investigate the role of shock persistence (ρ_b) and the monetary authority's policy rule (ϕ_π). Second, we provide an analogous analysis for the baseline environment of Section 2.1, without the added complications of flow adjustment costs and sticky information.

Finally, in the last part of this appendix, we use numerical simulations to investigate the generality of Proposition 6.

SHOCK PERSISTENCE AND POLICY RULE. Figure B.1 reports nCIR ratios in the estimated model, varying shock persistence and the Taylor rule coefficient.




Figure B.1: Left panel: Percentage gap between the nCIR to an aggregate demand shock b_0^a in an economy with the U.S. vs. Canada long-run expenditure shares, as a function of shock persistence (x -axis) and the Taylor rule coefficient (y -axis). Right panel: Percentage gap between the nCIR to demand shocks (b_0^a, b_0^d, b_0^s) inducing a composition of expenditure changes on impact as in a services-led vs. an average U.S. recession, again as a function of shock persistence (x -axis) and the Taylor rule coefficient (y -axis). The red cross in both figures indicates our preferred parameterization. All other parameters are set as in Table 4.1.

The key take-away is that pent-up demand effects remain quite strong throughout; in particular, even for quite transitory shocks and aggressive monetary responses, the estimated nCIR differences are still substantial. Thus, in a model consistent with the evidence in Section 3.1, the conclusion that sectoral incidence matters for recovery dynamics remains inescapable.

We provide another illustration – geared towards the ongoing COVID-19 pandemic – in Figure B.2. This figure repeats the analysis in the right panel of Figure 6, but now for transitory shocks ($\rho_b = 0.2$), with the persistence chosen to match the time it took U.S. total consumption expenditures to reach their pre-crisis level. In all cases, because the shock is transitory, the recovery is rather strong. Crucially, however, the recovery is even stronger for more durables-led contractions. Thus, applied to COVID-19, the pent-up demand logic suggests that if such a transitory shock had been concentrated on durables spending instead, the recovery would have been *even stronger*.

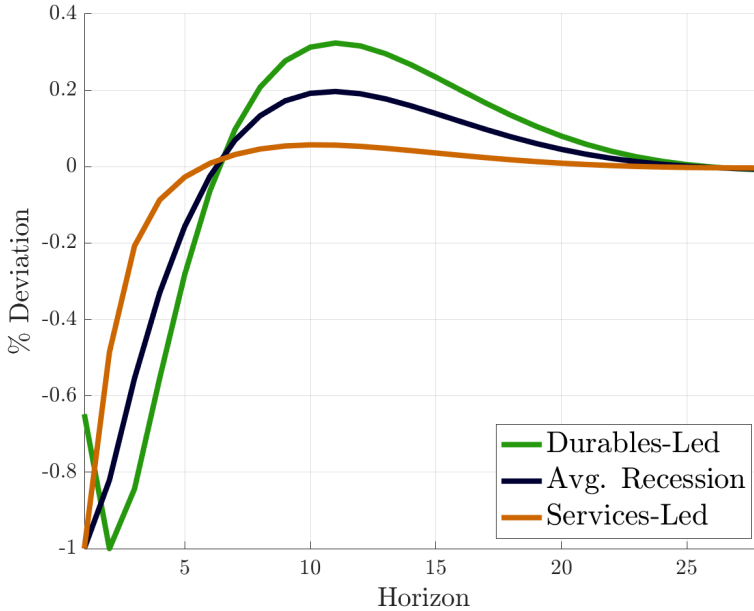


Figure B.2: See the caption for the right panel of Figure 6. The only difference is that now we consider shocks with persistence $\rho_b = 0.2$.

A SIMPLER MODEL. The principal appeal of our model in Section 4.3 is that it closely matches our empirically estimated impulse responses, and so mirrors our preferred approach to measurement — the semi-structural shift-share. While flow adjustment costs and sticky information matter for the impulse response fit, however, they play little role in our quantitative conclusions regarding ratios of nCIRs. We provide a visual illustration in Figure B.3, which presents the analogue of Figure 7 for a model calibrated exactly as in Table 4.1 but with $\theta = 0$ (no informational frictions), $\kappa_e = 0$ (no flow adjustment costs), and $\kappa \geq 0$ (adding quadratic stock adjustment costs). The preferred parameter estimate of the adjustment cost is $\kappa = 0.6$, which ensures durables and services nCIRs matching those implied by our empirically identified monetary policy shock impulse responses.³⁶

Figure B.3 paints almost exactly the same picture as Figure 7. The intuition is straightforward: changing the nature of adjustment costs and informational frictions changes the *shape* of impulse responses, but for Figure B.3 all that matters is the *level* of shock-specific nCIRs. Pent-up demand effects remain present and strong across a wide range of plausible model parameterizations, and so the percentage gaps in nCIRs across shocks are large

³⁶We have verified not only that the nCIRs match, but also that the corresponding conditional volatilities (i.e., without the trough normalization) are in line with the empirical evidence.




Figure B.3: Left panel: Percentage gap between the nCIR to an aggregate demand shock b_0^a in an economy with the U.S. vs. Canada long-run expenditure shares, as a function of adjustment costs (x -axis) and the NKPC slope (y -axis). Right panel: Percentage gap between the nCIR to demand shocks (b_0^a, b_0^d, b_0^s) inducing a composition of expenditure changes on impact as in a COVID-19 vs. an average U.S. recession, again as a function of adjustment costs (x -axis) and the NKPC slope (y -axis). The red cross in both figures indicates our preferred parameterization. All other parameters are set as in Table 4.1.

throughout — irrespective of whether impulse responses actually follow hump shapes (as in Figure 5) or not (as is the case here).

ROBUSTNESS OF PROPOSITION 6. While Proposition 6 is stated for the special case of no adjustment costs and transitory shocks, the result is numerically robust to changing both κ and ρ_b . In particular we have verified that, on a fine grid of $\kappa \in [0, 1]$ (as in Figure B.3) and $\rho_b \in [0, 0.95]$, it is always true that

$$\frac{\sum_{t=0}^H \widehat{r}_t(b_0^s)}{\max_t |\widehat{r}_t(b_0^s)|} > \frac{\sum_{t=0}^H \widehat{r}_t(b_0^d)}{\max_t |\widehat{r}_t(b_0^d)|}$$

In words, the cumulative real interest rate impulse response to a services demand shock (normalized by the peak response) is always strictly larger than the corresponding normalized cumulative impulse response for a durables demand shock, consistent with the basic message of Proposition 6.

