Demand Composition and the Strength of Recoveries†

Martin Beraja — MIT and NBER
Christian K. Wolf — University of Chicago

February 20, 2021

Abstract: We argue that recoveries from demand-driven recessions with expenditure cuts concentrated in services tend to be weaker than recoveries from recessions biased towards durables. Intuitively, the smaller the recession's bias towards durables, the less the subsequent recovery is buffeted by *pent-up demand*. We show that, in standard multi-sector business-cycle models, this result on recovery strength holds if and only if, following a contractionary monetary policy shock, durable expenditures revert back faster than services and non-durable expenditures. This condition receives ample support in aggregate U.S. time series data. We then use a semi-structural shift-share as well as a fully structural model to quantify our effect, asking how recovery strength varies with (i) differences in long-run expenditure shares across countries and (ii) the sectoral incidence of demand shocks across recessions. We find the effects to be large, and so 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@uchicago.edu`. We thank Marios Angeletos, Florin Bilbiie, Ricardo Caballero, Basile Grassi, Erik Hurst, Greg Kaplan, Andrea Lanteri, Simon Mongey, Matt Rognlie, Alp Simsek, Gianluca Violante, Iván Werning, Johannes Wieland (our discussant), Tom Winberry, and Nathan Zorzi for very helpful conversations, and Isabel Di Tella for outstanding research assistance.

1 Introduction

When a consumer decides against a car purchase in the midst of a recession, she simply postpones such expenditure for later (Mankiw, 1982; Caballero, 1993). This *pent-up demand* is likely to be absent or at least weaker in the case of services: when a consumer cuts down on a dinner away from home, she may not have two dinners out in the future — the lost services expenditure is simply foregone. In the aggregate, this logic would imply that durable expenditure cuts in a recession should reverse during the subsequent recovery, whereas the reversal in services (and non-durables) expenditures should be much weaker. Figure 1 documents precisely this pattern, here conditional on a contractionary monetary policy shock: durable expenditures exhibit a Z-shaped cycle, declining first and then overshooting, while services and non-durables expenditures follow a V-shape.

In this paper, we study how the composition of consumption expenditures during demand-driven recessions shapes subsequent recovery dynamics. We first show that standard multi-sector business-cycle models with demand-determined output can naturally generate the patterns in Figure 1. We then prove our main result: whenever such models are consistent with the documented sectoral expenditure patterns, they will invariably imply that recoveries from demand-driven recessions with expenditure cuts concentrated in services tend to be weaker than recoveries from recessions biased towards durables. Intuitively, the larger the recession's bias away from durables, the less the recovery is buffeted by pent-up demand effects. In practice, demand composition will differ across recessions chiefly because of differences in (i) long-run expenditure shares and (ii) the sectoral incidence of the underlying shocks.[1] We argue theoretically and empirically that the effect of both on recovery strength can be quantitatively meaningful. In light of this, we conclude the paper by discussing the implications of our results for the conduct of optimal stabilization policy.

To transparently illustrate the pent-up demand mechanism, our analysis begins with a stylized two-sector business-cycle model with perfectly transitory shocks and fully demand-determined output (e.g., due to perfectly rigid prices). A representative household derives utility from durable goods and services, with the durables stock depreciating at rate $\delta < 1$,

[1] Differences in long-run expenditure shares are large; for example, amongst OECD countries in 2017, the durables share ranged from 0.04 to 0.15 and the services share ranged from 0.3 to 0.68. Second, certain U.S. recessions featured particularly salient sectoral patterns due to the nature of the shocks. For example, following the oil crisis of 1973, durable expenditure declines (like cars) accounted for 165 percent of consumption expenditure declines (peak-to-trough), while in the COVID-19 recession services (like food at restaurants) and non-durable expenditures contributed around 85 percent.

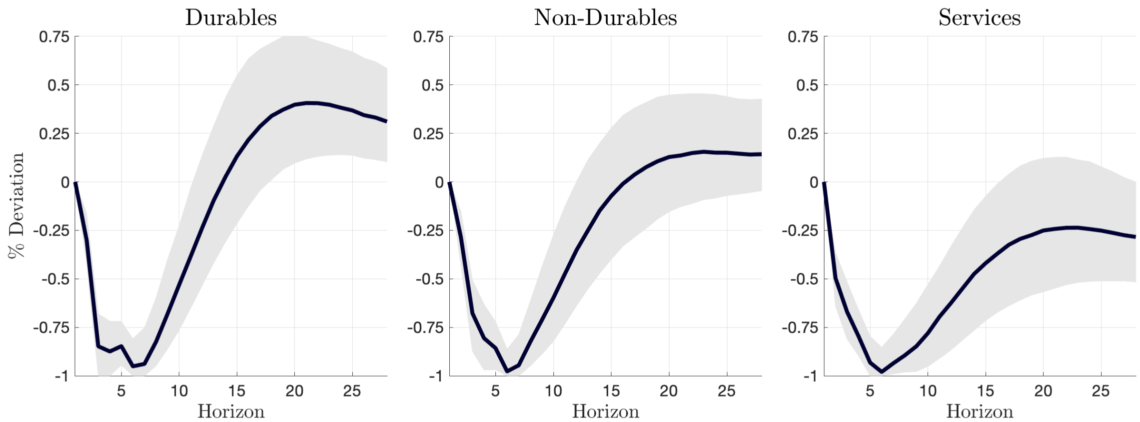

Figure 1: Quarterly impulse responses to a recursively identified monetary policy shock (as in Christiano et al. (1999)) by consumption spending category, all normalized to drop by -1% at the trough. The solid blue line is the posterior mean, while the shaded areas indicate 16th and 84th percentiles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

while services depreciate instantly. The marginal utility of household consumption is subject to three reduced-form demand shocks — one for each sector, and one to aggregate spending. 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 first establish that, following an arbitrary combination of aggregate and sectoral demand shocks, the impulse response of durable expenditures is 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 initially at $t=0$, and then just returns to baseline at $t=1$. Since the special case of an aggregate demand shock common to all sectors is equivalent to an ordinary monetary policy shock, we can conclude that the simple model is qualitatively consistent with the patterns in Figure 1. At the same time, the model predicts that recoveries from recessions concentrated in durables spending are stronger than those from recessions biased towards services: when services account for a share ω of the expenditure decline at $t=0$, aggregate output overshoots at $t=1$, with the overshoot equal to a fraction $(1-\omega)(1-\delta)$ of the initial drop. The cumulative impulse response (CIR) of output relative to its trough — a natural measure of persistence and so weakness of recovery — is then equal to $1-(1-\omega)(1-\delta)$. It follows that, as claimed, recoveries are weaker for a larger services share ω. In particular, the 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 sectoral 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; adjustment costs on durables; imperfectly sticky prices and wages; incomplete markets and hand-to-mouth households; supply shocks; and an arbitrary number of goods varying in their durability. We prove that, in this extended setting, our main result on the effects of demand composition on recovery strength continues to hold *if and only if*, conditional on a contractionary common demand shock, the CIR for durables spending (relative to its trough) is strictly smaller than the corresponding CIR for services and non-durables spending. Thus, through the lens of this class of models, Figure 1 provides strong evidence in favor of our central hypothesis. For further empirical support, we document similar patterns following: (i) uncertainty shocks (Basu & Bundick, 2017), (ii) oil shocks (Hamilton, 2003), and (iii) reduced-form forecast errors of sectoral output.