C Empirical appendix

This appendix provides further details for the empirical exercises in Section 3 as well as our measurement of expenditure share variation for Section 4.1.

C.1 Monetary policy

We estimate recursive VARs containing a sectoral measure of consumption, aggregate consumption, aggregate GDP (all real), the GDP deflator, and the federal funds rate, in this order. The sectoral consumption series are taken straight from NIPA tables (quarterly real series), while all other series are taken from the St. Louis Fed’s FRED database.

All of our VARs are estimated on a quarterly sample from 1960:Q1 — 2007:Q4, and contain four lags, a constant and a linear time trend. We furthermore impose the usual uniform-normal-inverse-Wishart prior over the orthogonal reduced-form parameterization (Arias et al., 2018). Throughout, we display confidence bands constructed through 10,000 draws from the model’s posterior. Finally, to construct a posterior credible set for the nCIR difference of durables- and services-led recession, we estimate a single VAR containing our three baseline consumption series (i.e., durables, non-durables, and services), compute sectoral nCIRs for each posterior draw and take ratios (to get (14)), and then report percentiles of that distribution.

ROBUSTNESS: HOUSING. The baseline NIPA measure of services also contains housing and utilities. Since imputed rent arguably has a durable component we here repeat our analysis with services less housing replacing our baseline services measure. To construct this cleaned measure we proceed as follows: We first collect nominal expenditure and price series for total services and for the housing/utility component, both again from the NIPA tables. We then construct the growth rate of the non-housing services basket as

$$\Delta p_{-H,t} = \frac{\Delta p_t - s_{H,t-1} \Delta p_{H,t}}{1 - s_{H,t-1}}$$

where Δp_t is the growth rate of the services price index, $\Delta p_{H,t}$ is the growth rate of the services housing/utilities price index, and $s_{H,t-1}$ is the $t - 1$ share of housing/utility expenditures in the total services basket. Given this price series for the non-housing services basket, we then deflate the nominal expenditure series for services less housing.

Results are reported in Figure C.1. The first two panels are the same as in Figure 2 and




Figure C.1: Quarterly impulse responses to a recursively identified monetary policy shock (as in Christiano et al. (1999)) by consumption spending category (with services excluding housing & utilities), all normalized to drop by -1% at the trough. The solid blue line is the posterior mean, while the shaded areas indicate 16th and 84th percentiles of the posterior distribution, respectively.

just repeated here for convenience. The third panel reveals that housing and utility services play a negligible role in our conclusions. This finding is of course not surprising given the finer sectoral decomposition shown in Table 3.1.

C.2 Uncertainty

Our analysis of uncertainty shocks closely follows Basu & Bundick (2017). We estimate recursive VARs in the VIX as a measure of uncertainty shocks, real GDP, the GDP deflator, and real measures of sectoral consumption (durables, non-durables, services). By the results in Plagborg-Møller & Wolf (2021), this specification is asymptotically equivalent to a local projection on innovations in the VIX. All series are taken from the replication files for Basu & Bundick (2017). We estimate the recursive VAR on a quarterly sample from 1986:Q1 — 2014:Q4, and include four lags.³⁷ As before we include a constant and a linear time trend, impose a uniform-normal-inverse-Wishart prior over the orthogonal reduced-form parameterization of the VAR, and draw 10,000 times from the model’s posterior.

Figure C.2 shows the sectoral consumption impulse responses, all scaled to show a peak drop in consumption of -1 per cent. As predicted by theory and as in our application to monetary policy transmission, we find that durables expenditures overshoot and then return to baseline, while non-durables and services expenditure return to baseline from below.

³⁷The results are unaffected with longer lag lengths, which reduce precision but ensure accurate projection at longer horizons.




Figure C.2: Quarterly impulse responses to an uncertainty shock (à la Basu & Bundick (2017)) by consumption spending category, all normalized to drop by -1% at the trough. The solid blue line is the posterior mean, while the shaded areas indicate 16th and 84th percentiles of the posterior distribution, respectively.

Figure C.3 uses these estimates to construct a shift-share evaluation of the two counterfactuals studied in Section 4, analogous to the analysis in Section 4.2. The results agree closely with our baseline estimates using monetary policy shocks.

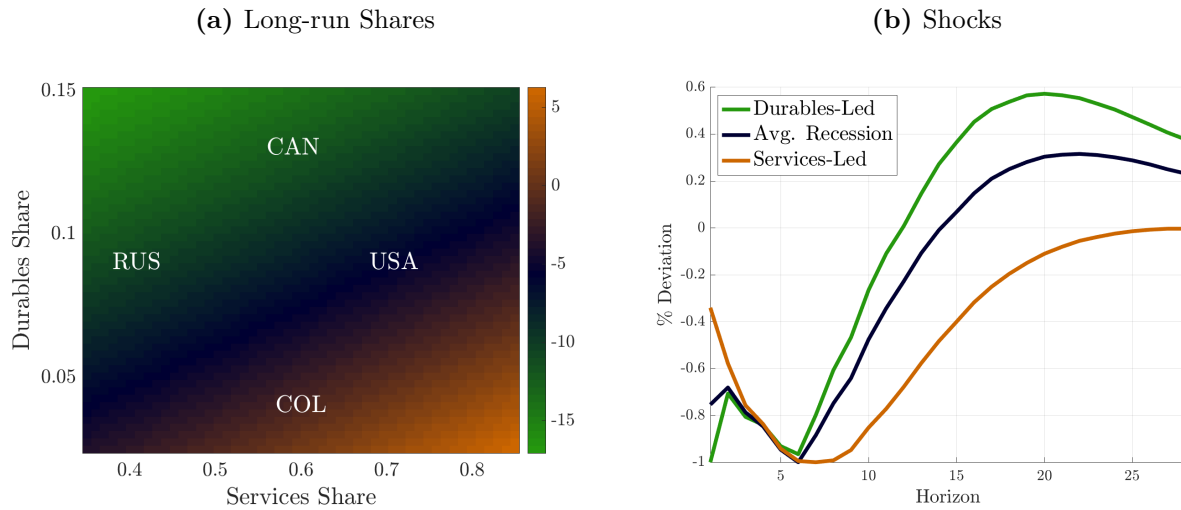


Figure C.3: Left panel: nCIR to an aggregate demand shock b_0^a as a function of long-run expenditure shares and relative to the U.S. Right panel: Impulse response of total consumption to shock combinations giving a durables/non-durables/services trough spending composition of: (i) [100%, 50%, -50%] for a durables-led recession, (ii) [62%, 45%, -7%] for the average U.S. recession, and (iii) [12.5%, 12.5%, 75%] for a services-led recession. All are normalized to lead to a peak-to-trough consumption contraction of -1 per cent. For both panels we use the posterior mode point estimates from Figure C.2.

C.3 Oil

For our analysis of oil price shocks we take the shock series from Hamilton (2003), and order it first in a recursive VAR containing the shock measure, real GDP, the GDP deflator, aggregate consumption, and sectoral measures of consumption. The model specification is largely as before: We estimate the VAR on a sample from 1970:Q1 — 2006:Q4 (dictated by data constraints), include 8 lags to ensure for accurate projection at long horizons, allow for a constant and a linear time trend, and use Bayesian estimation methods.

Since the oil price shock directly affects relative sectoral prices at a level finer than the durable/non-durable distinction considered in most the paper, we consider several granular measures of sectoral consumption. The results from a subset of our experiments are reported in Figure C.4. Durables show the expected overshoot. At a finer sectoral level, we see that expenditures on gas and transport show a similar overshoot. Intuitively, transport — in particular holiday travel — is arguably a memory good and so behaves like a durable good, explaining the overshoot in transport itself as well as the complementary gas expenditure (Hai et al., 2013). In contrast, expenditure on food and financial services decline in the initial recession, but then only recover gradually and without much of an overshoot.

C.4 Reduced-form dynamics

We estimate a reduced-form autoregressive representation for our three main sectoral consumption series (durables, non-durables, services) on the largest possible sample, from 1960:Q1 — 2019:Q4. To flexibly capture general Wold dynamics in each individual series we include six lags, with results largely unchanged for even more flexible lag specifications. We then compute nCIRs of total consumption to each of the three reduced-form Wold innovations, with the impact consumption response normalized to 1.³⁸

Our main conclusion is that innovations in non-durables and services spending are much more persistent than innovations in durables spending, thus yielding large differences in the implied nCIRs. While not tied to any particular structural shock interpretation, this reduced-form evidence is also in line with the predictions of our basic theory.

³⁸For these computations, we construct aggregate consumption as a weighted average of the sectoral series, with weights of 10 per cent for durables, 65 per cent for services, and 25 for non-durables. These weights are consistent with averages in the NIPA tables over the sample period.

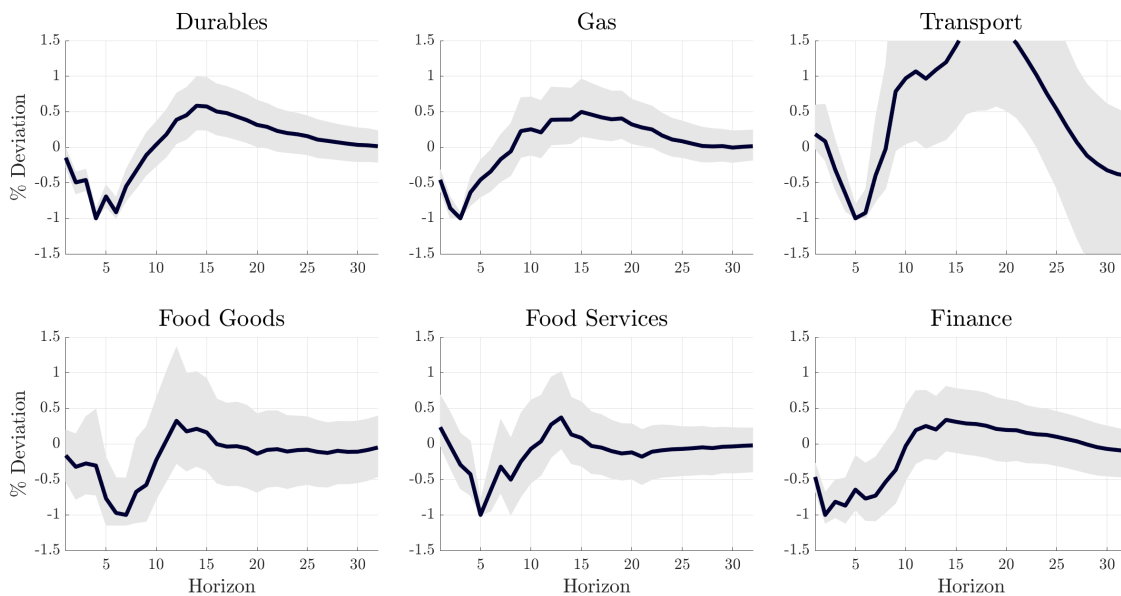


Figure C.4: Quarterly impulse responses to an oil shock (à la Hamilton (2003)) by consumption spending category, all normalized to drop by -1% at the trough. The solid blue line is the posterior mean, while the shaded areas indicate 16th and 84th percentiles of the posterior distribution, respectively.

C.5 Evidence from U.S. regional business cycles

This section complements the discussion in Section 3.3. We first give further details on our data construction for Figure 3, then present results for alternative definitions of the share ω , and finally provide supplementary evidence for other recessions.

C.5.1 Variables construction

The data comes from the BEA, Table SQGDP11. For each state, we are given the quarter-to-quarter percent change in Private Industry output, as well as the contribution of each sector (e.g., durables manufacturing) to that total change. We first add-up the contributions of all sectors that we include in our definition of $1 - \omega$. For example, for our benchmark in Figure 3, we compute the total sectoral output contribution by adding up the contributions of durables manufacturing, non-durables manufacturing and wholesale retail. We then convert all percent changes to log-changes, as our definitions of \hat{y} and ω require. For example, given the private industry output percent changes $\frac{y_t - y_{t-1}}{y_{t-1}} \times 100$, we compute $\log\left(\frac{y_t}{y_{t-1}}\right) \times 100 = \log\left(1 + \frac{y_t - y_{t-1}}{y_{t-1}}\right) \times 100$.

Next, for each episode, we identify the state-specific peak and trough of private industry output. To obtain the log-change from the peak (P), we cumulate the quarter-to-quarter log-changes in total sectoral output contributions and in private industry output. We call this $\Delta \log(x_{t,P})$ and $\Delta \log(y_{t,P})$, respectively. Given trough T , we then compute the recession share ω as $1 - \Delta \log(x_{T,P}) / \Delta \log(y_{T,P})$. Finally, we compute the nCIR of output \hat{y} up to four quarters after the trough as $1 + \sum_{j=1}^4 \Delta \log(y_{T+j,P}) / \Delta \log(y_{T,P})$.