In the second part of the paper, we quantify the effects of demand composition on the strength of recoveries. We do so in two ways. The first approach is a simple shift-share. We prove that, in the class of models described above, the behavior of aggregate consumption in a demand-driven recession of arbitrary sectoral composition can be estimated semi-structurally, by suitably weighting and then summing the category-specific consumption responses to a *common* demand shock. We do so using the impulse responses displayed in Figure 1, with the weights chosen in line with (i) observed cross-country variation in expenditure shares and (ii) observed cross-recession variation in sectoral incidence. Our second approach is fully structural, and relies on an extended model that violates the conditions required by the shift-share. We calibrate this model and then, mirroring the shift-share, compute output CIRs in model economies with: (i) different long-run expenditure shares and (ii) different mixes of sectoral shocks. The two approaches paint a consistent picture: the effects of sectoral spending composition on recovery strength are estimated to be large. For example, the CIR of output in a U.S. recession as biased towards services as COVID-19 is estimated to be about 70 to 90 per center larger than that of an average durables-led recession. Similarly, moving from an economy like the U.S. to one with the high durable expenditure share of Canada, the output CIR to a given common aggregate demand shock decreases by about 15 per cent.

In light of this quantitative relevance, we conclude with a discussion of (optimal) stabilization policy. Our main finding is that our two main sources of heterogeneity in sectoral composition — differences in long-run expenditure shares and sectoral shock incidence — actually have very different implications for optimal policy design. First, in an economy subject only to common (i.e., not sectoral) demand disturbances, optimal policy turns out to be completely independent of long-run expenditure shares. Intuitively, changes in shares affect not only the transmission of exogenous demand shocks, but also that of the stabilization policy itself; in our model, these two effects exactly offset, leaving optimal monetary policy unaffected. It follows that, at least in our setting, the presence of a durables good sector *per se* is irrelevant for the conduct of optimal stabilization policy. 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 service sector, and thus the longer the expected recession in the absence of monetary stabilization.

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 aggregate business-cycle dynamics. So far, most work has emphasized the effects of durables on recession severity (Barsky et al., 2007) and state-dependent shock elasticities (Berger & Vavra, 2015). Similar to our Figure 1, Erceg & Levin (2006) and McKay & Wieland (2020) highlight that durables spending tends to reverse over time after monetary policy shocks.[2] Our analysis offers additional insights by discussing the implications of this observation for how demand composition affects recovery dynamics in general, and for the design of *optimal* monetary policy in particular.

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; 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 the strength and shape of recoveries. The mechanisms discussed in previous work 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, driven by either (i) structural forces leading to differences

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

in long-run expenditure shares or (ii) the nature of shocks. In fact, our results regarding changes in long-run expenditure shares are consistent with the empirical results in Olney & Pacitti (2017), who show that U.S. states with higher shares of non-tradable services tend to have slower employment recoveries.

Finally, we 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 possible 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 highlight one very particular mechanism – pent-up demand — that is may well be weaker during this recovery than in previous ones.

Outline. Section 2 provides analytical characterizations of business-cycle dynamics in a multi-sector general equilibrium model with demand-determined output. Section 3 connects the predictions of our theory to time series evidence on the propagation of shocks to household spending. Section 4 blends theory and empirics to quantify 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.

2 Pent-up Demand and Recovery Dynamics

This section presents our main theoretical results on recovery dynamics in an economy with durables and services. Section 2.1 outlines the model. Sections 2.2 and 2.3 then illustrate the pent-up demand mechanism in a stripped-down variant and discuss implications for recovery dynamics. Finally Sections 2.4 and 2.5 extend those insights back to the full model.

2.1 Model

We consider a discrete-time, infinite-horizon economy populated by a representative household, monopolistically competitive retailers, and a government. Households consume services and durables, and the only source of aggregate risk are shocks to household preferences over consumption bundles.[3]

[3] In Section 2.5, we consider an extended variant of this economy in which households consume N goods with different durability (instead of only services and durables), some households are hand-to-mouth (instead of there being a representative agent), and there are sectoral productivity shocks (in addition to household demand shocks).

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

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

where we assume

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

b_t^c is a common shock to aggregate demand, 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 household spending — e.g., increased precautionary savings due to greater income risk ($b^c < 0$) or increased fear of consuming certain services during a pandemic due to greater infection risk ($b^s < 0$). The scaling factors $\{\alpha, \varsigma_c\}$ are chosen to ensure that, in the flexible-price limit of our economy, the aggregate demand shock b_t^c has no real effects on equilibrium quantities (to first order), instead only moving the path of real interest rates. $\{\varsigma_s, \varsigma_d\}$ are then pinned down by the relative sizes of the services and durables sectors, ensuring that a combined shock $b_t^d = b_t^s$ is isomorphic to a common aggregate demand shock of the same magnitude.[4]

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, 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 restrict attention

[4]See Appendix A.1 for the expressions. We think that the neutrality property for the common aggregate shock b_t^c is desirable because it holds in the textbook New Keynesian model with only non-durables. Our definition of b_t^c is the natural extension of this notion of an "aggregate demand shock" to a multisector economy with durables; in particular, it is isomorphic to a shock to the shadow price of the total household consumption bundle, and so readily seen to be equivalent to standard monetary policy shocks (see Proposition 3). However, we emphasize that our results on recovery dynamics are largely invariant to reasonable alternative definitions of "common" aggregate demand shocks. For a detailed discussion, please see Appendix B.2. We thank our discussant Johannes Wieland for raising this point.

to a standard quadratic specification:

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

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 be[5]

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

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.

Production. Both services and durable goods are produced by aggregating varieties sold by monopolistically competitive retailers. Production 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. Consistent with the empirically documented absence of significant short-run relative price movements (House & Shapiro, 2008; McKay & Wieland, 2020), we assume that the 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 real relative prices.

In equilibrium, aggregate output y_t must equal total consumption expenditures. In log-deviations from the steady state (denoted by $\widehat{\ }$) aggregate output then satisfies[6]

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

Policy. The monetary authority sets the nominal rate of interest on bonds, r_t^n. For our quantitative explorations in Section 4.3 we will consider a standard rule of the form

$$\widehat{r}_t^n = \phi_\pi \widehat{\pi}_t \tag{2}$$

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

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

For much of the remainder of this section, we will instead consider a monetary rule that fixes the (expected) real rate of interest.

SHOCKS. The disturbances b_t^c, b_t^s and b_t^d follow exogenous AR(1) processes with common persistence ρ_b and innovation volatilities $\{\sigma_b^c, \sigma_b^s, \sigma_b^d\}$, respectively.

2.2 The Pent-Up Demand Mechanism

We use a stripped-down version of the baseline model above to cleanly illustrate the pent-up demand mechanism. Specifically, we assume that: (i) all shocks are perfectly transitory ($\rho_b = 0$), (ii) there are no adjustment costs ($\kappa = 0$), (iii) durables and services are neither complements nor substitutes ($\zeta = \gamma$), and (iv) prices and wages are fully rigid and the nominal interest rate is fixed.