C.5.2 Alternative empirical definitions of the share ω

Figure C.5 shows that our results from Section 3.3 are robust to defining ω using the Manufacturing share of the recession or the Manufacturing plus Retail share.

C.5.3 Evidence for other recessions

Our analysis requires two inputs: regional recession output nCIRs as a measure of recovery strength, and the composition of output changes. Unfortunately, these are not available for recessions before the Great Recession, so in the main text we focus on the Great Recession and the COVID-19 pandemic. We here use alternative data sources to extend the analysis further into the past.

Our analysis relies on regional employment data (from the BLS) to compute recession nCIRs. We define state-specific recessions following Olney & Pacitti (2017), and the state-specific peak as the month where state employment attains a maximum within a window that starts twelve months before and extends to twelve months after the NBER business cycle peak dates. We drop the 1982 recession because by then many states still had not recovered from the 1980 recession. Then, we define the state-specific trough as the month with the lowest employment in a window that starts at the state-specific peak and extends until the next state-specific peak. Finally, we calculate the nCIR as the cumulative sum of the percent change in employment from the state-specific peak for the twelve months following the state-specific trough, and normalize it by the change from peak to trough. In both figures we drop observations where the nCIR was below -20. To measure recession demand composition, we use BEA data to compute the state-specific share of non-tradable services in GDP in the year of the recession. More precisely, following Olney & Pacitti (2017), we define this measure as total services minus accommodations and finance. For recessions prior to 1997 that use SIC industry codes, we subtract the categories ‘Hotels and other Lodging Places’ (major group 70), ‘Finance, Insurance and Real Estate’ (division H), and add ‘Real Estate’

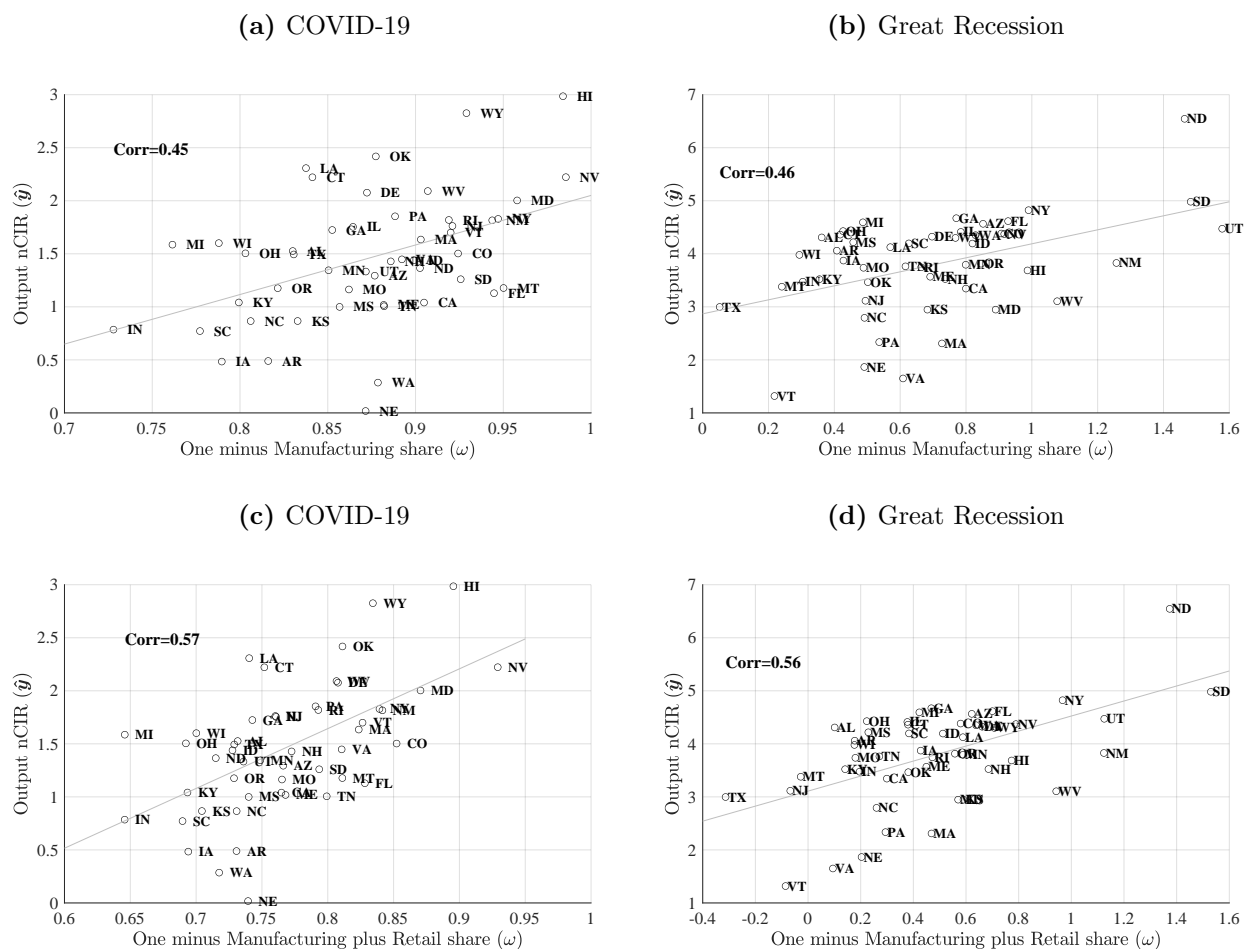



Figure C.5: Recovery strength and recession sectoral composition. The nCIR of output (y) is the cumulative percent change in private industry real GDP from peak to four quarters after the trough, normalized by the peak-to-trough change. The share ($1-\omega$) is the ratio of the percent change in manufacturing output only or manufacturing plus retail output over the percent change in private industry output from peak to trough. Data is from the BEA, Table SQGDP11.

Figure C.6: Regional nCIR of employment vs. GDP share in nontradable services across recessions



Note: The normalized cumulative impulse response (nCIR) of employment one year out is measured as the cumulated percent change in employment from peak divided by the percent change in employment from peak to trough. Observations where the nCIR was below 20 were dropped. The x-axes corresponds to the share of nontradable services in GDP the year of the recession.