In this economy, we characterize sectoral and aggregate output dynamics conditional on an arbitrary vector of time-0 shocks $\{b_0^c, b_0^s, b_0^d\}$. To ensure equilibrium determinacy given assumption (iv), we impose that output ultimately reverts back to steady-state:

$$\lim_{t\to\infty} \widehat{y}_t = 0 \tag{3}$$

Given the equilibrium selection in (3), we arrive at the following characterization of aggregate impulse response functions.[7]

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

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

and

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

The impulse response of aggregate output is thus

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

[7] Equivalently, those impulse responses can be interpreted as applying to an economy where monetary policy is neutral, in the sense that it fixes the expected real rate, i.e., $\widehat{r}_t^n = \phi_\pi \mathbb{E}_t[\widehat{\pi}_{t+1}]$, with $\phi_\pi = 1$. This equilibrium selection can be formally justified with the continuity argument of Lubik & Schorfheide (2004): For $\phi_\pi \to 1^+$, our equilibrium selection delivers continuity in ϕ_π.

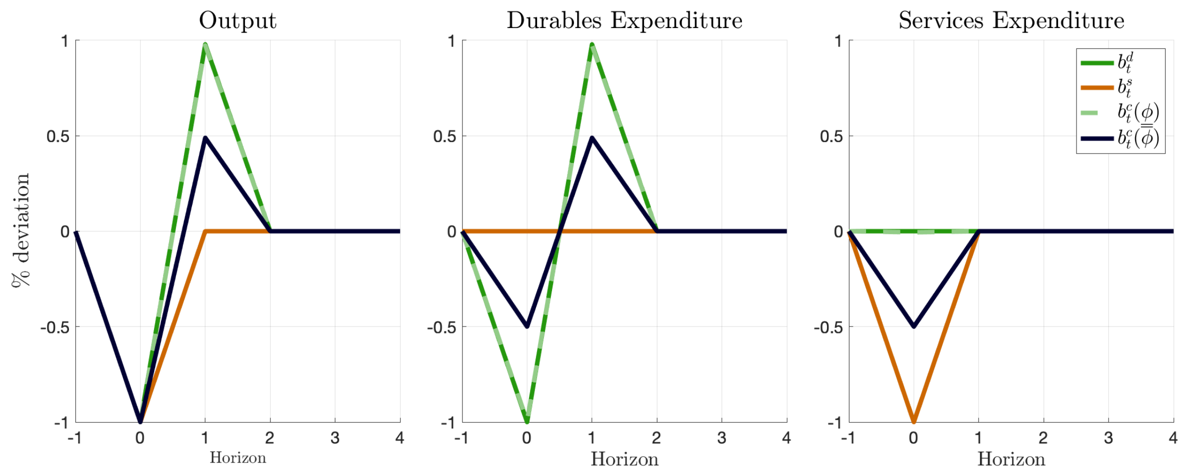

Figure 2: Recession dynamics in the stripped-down model. Responses for: a pure durables shock (green), a pure services shock (orange), and a common demand shock in an economy with a low services share $\underline{\phi}$ (dashed green) and a high services share $\bar{\phi}$ (dark blue). For details on the model parameterization see Appendix A.1.

Figure 2 shows impulse responses to three possible sets of time-0 shock vectors $\{b_0^c, b_0^s, b_0^d\}$, 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 response dynamics — 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.

First, the solid green lines depict impulse responses to a pure durables demand shock ($b_0^d < 0$) — or equivalently, impulse responses to a common demand shock ($b_0^c < 0$) in an economy with only durables ($\phi = 0$). Consumption demand and so equilibrium output decline on impact. Following the contraction in durables spending, the household durable stock at the beginning of the recovery is below target, so there is *pent-up demand* for durables. As a result, durable expenditures overshoot their steady-state at $t = 1$, and so does aggregate consumption demand. But since output is demand-determined, output also overshoots at $t = 1$ — a Z-shaped cycle. Second, the solid orange lines depict impulse responses to a pure services demand shock ($b_0^s < 0$) — or equivalently, impulse responses to a common demand shock ($b_0^c < 0$) in an economy with only services ($\phi = 1$). In this case 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 the cycle is V-shaped. Third, the dashed green and solid blue lines show impulse responses to a common demand shock ($b_0^c < 0$) 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.

Relation to empirical evidence. The results in Figure 2 are qualitatively consistent with the empirical impulse response estimates presented in Figure 1: in both cases, conditional on a common aggregate demand shock at $t = 0$, durables expenditures show a sharp overshoot, while services expenditures return to baseline from below.[8] Thus, as soon as consumption goods are heterogeneous in their durability, a simple multi-sector New Keynesian model will invariably generate sectoral heterogeneity in impulse responses of the sort documented in aggregate time series data.

The next subsection explores implications of this observation for aggregate recovery dynamics, again within the confines of our stripped-down model. Sections 2.4 and 2.5 extend all results back to our rich baseline model (and beyond), and Section 3 formalizes the connection between those theoretical results and the empirical evidence of Figure 1.

2.3 Implications for Recovery Dynamics

The model of Section 2.2 makes strong predictions about how the sectoral composition of spending declines in a recession affects recovery dynamics. To show this we begin by defining two objects. First, we denote the share of services expenditures in time-0 aggregate consumption expenditure changes by ω:

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

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

$$\widehat{\mathbf{y}} \equiv \frac{\sum_{t=0}^{\infty}\widehat{y}_t}{\widehat{y}_0} \tag{8}$$

The normalized CIR measures the weakness of the reversal of output in the recovery phase; given a recession at $t = 0$, the CIR is smaller when output reverts to steady state faster (or

[8] Our choice of the scaling factors $\{\alpha, \varsigma_c\}$ ensures that, in our setting, common aggregate demand shocks and conventional monetary policy shocks are equivalent. We state the formal result in Section 3.

overshoots). Therefore, we will say that *a recovery is stronger* whenever $\widehat{\mathbf{y}}$ is smaller.[9]

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

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

$$\widehat{\mathbf{y}} \quad = \quad 1 \quad - \quad (1-\omega)(1-\delta). \tag{9}$$

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 2 and the discussion surrounding it: the larger the services share ω, the smaller pent-up demand effects, and so the weaker the subsequent recovery.

In practice, there are at least two reasons to expect ω to vary across recessions. First, across countries (or in the same country over time), changes in ϕ imply changes in ω for any given set of shocks. Our results imply that, the larger an economy's ϕ, the slower its recovery from any given common aggregate demand shock b_0^c. Second, ω may differ across recessions because recessions may be heterogeneous in their shock incidence $\{b_0^c, b_0^s, b_0^d\}$. By (9), recoveries from recessions driven by shocks to services demand (b_0^s) will tend to be more gradual than recoveries following shocks to durables demand (b_0^d). We assess both of these channels quantitatively in Section 4.

2.4 Back to the Full Model

We now show that the pent-up demand mechanism and its implications for recovery dynamics extend to the general model of Section 2.1.

We begin by considering a variant of this general model with separable preferences ($\gamma = \zeta$) and a passive monetary policy rule that fixes the (expected) real rate of interest.[10] In the end we briefly explore the effects of non-separabilities in household preferences and of alternative monetary policy rules.

[9] An 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 normalized CIR.

[10] Note that a rule of this sort is consistent with any degree of price stickiness except for the limit case of *perfect* price flexibility. As before, equilibrium selection given this rule will rely on (3).

Impulse responses. We proceed exactly as before: first characterizing sectoral and aggregate impulse response paths for arbitrary shock mixtures $\{b_0^c, b_0^s, b_0^d\}$, and then discussing implications for recovery dynamics.