(major group 65) back in. For the recessions that fall under the NAICS system, we subtract ‘accommodation’ (subsector 721) and ‘finance and insurance’ (subsector 521).

The left panel of Figure C.6 shows how the our recession employment nCIRs vary across U.S. states with the share on non-tradable services in state GDP. For the past six recessions (excluding 1982), the nCIR one year out from the state-specific trough has been larger in states that have a higher share of non-tradable services. The right panel differences-out the median trends across states. The slope coefficients from the OLS regression lines are 29.328 (4.275) and 15.452 (5.447), and are significant at a 1 percent level. The evidence in Figure C.6 is consistent with our main hypothesis if we interpret each variable as corresponding to the theoretical ω and \hat{y} .

C.6 Expenditure shares across countries and U.S. recessions

Figure C.7 shows expenditure shares across countries in 2017 (left panel) and across post-war U.S. recessions (right panel).

The country-specific expenditure shares in 2017 come directly from stats.oecd.org. For U.S recessions, the contribution of category X to the change in real personal consumption

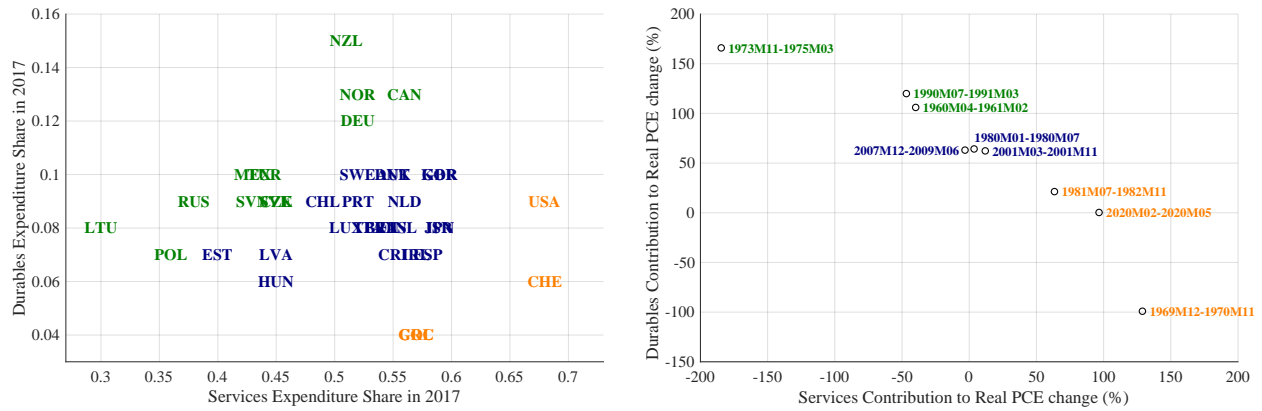


Figure C.7: Left panel: Durables and services expenditure shares across OECD countries in 2017. Source: stats.oecd.org. Right panel: Contributions of durables and services expenditures changes to real personal consumption expenditures (PCE) changes (peak-to-trough, monthly) in U.S. recessions. Source: bea.gov

Y , c_x , is calculated using the Tornqvist formula approximation:

$$c_x = \theta_t \frac{\log(X_t/X_{t-1})}{\log(Y_t/Y_{t-1})},$$

where X and Y are measured by the Fisher Quantity Index (U.S. BEA, “Table 2.4.3U Real Personal Consumption Expenditures”), and θ_t is the average of the ratio of nominal X to nominal Y (U.S. BEA, “Table 2.4.5U Personal Consumption Expenditures”) in periods t and $t - 1$ (trough and peak months, respectively). For the ‘Average of the past U.S. recessions’ mentioned in the text, we take the simple average of the change from peak to trough following NBER business cycle reference dates starting in April 1960. For the ‘COVID-19 US recession’ we set February 2020 as the peak and May 2020 as the trough (at the time of writing).

D Proofs

D.1 Proof of Lemmas 1 and B.1

With a fixed (expected) real rate of interest and equilibrium selection from Section 2.2, it follows that $\widehat{\lambda}_t = 0$ for all t , and so we can solve for the impulse responses of services and durables consumption by solving the system (A.3), (A.5) and (A.6). Durable expenditures are then given by

$$\widehat{e}_t = \frac{1}{\delta}(\widehat{d}_t - (1 - \delta)\widehat{d}_{t-1}) \quad (\text{D.1})$$

and aggregate output in equilibrium is fully demand-determined and given by

$$\widehat{y}_t = \phi\widehat{s}_t + (1 - \phi)\widehat{e}_t \quad (\text{D.2})$$

Since all shocks have common persistence ρ_b , we use the method of undetermined coefficients to obtain the dynamic responses

$$\widehat{s}_t = \frac{1}{\gamma}(b_t^a + b_t^s) \quad (\text{D.3})$$

$$\widehat{d}_t = \frac{1}{\gamma}\theta_b(b_t^a + b_t^d) + \theta_d\widehat{d}_{t-1} \quad (\text{D.4})$$

where $\theta_d \in [0, 1)$ is the smallest solution to

$$0 = \beta\kappa(\theta_d)^2 - ((1 + \beta)\kappa + \gamma(1 - \beta(1 - \delta)))\theta_d + \kappa$$

and

$$\theta_b = \frac{1 - \rho_b\beta(1 - \delta)}{1 - \beta(1 - \delta) + \frac{\kappa}{\gamma}(1 + \beta(1 - \rho_b - \theta_d))}$$

Using that $b_t^a + b_t^s = (b_0^a + b_0^s)\rho_b^t$, the above directly implies the dynamic response for \widehat{s}_t in Lemma B.1. Moreover, iterating \widehat{d}_t forward, we obtain

$$\widehat{d}_t = \frac{1}{\gamma}(b_0^a + b_0^d)\theta_b \sum_{j=0}^t (\theta_d)^{t-j} (\rho_b)^j = \frac{1}{\gamma}(b_0^a + b_0^d)\theta_b \frac{(\theta_d)^{t+1} - (\rho_b)^{t+1}}{\theta_d - \rho_b} \quad (\text{D.5})$$

and thus we obtain the dynamic response for \widehat{e}_t in Lemma B.1

$$\widehat{e}_t = \frac{1}{\gamma}(b_0^a + b_0^d) \frac{\theta_b}{\delta} \left(\rho_b^t - (1 - \delta - \theta_d) \frac{\theta_d^t - \rho_b^t}{\theta_d - \rho_b} \right) \quad (\text{D.6})$$