Lemma 2. *Suppose that the monetary authority fixes the real rate of interest and that $\gamma = \zeta$. Then the impulse responses of services and durables consumption expenditures to a vector of time-0 shocks $\{b_0^c, b_0^s, b_0^d\}$ satisfy*

$$\widehat{s}_t = \frac{1}{\gamma}(b_0^c + b_0^s)\rho_b^t \tag{10}$$

and

$$\widehat{e}_t = \frac{1}{\gamma}(b_0^c + b_0^d)\frac{\theta_b}{\delta}\left(\rho_b^t - (1-\delta-\theta_d)\frac{\theta_d^t - \rho_b^t}{\theta_d - \rho_b}\right) \tag{11}$$

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

$$\widehat{y}_t = \phi\widehat{s}_0\rho_b^t + (1-\phi)\widehat{e}_0\left(\rho_b^t - (1-\delta-\theta_d)\frac{\theta_d^t - \rho_b^t}{\theta_d - \rho_b}\right) \tag{12}$$

Lemma 2 reveals that the pent-up demand logic at the heart of our argument 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. Second, for $\kappa > 0$, adjustments in the durables stock are slowed down, adding mechanical endogenous persistence that offsets pent-up demand effects. In this case, the pent-up demand effects will continue to dominate if and only if $\theta_d < 1 - \delta$.

Demand composition and recovery dynamics. We can now as before translate Lemma 2 into a result relating demand composition and the strength of the recovery.

Proposition 2. *Suppose that the monetary authority fixes the real rate of interest and that $\gamma = \zeta$, and consider a vector of time-0 shocks $\{b_0^c, b_0^s, b_0^d\}$ with a services share ω. Then, the normalized cumulative impulse response of aggregate output satisfies*

$$\widehat{\mathbf{y}} \quad = \quad \frac{1}{1-\rho_b}\left[1 \quad - \quad (1-\omega)(1 - \frac{\delta}{1-\theta_d})\right] \tag{13}$$

Proposition 2 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 $\theta_d < 1 - \delta$. Fortunately, this abstract condition on model primitives can be translated into a simple-to-interpret condition on objects which can be measured in the data. The following theorem does so, stating a necessary and sufficient condition for our results on recovery strength to go through.

Theorem 1. *Suppose that the monetary authority fixes the real rate of interest and that* $\gamma = \zeta$*. Let* $\widehat{\mathbf{s}}^c$ *and* $\widehat{\mathbf{e}}^c$ *denote the normalized cumulative impulse responses of services and durables expenditure to a recessionary common demand shock* $b_0^c < 0$*, defined as in* (8).

Then, the normalized cumulative impulse response of aggregate output $\widehat{\mathbf{y}}$ *in* (13) *is increasing in the services share* ω *if and only if*

$$\widehat{\mathbf{s}}^c \quad > \quad \widehat{\mathbf{e}}^c \tag{14}$$

Theorem 1 links the sectoral CIRs to a *particular* type of shock (the common shock b_0^c) to how the strength of recovery varies with the services bias in demand composition ω. Again, this result holds regardless of whether such variation in ω resulted from (i) changes in the steady-state share ϕ in an economy subject to that *same* common demand shock alone or (ii) the realization of *other* sector-specific shocks $\{b_0^s, b_0^d\}$.

Non-separability, sticky prices, and other monetary rules. In Appendix B.1, we relax the simplifying assumptions of separability ($\gamma = \zeta$) and a passive monetary rule. There, we provide a generalized version of the condition (14). The expression reveals that: (i) in the empirically relevant case of net substitutability, (14) is likely to remain sufficient, thus even further strengthening our results; (ii) with flexible wages, arbitrarily sticky prices and a monetary rule of the form, (14) is generally only necessary, not sufficient. However, as we show through model simulations in Section 4.3, reasonable model calibrations satisfying (14) also robustly imply that $\widehat{\mathbf{y}}$ is increasing in ω.

Outlook. In Section 3 we take the condition (14) to the data. By Theorem 1, testing (14) is equivalent — at least through the lens of our model — to testing our predictions on recovery dynamics. Before doing so, however, we briefly present generalizations of (14) beyond the baseline model of Section 2.1.

2.5 Further Generalizations

We provide a summary discussion of further model extensions here, and relegate details to Appendices A.2 and B.2.

Incomplete markets. Proposition 2 and Theorem 1 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 cyclicality of income for hand-to-mouth households, the impulse responses in Lemma 2 are scaled up or down. Impulse response *shapes*, however, are unaffected by this scaling, and so our conclusions on recovery dynamics are entirely unaffected.

Many sectors. We 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 proofs of Proposition 2 and Theorem 1, we can show that the normalized output CIR $\widehat{\mathbf{y}}$ for an arbitrary shock mix $\{b_0^c, \{b_0^i\}_{i=1}^N\}$ that results in shares $\{\omega_i = \frac{\phi_i \widehat{e}_0^i}{\widehat{y}_0^i}\}_{i=1}^N$ is given by

$$\widehat{\mathbf{y}} \quad = \quad \sum_{i=1}^{N} \omega_i \frac{\delta_i}{1-\theta_d^i} \quad = \quad \sum_{i=1}^{N} \omega_i \widehat{\mathbf{e}}_i^c \tag{15}$$

Thus, equation (15) is a natural extension of the two-sector expressions in (13) and (14).

General adjustment costs. Our baseline model considered a very particular (convenient) form of quadratic adjustment costs in the durable stock. Consider instead a general adjustment cost function of the form

$$\psi(\{d_{t-\ell}\}_{\ell=0}^{\infty}) \tag{16}$$

Importantly, (16) is general enough to nest arbitrary forms of non-quadratic adjustment costs as well as 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 $\gamma = \zeta$ and the path of real rates is fixed, it is still true that

$$\widehat{\mathbf{y}} \quad = \quad \omega \widehat{\mathbf{s}}^c \ + \ (1-\omega)\widehat{\mathbf{e}}^c,$$

for any vector of shocks $\{b_0^c, b_0^s, b_0^d\}$ resulting in services share ω.[11] Thus (14) still applies. Intuitively, the crucial restriction is that the system of equations characterizing the equilibrium remains separable in s_t and d_t.

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 passive 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.[12]

3 Pent-Up Demand in Time Series Data

The main hypothesis of this paper is that recoveries from demand-driven recessions concentrated in services tend to be weaker than recoveries from recessions biased towards durables. In Section 2 we have shown that, in standard structural multi-sector macro models, this hypothesis is true if and only if durable expenditures exhibit a stronger reversal than services (and non-durables) expenditures *conditional on a common aggregate demand shock*.

In this section, we test the validity of our hypothesis by testing this condition. We proceed in two steps. First, in Section 3.1, we revisit Figure 1 and study sectoral expenditure dynamics conditional on monetary policy shocks. Second, in Section 3.2, we discuss supporting evidence from several other experiments.

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

[11] The scaling coefficient α in household preferences, however, may change, adjusting to ensure that the demand shocks $\{b_t^c, b_t^d, b_t^s\}$ enter all first-order conditions exactly additively with the marginal utility term $\widehat{\lambda}_t$ (e.g., as in (A.10) - (A.11)).

[12] Of course, by the 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.

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 a common aggregate demand shock b_t^c, and so directly map into the empirical test of Theorem 1. 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.