The dynamic response of \widehat{y}_t follow directly from replacing the above in (D.2). Finally, the responses in Lemma 1 are the special case when $\kappa = \rho_b = 0$, so that $\theta_d = 0$ and $\theta_b = \frac{1}{1-\beta(1-\delta)}$. \square

D.2 Proof of Propositions 1 and B.1

From the expression for output in (D.2) and that $\omega \equiv \frac{\phi \widehat{s}_0}{\widehat{y}_0}$, we immediately obtain that

$$\widehat{y} \equiv \frac{\sum_{t=0}^{\infty} \widehat{y}_t}{\widehat{y}_0} = \omega \frac{\sum_{t=0}^{\infty} \widehat{s}_t}{\widehat{s}_0} + (1 - \omega) \frac{\sum_{t=0}^{\infty} \widehat{e}_t}{\widehat{e}_0}. \quad (\text{D.7})$$

Using the dynamic responses in Lemma B.1, we have that

$$\frac{\sum_{t=0}^{\infty} \widehat{s}_t}{\widehat{s}_0} = \frac{1}{1 - \rho_b} \quad (\text{D.8})$$

$$\frac{\sum_{t=0}^{\infty} \widehat{e}_t}{\widehat{e}_0} = \frac{\delta}{1 - \theta_d} \frac{1}{1 - \rho_b}. \quad (\text{D.9})$$

Thus, replacing above, we obtain the expression for \widehat{y} in equation (B.4) of Proposition B.1

$$\widehat{y} = \frac{1}{1 - \rho_b} \left[1 - (1 - \omega) \left(1 - \frac{\delta}{1 - \theta_d} \right) \right].$$

Finally, in the special case of $\kappa = \rho_b = 0$ (and thus $\theta_d = 0$) we immediately obtain equation (9) in Proposition 1. \square

D.3 Proof of Proposition 2

First, note that \widehat{y} in equation (B.4) of Proposition B.1 is decreasing in ω if and only if $\delta < 1 - \theta_d$.

Second, using the dynamic responses in Lemma B.1, we have that

$$\widehat{s}^a \equiv \frac{\sum_{t=0}^{\infty} \widehat{s}_t}{\widehat{s}_0} \Big|_{b_0^a \neq 0, b_0^s = b_0^d = 0} = \frac{1}{1 - \rho_b}$$

$$\widehat{\mathbf{e}}^a \equiv \frac{\sum_{t=0}^{\infty} \widehat{e}_t}{\widehat{e}_0} \Big|_{b_0^a \neq 0, b_0^s = b_0^d = 0} = \frac{\delta}{1 - \theta_d} \frac{1}{1 - \rho_b}.$$

Then, we can readily see that $\delta < 1 - \theta_d$ if and only if $\widehat{\mathbf{s}}^a > \widehat{\mathbf{e}}^a$. □

D.4 Proof of Proposition 3

Plugging the policy rule (3) into the bond FOC (A.4), we get

$$\widehat{\lambda}_t = \phi_\pi \widehat{\pi}_t + m_t - \mathbb{E}_t [\widehat{\pi}_{t+1}] + \mathbb{E}_t [\widehat{\lambda}_{t+1}]$$

Now let $\widehat{\lambda}_t \equiv \widehat{\lambda}_t - \frac{1}{1 - \rho_m} m_t$. Plugging this in, the bond FOC becomes

$$\widehat{\lambda}_t = \phi_\pi \widehat{\pi}_t - \mathbb{E}_t [\widehat{\pi}_{t+1}] + \mathbb{E}_t [\widehat{\lambda}_{t+1}]$$

Similarly, the services and durables FOCs become

$$\begin{aligned} -\gamma \widehat{s}_t &= \widehat{\lambda}_t + \frac{1}{1 - \rho_m} m_t \\ -\gamma \widehat{d}_t &= \frac{1}{1 - \beta(1 - \delta)} \left[\widehat{\lambda}_t + \frac{1}{1 - \rho_m} m_t + \kappa(\widehat{d}_t - \widehat{d}_{t-1}) \right] \\ &\quad - \frac{\beta(1 - \delta)}{1 - \beta(1 - \delta)} \mathbb{E}_t \left[\widehat{\lambda}_{t+1} + \frac{1}{1 - \rho_m} m_{t+1} + \frac{\kappa}{1 - \delta} (\widehat{d}_{t+1} - \widehat{d}_t) \right] \end{aligned}$$

and the wage-NKPC becomes

$$\widehat{\pi}_t^w = \frac{(1 - \beta\phi_w)(1 - \phi_w)}{\phi_w(\frac{\varepsilon_w}{\varphi} + 1)} \left[\frac{1}{\varphi} \widehat{\ell}_t - \left(\widehat{w}_t + \widehat{\lambda}_t + \frac{1}{1 - \rho_m} m_t \right) \right] + \beta \mathbb{E}_t [\widehat{\pi}_{t+1}^w]$$

Therefore, for shocks m_t with persistence ρ_m and size m_0 , these equations are identical to those in a model with aggregate demand shocks b_t^a with persistence $\rho_b = \rho_m$ and size $b_0^a = -\frac{1}{1 - \rho_m} m_0$, completing the argument. □

D.5 Proof of Proposition 4

By Proposition 3, we can equivalently prove the result for impulse responses \widehat{s}_t^a and \widehat{e}_t^a to an aggregate demand shock.

³⁹The conclusion is immediate by comparing to (A.7) - (A.8).

1. We know from the expressions in Lemma B.1 that the impulse responses \widehat{s}_t^a and \widehat{e}_t^a are independent of ϕ . The claim then follows immediately from the market-clearing condition

$$\widehat{y}_t = \phi \widehat{s}_t + (1 - \phi) \widehat{e}_t$$

together with the restriction that $\widehat{y}_0 = -1$.