Proposition 3. *Consider the model of Section 2.1, extended to feature innovations m_t to the central bank's rule (2). The impulse responses of all real aggregates $x \in \{s, e, d, y\}$ to (i) to a recessionary common demand shock $b_0^c < 0$ with persistence ρ_b, and (ii) a contractionary monetary shock $m_0 = -(1-\rho_b)\varsigma_c b_0^c$ with persistence $\rho_m = \rho_b$ are identical:*

$$\widehat{x}_t^c \;=\; \widehat{x}_t^m$$

Intuitively, equivalence obtains because both our common aggregate demand shock as well as conventional monetary shocks move the shadow price of the household consumption bundle. We can thus test the key condition (14) using sectoral impulse responses to monetary policy shocks.

Empirical framework. Our analysis of monetary policy transmission closely follows the seminal contribution of Christiano et al. (1999): We estimate a reduced-form Vector Autoregression (VAR) in measures of consumption, output, prices and the federal funds rate, and identify monetary policy shocks as the innovation to the federal funds rate under a recursive ordering, with the policy rate ordered last.

We estimate our VARs on 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 categories of household spending — durables, non-durables, and services.[13] We include four lags throughout, and estimate the models using standard Bayesian techniques. Details are provided in Appendix C.1.

[13] As shown in Plagborg-Møller & Wolf (2020), the econometric estimands of all three specifications would be identical if the different measures of sectoral consumption did not affect the forecast errors in the non-consumption equations. Since the additional explanatory power (in a Granger-causal sense) of sectoral consumption measures for other macroeconomic aggregates is relatively small in our set-up, all three specifications are effectively projecting on the same shocks.

Results. Consistent with previous work, we find that a contractionary monetary policy shock lowers output and consumption.[14] Figure 1 — our motivating figure from the introduction — 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.

To facilitate the comparison of empirical estimates with the theoretical predictions in Proposition 2 and Theorem 1, we scale the impulse response of each component to drop by -1 per cent at the trough. To test (14), we compute the posterior distribution of[15]

$$\frac{\mathbf{s}^c}{\mathbf{e}^c} - 1$$

We find that, at the posterior mode, the normalized services CIR is 88 per cent larger than the durables CIR. This difference is also statistically significant, with the 68 per cent posterior credible set ranging from 10 per cent to 250 per cent. Similarly, we find that the non-durables spending CIR is between the two, around 22 per cent larger than the durables CIR. We conclude that the empirical evidence is consistent with (14) and thus with our main hypothesis about the effects of demand composition on recovery dynamics. In Section 4 we go beyond such qualitative statements and proceed to *quantify* this effect. Before doing so, however, we review other, complementary evidence.

3.2 Other Experiments

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

[14] In our baseline specification, prices increase — the price puzzle. Augmenting our model to include a measure of commodity prices ameliorates the price puzzle, without materially affecting other responses.

[15] In computing the CIRs, 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 CIR ratio for each draw from the posterior, and then report percentiles.

Uncertainty. Uncertainty shocks are a natural structural candidate for the common reduced-form demand shocks b_t^c, 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. Consistent with Plagborg-Møller & Wolf (2020), our VAR-based implementation controls for a large number of shock lags, ensuring consistent projections even at medium horizons.

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

Oil. As a third test, we study oil price shocks, identified as in Hamilton (2003) and embedded 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 CIRs at a finer sectoral level, as in (15).

Again, the results support our main hypothesis. 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, clothes and finance expenditures instead all dip in the initial recession, but then simply return to baseline, without any further overshoot. We discuss several additional sectoral impulse responses in Appendix C.3.

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 Theorem 1. We here complement these shock-specific results by instead looking at *unconditional* sectoral expenditure dynamics. Implicitly, in looking at such reduced-form forecasts, we are assuming that sectoral dynamics are largely driven by common, aggregate shocks; in that case, unconditional forecasts can also be used for the test in (14).

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

4 Quantifying the Effects of Demand Composition

Having documented qualitative support for our main claim on the effects of sectoral demand composition on recovery dynamics, we now turn to quantification. Section 4.1 describes and motivates our counterfactual exercises. Section 4.2 shows that, even in relatively general variants of our structural model in Section 2.1, the desired counterfactual impulse responses can be estimated directly through a simple shift-share design on the impulse responses to a common aggregate demand shock — i.e., our estimates from Section 3. In Section 4.3, we instead use a calibrated structural model to recover the desired counterfactuals, and then consider the sensitivity of our results to a wide range of plausible model parameterizations, in particular on the degree of price stickiness and adjustment costs.

4.1 Sources of Variation in Demand Composition

We will consider two kinds of counterfactual exercises.

The first exercise is motivated by the observed differences in long-run expenditures shares across countries, possibly due to structural forces. Results are displayed in the left panel of Figure 3. The figure reveals that economies differ widely in their sectoral make-up. We thus ask: fixing a common shock to aggregate household demand, how different would the recovery look like in a high-durables economy (e.g., Canada) vs. a low-durables economy (e.g., Colombia) or a low-services economy (e.g, Russia)?

The second exercise is motivated by the stark sectoral patterns observed in some past U.S. recessions, reflecting heterogeneity in the sectoral incidence of shocks. The right panel of Figure 3 shows three examples. As is well known, real expenditure declines in a typical U.S. recession tend to be more biased towards durable expenditures. An extreme example of this general pattern is 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

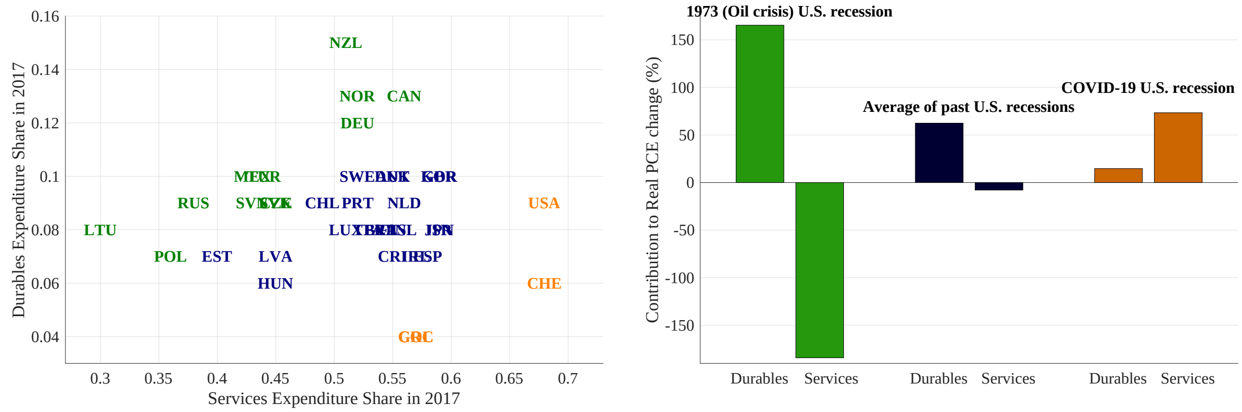

Figure 3: Left panel: Durables and services expenditure shares across OECD countries in 2017. Source: stats.oecd.org. Right panel: Contributions of durables and services expenditures changes to real personal consumption expenditures (PCE) changes in a recession. Average of past U.S. recessions (average of peak-to-trough changes from 1960 to 2019), 1973 oil crisis recession (peak-to-trough), and COVID-19 recession (February to May 2020). Source: bea.gov

overall accounted for more than 100 percent of the total expenditure decline. At the other extreme, the COVID-19 pandemic triggered a recession in which services and non-durables spending cuts accounted for almost all of the total expenditure decline — fearing infection, consumers mostly cut down on food away-from-home as well as travel- and health-related services. We thus ask: how different would the recovery be following combinations of shocks that induced a spending composition as in the average U.S. recession vs. the one observed during the 1973 oil recession or the COVID-19 recession?

4.2 Shift-Share Design

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 a common demand shock b_t^c. To quantify the effect of demand composition on the strength of the recovery, Proposition 4 gives sufficient conditions under which the response of total consumption to (i) a common shock b_t^c in an economy with arbitrary sectoral composition or (ii) an arbitrary combination of sectoral shocks $\{b_t^c, b_t^s, b_t^d\}$ in the baseline economy can be recovered through a simple shift-share based on the estimated sectoral responses to b_t^c.[16]

[16]For consistency, we present Proposition 4 in the context of the model of Section 2.1. However, as the proof makes clear, the result does not hinge on our particular parametric form (1) of the adjustment cost

Proposition 4. *Consider the model of Section 2.1 with $\gamma = \zeta$, and suppose that the monetary authority fixes the expected real rate of interest, 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 a common demand shock b_0^c 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 $\{b_0^c, b_0^s, b_0^d\}$ with persistence ρ_m and such that $\{\widehat{y}_0 = -1, \phi\widehat{s}_0 = -\omega, (1-\phi)\widehat{e}_0 = -(1-\omega)\}$ is*

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

Note that Proposition 4 is derived in the context of the baseline two-sector structural model of Section 2.1. Since our empirical estimates in Section 3.1 split spending into three categories, we use the natural three-sector extension of Proposition 4, derived easily from our general multi-sector characterizations in Section 2.5.

Results. Under the conditions of Proposition 4, we can use the sectoral monetary policy impulse responses from Figure 1 to construct our two desired counterfactuals. The left panel of Figure 4 shows CIRs for a common demand shock b_0^c as a function of the durables and services share — our first counterfactual.[17] In the figure, we have normalized the CIR of an economy with the sectoral composition of the U.S. to 1. The color shadings reveal that, as sectoral shares are adjusted, the strength of recoveries as measured by the normalized CIR 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 CIR is around 15 per cent smaller; on the other hand, for economies with a durables share as low as that in Colombia, the CIR can be around 5 per cent larger.

function. In particular, the result applies unchanged for adjustment costs on the *flow* of durable expenditures.

[17] Note that the non- or semi-durables share is then recovered residually.

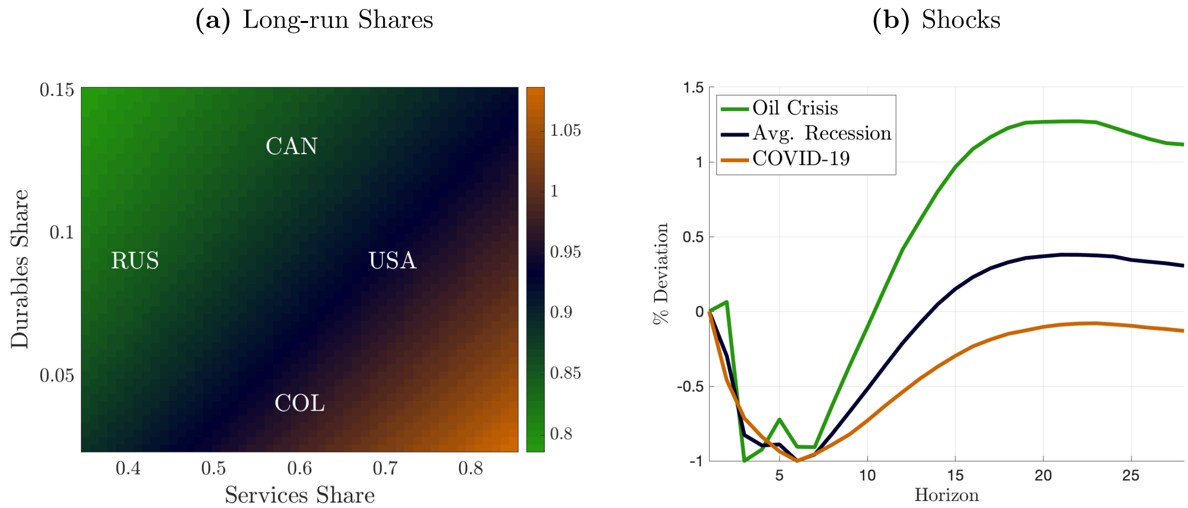

Figure 4: Left panel: CIR to a common demand shock b_0^c as a function of long-run expenditure shares, with the U.S. CIR normalized to 1, computed using the posterior mode point estimates from Figure 1. Right panel: Impulse response of total consumption to sectoral demand shocks reproducing expenditure composition changes in (i) ordinary recessions, (ii) the 1973 oil crisis, and (iii) the COVID-19 recession, all normalized to lead to a peak-to-trough consumption contraction of -1 per cent and evaluated again using the posterior mode point estimates from Figure 1.

The right panel shows entire impulse response paths for a vector of sectoral demand shocks with a peak effect on consumption of -1 per cent and sectoral composition of expenditure changes from peak-to-trough as in (i) an average U.S. recession, (ii) the oil crisis of 1973, and (iii) the COVID-19 recession — our second counterfactual. As expected, the durables-biased oil crisis shows a fast reversal, while the recovery from an ordinary recession is more gradual, and the recovery from a heavily services-biased recession (like COVID-19) is even weaker. In CIR terms, the implied effects are very large; for example, at the point estimates displayed in Figure 4, the CIR of output in a recession as biased towards services as COVID-19 is 67.8 per cent larger compared to an average, more durables-led recession, with the difference strongly statistically significant.[18]

4.3 Structural Counterfactuals

In this section we instead compute our two counterfactuals in fully parameterized, explicit structural models. We return to the baseline model of Section 2.1, and then depart from the analysis in Section 2.4 by allowing for imperfectly sticky prices and wages in conjunction

[18] The 68 per cent posterior credible set here ranges from 20 per cent to 170 per cent.

Parameter	Description	Value	Source/Target
Preferences			
β	Discount Rate	0.99	Annual Real FFR
γ	Inverse EIS	1	Literature
ζ	Elasticity of Substitution	1	= EIS
ϕ	Durables Consumption Share	0.1	NIPA
Technology			
ε_w	Labor Substitutability	10	Literature
δ	Depreciation Rate	0.021	BEA Fixed Asset
ϕ_w	Wage Re-Set Probability	0.2	Literature
Policy			
ϕ_π	Inflation Response	1.5	Literature
Shocks			
ρ_b	Demand Shock Persistence	0.83	Lubik & Schorfheide (2004)

Table 4.1: Calibration of fixed parameters for the quantitative structural model.

with a conventional monetary policy rule as in (2). Imperfect price and wage stickiness together with a non-passive monetary policy breaks the neat mapping between sectoral spending impulse responses to common shocks b_t^c and to sectoral shocks $\{b_t^s, b_t^d\}$ at the heart of Proposition 4, thus forcing us to rely on numerical simulations. Rather than focussing on a particular baseline parameterization, we will show that both counterfactuals remain quantitatively meaningful over a very large *range* of plausible parameterizations.