2. Since services spending impulse responses to any shock vector (b_0^a, b_0^s, b_0^d) scale in $b_0^a + b_0^s$, while durables spending impulse responses scale in $b_0^a + b_0^d$, it follows that $\widehat{s}_t \propto \widehat{s}_t^a$ and $\widehat{e}_t \propto \widehat{e}_t^a$. The statement then follows from the assumed shares $\{\omega, 1 - \omega\}$ and the normalization that $\widehat{y}_0 = -1$. \square

D.6 Proof of Proposition 5

It is straightforward to show using guess-and-verify that, in the unique equilibrium of the flexible-price economy, the real interest rate satisfies

$$\widehat{r}_t = (1 - \rho_b) b_t^a$$

But the flexible-price economy is efficient, so the monetary authority indeed optimally sets nominal interest rates as

$$\widehat{r}_t^n = (1 - \rho_b) b_t^a,$$

replicating the flexible-price allocation. \square

D.7 Proof of Proposition 6

In the flexible-price analogue of the economy of Section 2.1, the equilibrium sequences of $\{\widehat{s}_t, \widehat{d}_t, \widehat{y}_t, \widehat{r}_t\}$ given sectoral shocks $\{b_{1t}, b_{2t}\}$ are fully characterized by the following system of equations:

$$\begin{aligned} -\gamma \widehat{s}_t &= \frac{1}{\varphi} \widehat{y}_t - (1 - \phi)(b_t^s - b_t^d) \\ -\gamma \widehat{d}_t &= \frac{1}{1 - \beta(1 - \delta)} \left[\frac{1}{\varphi} \widehat{y}_t + \phi(b_t^s - b_t^d) + \kappa(\widehat{d}_t - \widehat{d}_{t-1}) \right] \\ &\quad - \frac{\beta}{1 - \beta(1 - \delta)} \left\{ (1 - \delta) \left[\frac{1}{\varphi} y_{t+1} + \phi \rho_b (b_t^s - b_t^d) \right] + \kappa(\widehat{d}_{t+1} - \widehat{d}_t) \right\} \\ \widehat{y}_t &= \phi \widehat{s}_t + (1 - \phi) \frac{1}{\delta} (\widehat{d}_t - (1 - \delta) \widehat{d}_{t-1}) \end{aligned}$$

$$\widehat{r}_t = \frac{1}{\varphi}[\widehat{y}_t - y_{t+1}] + (1 - \rho_b)(\phi b_t^s + (1 - \phi)b_t^d)$$

We guess and verify that this system admits a solution with the lagged durables stock as the only endogenous state variable. Plugging in this guess and matching coefficients, we get the following system of eight equations in eight unknowns:

$$\begin{aligned} -\gamma\theta_{sd} &= \frac{1}{\varphi}\theta_{yd} \\ -\gamma\theta_{sb} &= \frac{1}{\varphi}\theta_{yb} - (1 - \phi) \\ \theta_{yd} &= \phi\theta_{sd} + (1 - \phi)\frac{1}{\gamma}[\theta_{dd} - (1 - \delta)] \\ \theta_{yb} &= \phi\theta_{sb} + (1 - \phi)\frac{1}{\delta}\theta_{db} \\ -\gamma\theta_{dd} &= \frac{1}{1 - \beta(1 - \delta)} \left\{ \frac{1}{\varphi}\theta_{yd} + \kappa(\theta_{dd} - 1) \right\} \\ &\quad - \frac{\beta}{1 - \beta(1 - \delta)} \left\{ (1 - \delta)\frac{1}{\varphi}\theta_{yd}\theta_{dd} + \kappa\theta_{dd}(\theta_{dd} - 1) \right\} \\ -\gamma\theta_{db} &= \frac{1}{1 - \beta(1 - \delta)} \left\{ \frac{1}{\varphi}\theta_{yb} + \phi + \kappa\theta_{db} \right\} \\ &\quad - \frac{\beta}{1 - \beta(1 - \delta)} \left\{ (1 - \delta)\left[\frac{1}{\varphi}(\theta_{yd}\theta_{db} + \theta_{yb}\rho) + \phi\rho_b\right] + \kappa\theta_{db}[(\theta_{dd} - 1) + \rho_b] \right\} \\ \theta_{rd} &= \frac{1}{\varphi}\theta_{yd}[1 - \theta_{dd}] \\ \theta_{rb} &= \frac{1}{\varphi}[\theta_{yb} - \theta_{yd}\theta_{db} - \theta_{yb}\rho_b] + (1 - \rho_b)\phi \end{aligned}$$

where the law of motion for $x \in \{s, d, y\}$ is $\widehat{x}_t = \theta_{xd}\widehat{d}_{t-1} + \theta_{xb}(b_t^s - b_t^d)$, while for r_t we have $\widehat{r}_t = \theta_{rd}\widehat{d}_{t-1} + \theta_{rb}b_t^s + (1 - \rho_b - \theta_{rb})b_t^d$.

To prove Proposition 6, it suffices to show that $\theta_{dd} > 0$, $\theta_{db} < 0$, $\theta_{rd} < 0$ and $\theta_{rb} > 0$. We have obtained closed-form solutions and verified that, in the special case of $\rho_b = \kappa = 0$, these inequalities all hold for $\delta \in (0, 1)$, $\gamma > 0$, $\varphi > 0$, $\phi \in (0, 1)$ and $\beta \in (0, 1)$. The expressions are unwieldy and thus omitted, but available upon request. \square

D.8 Proof of Propositions B.3 and B.4

Consider the equations characterizing the equilibrium in the model in Appendix A.1 for arbitrary degrees of price stickiness, but flexible nominal wages ($\phi_w = 0$).

Using the method of undetermined coefficients, we can show that the recursive representation of the equilibrium dynamics of $\widehat{s}_t, \widehat{d}_t$ takes the form of

$$\begin{aligned}\widehat{s}_t &= \theta_s \widehat{d}_{t-1} + \vartheta_1^s (\widehat{b}_t^a + \widehat{b}_t^s) + \vartheta_2^s (\widehat{b}_t^a + \widehat{b}_t^d) \\ \widehat{d}_t &= \theta_d \widehat{d}_{t-1} + \vartheta_1^d (\widehat{b}_t^a + \widehat{b}_t^s) + \vartheta_2^d (\widehat{b}_t^a + \widehat{b}_t^d)\end{aligned}$$

Following the same steps as in Appendix D.1, we can then show that

$$\begin{aligned}\widehat{s}_t &= \rho_b^t \widehat{s}_0 + \frac{\theta_d^t - \rho_b^t}{\theta_d - \rho_b} \theta_s \delta \widehat{e}_0 \\ \widehat{e}_t &= \widehat{e}_0 \left(\rho_b^t - (1 - \delta - \theta_d) \frac{\theta_d^t - \rho_b^t}{\theta_d - \rho_b} \right) \\ \frac{\widehat{y}_t}{\widehat{y}_0} &= \omega \frac{\widehat{s}_t}{\widehat{s}_0} + (1 - \omega) \frac{\widehat{e}_t}{\widehat{e}_0}\end{aligned}$$