Calibration: fixed parameters. Table 4.1 presents our calibration of a set of baseline parameters that will be kept fixed across experiments.

The three preference parameters (β, ζ, γ) are standard; in particular, we continue to set $\zeta = \gamma$, so durables and services are neither net complements nor net substitutes. We consider a broad notion of durables, and thus set the depreciation rate δ as annual durable depreciation divided by the total durable stock in the BEA Fixed Asset tables, exactly as in McKay & Wieland (2020). 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). Next, for monetary policy, we consider the conventional Taylor rule in (2). Our policy rule is active, so real interest rates now drop following negative demand shocks, thus feeding back into spending on both durables and services, and breaking the separability at the heart of the shift-share. Finally, we take the persistence ρ_b of demand shocks from Lubik & Schorfheide (2004).

Calibration: parameter ranges. Two parameters have so far been left unrestricted — the durables adjustment cost κ and the slope of the New Keynesian Phillips curve ζ_p. Since our conclusions are most sensitive to these two parameters, we illustrate a range of outcomes corresponding to a large joint support for $\{\kappa, \zeta_p\}$.

For reference, Ajello et al. (2020) estimate $\zeta_p \approx 0.02$; given this estimate, and given all other parameter values, a durable adjustment cost of $\kappa \approx 0.25$ matches the spending shares for average U.S. recessions displayed in the right panel of Figure 3.[19] To illustrate the robustness and quantitative significance of the pent-up demand logic, we consider a range of outcomes for $\zeta_p \in (0, 0.1)$ and $\kappa \in (0, 0.5)$.

Results. For any given parameterization of our economy, we can (i) compute CIRs for a common demand shock b_t^c, changing only ϕ, and (ii) compute CIRs for a vector of demand shocks $\{b_t^c, b_t^s, b_t^d\}$ generating any given sectoral incidence. While Figure 4 used a single shift-share for several possible shares ϕ and shock combinations $\{b_t^c, b_t^s, b_t^d\}$, we here instead use a large range of possible models to estimate a single counterfactual in (i) and (ii). In particular, we compute CIRs for (i) common 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 the COVID-19 recession. Results are displayed in Figure 5.

Both panels show that — across the entire parameter range that we entertain — recessions more biased towards services, either because of the economy's sectoral make-up or because of shock incidence, induce weaker recoveries. Quantitatively, around our preferred estimates of $\zeta_p = 0.02$ and $\kappa = 0.25$ (marked with the red cross), the results align remarkably well with those of the semi-structural shift-share. The discussion in Appendix B.1

[19] Formally, we consider an economy subject only to the common shock b_t^c, and compute the share of output fluctuations across business-cycle frequencies (i.e., 6 to 32 quarters) attributable to durables and services spending. We set this share to match the share in Figure 3.

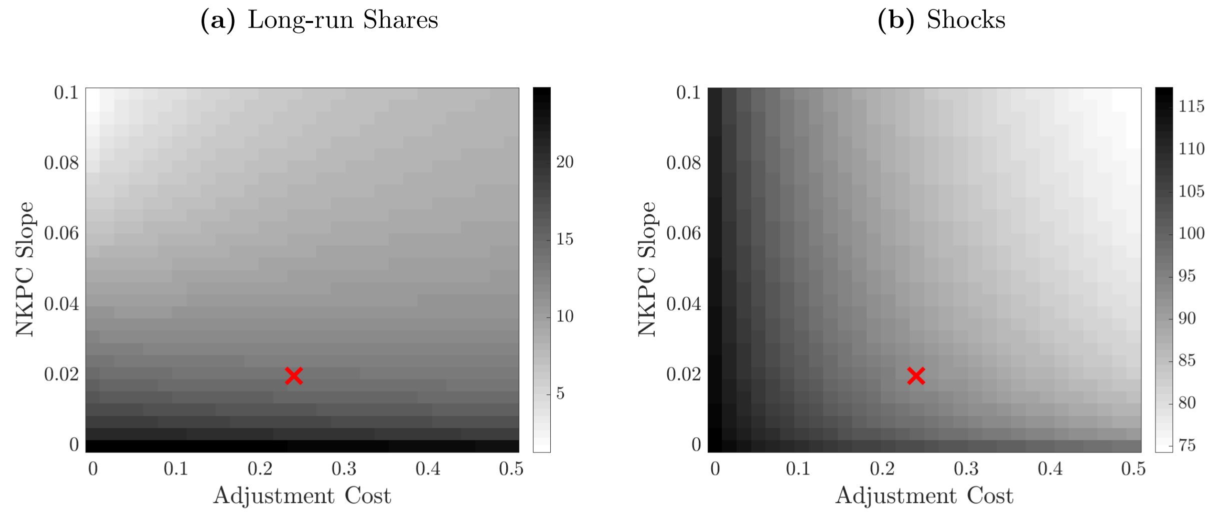

Figure 5: Left panel: Percentage gap between the CIR to a common demand shock b_0^c in an economy with the U.S. vs. Canada long-run expenditure shares, as a function of adjustment costs (x-axis) and the NKPC slope (y-axis). Right panel: Percentage gap between the CIR to demand shocks (b_0^c, b_0^d, b_0^s) 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.

helps to shed further light on those quantitative findings. We establish two results. First, for adjustment costs κ, we show that the condition $\theta_d < 1 - \delta$ — that is, pent-up demand effects outweighing adjustment costs — holds if and only if a common demand shock b_t^c moves durables expenditure by more than services expenditure. This condition is naturally satisfied in any sensible model calibration, explaining why pent-up demand effects remain dominant across the parameter range for κ entertained in Figure 5. Second, for the special case of flexible wages, we show that the normalized CIR of output can be written as

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

where θ_s is the response of services consumption to past changes in the durables stock $\widehat{d}_{t-1}$. For the wide range of parameterizations we consider, it turns out that $\mathbf{s}^c$ is always above $\mathbf{e}^c$ — consistent with the evidence in Section 3 — and that θ_s is relatively small (or even negative in some cases). Therefore, while the condition in Theorem 1 is not strictly speaking sufficient, $\mathbf{s}^c$ is sufficiently above $\mathbf{e}^c$ under the considered parameterizations so that the claim in Theorem 1 on the effects of demand composition on recovery strength still goes through.

5 Policy Implications

We have argued that the sectoral expenditure composition during demand-driven recessions is likely to have a large effect on recovery dynamics. Our conclusions so far, however, were conditional on a given monetary policy rule. In this section we explore the implications of pent-up demand and expenditure composition for the conduct of *optimal* stabilization policy.

5.1 Optimal Policy under Aggregate Shocks

We consider the general framework of Section 2.1. For now, however, we restrict the model to feature only *aggregate* demand shocks b_t^c, and rule out any sectoral shocks b_t^s or b_t^d. In this setting, the flexible-price allocation — and so the first-best policy — is straightforward to characterize.

Proposition 5. *Consider the model of Section 2.1 with $\gamma = \zeta$, simplified to feature only shocks to aggregate demand b_t^c. Then the first-best monetary policy sets*

$$\widehat{r}_t \quad = \quad (1 - \rho_b) b_t^c \tag{17}$$

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

The intuition is simple: changes in the durables share $1 - \phi$ affect the transmission of *both* common demand shocks b_t^c and conventional interest rate policy. With our definition of a common aggregate demand shock b_t^c, these two effects exactly offset, leaving optimal monetary policy as a function of b_t^c 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 also 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, in our setting, there is no such trade-off in *optimal* policymaking: while interest rate cuts today indeed lead to deficient demand tomorrow, negative fundamental shocks today lead to excess demand tomorrow, thus overall leaving optimal policy unaffected.