Computing the normalized CIR of output and using the fact that $\widehat{e}_0 \equiv \widehat{s}_0 \frac{\phi}{1-\phi} \frac{1-\omega}{\omega}$, we obtain the expression in the proposition

$$\widehat{\mathbf{y}} = \frac{1}{1 - \rho_b} \left[1 - (1 - \omega) \left(1 - \frac{\delta}{1 - \theta_d} \left(1 + \frac{\phi}{1 - \phi} \theta_s \right) \right) \right]$$

Moreover, we have that

$$\begin{aligned}\widehat{\mathbf{s}}^a &= \frac{1}{1 - \rho_b} + \frac{\phi}{1 - \phi} \frac{1 - \omega^a}{\omega^a} \theta_s \widehat{\mathbf{e}}^a \\ \widehat{\mathbf{e}}^a &= \frac{\delta}{1 - \theta_d} \frac{1}{1 - \rho_b}\end{aligned}$$

Replacing these expressions above, we obtain the expression for \mathbf{y}

$$\widehat{\mathbf{y}} = \omega \left(\widehat{\mathbf{s}}^a - \widehat{\mathbf{e}}^a \left(1 + \frac{\phi}{1 - \phi} \frac{1}{\omega^a} \theta_s \right) \right) + \widehat{\mathbf{e}}^a \left(1 + \frac{\phi}{1 - \phi} \theta_s \right)$$

Thus, $\widehat{\mathbf{y}}$ is increasing in ω if and only if $\widehat{\mathbf{s}}^a > \widehat{\mathbf{e}}^a \left(1 + \frac{\phi}{1 - \phi} \frac{1}{\omega^a} \theta_s \right)$.

D.9 Proof of Proposition B.2

Set $\theta_d = 1 - \delta$ in the expressions for $\{\theta_b, \theta_d\}$ in Appendix D.1. Solving the system for (κ, θ_b) gives

$$\kappa = \gamma \frac{1 - \delta}{\delta}$$

and so

$$\theta_b = \frac{\delta}{\gamma}$$

as claimed. It thus follows that the impulse responses to an aggregate demand shock are

$$\begin{aligned} \widehat{s}_t &= -\frac{1}{\gamma} \times \rho_b^t \\ \widehat{e}_t &= -\frac{1}{\gamma} \times \rho_b^t \end{aligned}$$

establishing the proposition. □

D.10 Proof of Proposition B.5

We now for each good i get the optimality condition

$$\begin{aligned} -\gamma \widehat{d}_{it} &= \frac{1}{1 - \beta(1 - \delta_i)} \left[\widehat{\lambda}_t - (b_t^a + b_t^i) + \kappa_i (\widehat{d}_{it} - \widehat{d}_{it-1}) \right] \\ &\quad - \frac{\beta(1 - \delta_i)}{1 - \beta(1 - \delta_i)} \mathbb{E}_t \left[\widehat{\lambda}_{t+1} - (b_{t+1}^a + b_{t+1}^i) + \frac{\kappa_i}{1 - \delta_i} (\widehat{d}_{it+1} - \widehat{d}_{it}) \right] \end{aligned}$$

Following the same steps as in Appendix D.1, we find policy functions

$$\widehat{d}_{it} = \theta_i \widehat{d}_{it-1} + \theta_b^i (b_t^a + b_t^i)$$

where $\{\theta_i, \theta_b^i\}$ are the same as before, but with κ_i instead. Given those policy functions, the derivations of extended versions of Lemma B.1 and Proposition B.1 can proceed exactly as in the baseline case. □

D.11 Proof of Proposition B.6

The only decision of hand-to-mouth households is how to split their income at each time t between durable and non-durable consumption. Optimal behavior is fully characterized by

the optimality condition

$$-\gamma \widehat{d}_t^H = \frac{1}{1 - \beta(1 - \delta)} (-\gamma \widehat{s}_t^H + b_t^s - b_t^d) - \frac{\beta(1 - \delta)}{1 - \beta(1 - \delta)} \mathbb{E}_t [-\gamma \widehat{s}_{t+1}^H + b_{t+1}^s - b_{t+1}^d]$$

Aggregating across constrained households H and unconstrained households R :

$$\begin{aligned} \widehat{s}_t &= (1 - \mu) \widehat{s}_t^R + \mu \widehat{s}_t^H \\ \widehat{e}_t &= (1 - \mu) \widehat{e}_t^R + \mu \widehat{e}_t^H \end{aligned}$$

Unconstrained R types behave exactly as in the baseline model. Thus, together with the income exposure rule (A.2), we have arrived at the full set of equations characterizing the model equilibrium.

Then, to show the proposition, consider first setting $\eta = 1$. It is then straightforward to verify that all equilibrium relations are satisfied for $\widehat{x}_t = \widehat{x}_t^R = \widehat{x}_t^H$ for $x \in \{s, d, e, c\}$. Now consider arbitrary η . Then, following the same steps as in Bilbiie (2020), we can easily verify that the total response of output is scaled by a factor of $\frac{1-\mu}{1-\mu\eta}$, with unchanged shape. This completes the proof. \square

D.12 Proof of Proposition B.7

Note that, with time-varying real relative sectoral prices, the optimality conditions characterizing household consumption expenditure become

$$\begin{aligned} -\gamma \widehat{s}_t &= \widehat{\lambda}_t + \widehat{p}_t^s \\ -\gamma \widehat{d}_t &= \frac{1}{1 - \beta(1 - \delta)} \left[\widehat{\lambda}_t + \widehat{p}_t^d + \kappa(\widehat{d}_t - \widehat{d}_{t-1}) \right] \\ &\quad - \frac{\beta(1 - \delta)}{1 - \beta(1 - \delta)} \mathbb{E}_t \left[\widehat{\lambda}_{t+1} + \widehat{p}_{t+1}^d + \frac{\kappa}{1 - \delta} (\widehat{d}_{t+1} - \widehat{d}_t) \right] \end{aligned}$$

With $\widehat{\lambda}_t = 0$ from our assumptions on equilibrium selection, it follows immediately from (A.16)-(A.17) that equilibrium consumption and so sectoral output are exactly as in the baseline model with demand shocks. The only difference is in total hours worked, which are derived residually from (A.18) to ensure overall output market clearing. It follows that all our main results apply without change to supply shocks. \square