5.2 Optimal Policy under Sectoral Shocks

We now return to the full model of Section 2.1, again allowing for sectoral demand shocks. In this general setting, optimal monetary policy depends on the sectoral incidence of shocks. We state our main result for the special case of transitory shocks ($\rho_b = 0$) and no adjustment costs ($\kappa = 0$); numerical explorations, however, reveal that the result also holds for our quantitative model in Section 4.3, evaluated at the parameters of Table 4.1 and for all $\kappa \in [0, 0.5]$, as in Figure 5.

Proposition 6. *Consider the model of Section 2.1 with $\gamma = \zeta$, $\kappa = 0$ and $\rho_b = 0$, and let $\widehat{r}_t(b_0^i)$ 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 $\widehat{r}_0(b_0^s) = \widehat{r}_0(b_0^d) < 0$, we have*

$$\widehat{r}_t(b_0^s) \quad < \quad \widehat{r}_t(b_0^d), \quad \forall t \geq 2 \tag{18}$$

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

Without monetary accommodation, a services demand shock leads to a persistent recession, while a durables demand shock leads to a relatively short-lived contraction. If the monetary authority cuts nominal rates in the face of such sectoral demand shocks, it invariably stimulates the initially unaffected sector. Proposition 6 reveals what this stimulus — written in terms of equilibrium real rates — should look like: persistent in the case of a recession biased towards services, and short-lived after a durables-led contraction. Intuitively, given a negative one-off services shock, the monetary authority optimally 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 one-off durables shock, future real interest rates are relatively high to depress services expenditures and allow the durables stock to be re-built gradually.

Figure 6 provides a graphical illustration, displaying optimal nominal interest rate paths in response to aggregate and sectoral demand shocks, all normalized to give an initial rate response of -1 per cent. Consistent with our results in Proposition 5, the blue line (for the common shock b_t^c) is simply given as

$$\widehat{r}_t(b_0^c) = -\rho_b^t$$

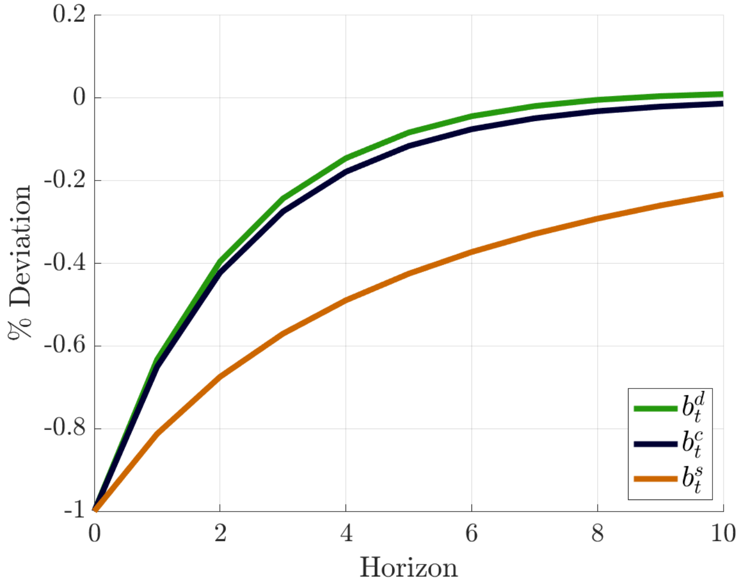

Figure 6: Optimal monetary policy following aggregate and sectoral demand shocks in the structural model of Section 2.1, with all shocks normalized to give $\widehat{r}_t(b_0^i) = -1$. For details on the model parameterization see Appendix A.1.

For the two sectoral shocks we instead have

$$\widehat{r}_t(b_0^s) = -\rho_b^t - \zeta_s \sum_{q=0}^{t-1} \rho_b^{t-q} \vartheta^q \tag{19}$$

$$\widehat{r}_t(b_0^d) = -\rho_b^t + \zeta_d \sum_{q=0}^{t-1} \rho_b^{t-q} \vartheta^q \tag{20}$$

where the parameters $\{\zeta_s, \zeta_d, \vartheta\}$ are functions of primitive model parameters, and (19) - (20) hold even in a model with adjustment costs $\kappa > 0$ and persistent shocks $\rho_b > 0$. In the special case covered by Proposition 6 we can prove that $\{\zeta_s, \zeta_d, \vartheta\}$ are all strictly positive, establishing the desired result; numerically, we find that they remain positive for the values of shock persistence and adjustment costs entertained in Section 4.3. In both cases it follows that, relative to the baseline equilibrium rate of interest for common demand shocks, the Wicksellian rate paths for pure services and durables demand shocks are tilted down and up, respectively. Consistent with our results on large effects of demand composition on recovery dynamics, Figure 6 reveals that, in our preferred model calibration, the differences in implied interest rate paths — here, the green line versus the orange line — are large.

6 Conclusions

We have argued that recoveries from demand-driven recessions with expenditure cuts concentrated in services tend to be weaker than recoveries from recessions biased towards durables. This prediction follows from standard consumer theory together with output being demand-determined, and we have documented strong empirical support for its key testable implication in aggregate U.S. time series data.

Our quantitative analysis suggests that the effect of expenditure composition on recovery strength can be meaningful, in particular for a recession as services-led as the ongoing COVID-19 pandemic. Moving from positive to normative analysis we also show that, if a policymaker were to ignore the sectoral incidence of shocks and instead applied a simple one-size-fits-all policy to all recessions, then monetary easing in services recessions would be too short-lived, and output would remain depressed for longer.

References

Ajello, A., Cairó, I., Cúrdia, V., Lubik, T., & Queralto, A. (2020). Monetary policy tradeoffs and the federal reserve's dual mandate.

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.

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'o, 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.

Bilbiie, F. O. (2019). The new keynesian cross. *Journal of Monetary Economics*.

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.

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.

Coibion, O., Gorodnichenko, Y., & Koustas, D. (2013). Amerisclerosis? the puzzle of rising us unemployment persistence. *Brookings Papers on Economic Activity*, 226.

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.

Fernald, J. G., Hall, R. E., Stock, J. H., & Watson, M. W. (2017). The disappointing recovery of output after 2009. *Brookings Papers on Economic Activity*, 1.

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.

House, C. L. & Shapiro, M. D. (2008). Temporary investment tax incentives: Theory with evidence from bonus depreciation. *American Economic Review*, *98*(3), 737–68.

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.

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. (2020). Local projections and vars estimate the same impulse responses. *Working paper*, *1*.

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. *Unpublished manuscript*.

Schmitt-Grohé, S. & Uribe, M. (2017). Liquidity traps and jobless recoveries. *American Economic Journal: Macroeconomics*, *9*(1), 165–204.

Wolf, C. K. (2020). SVAR (Mis-)Identification and the Real Effects of Monetary Policy Shocks. *American Economic Journal: Macroeconomics*. Forthcoming.

Woodford, M. (2011). *Interest and prices: Foundations of a theory of monetary policy*. princeton university press.

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, and the scaling factors $\{\alpha, \varsigma_c, \varsigma_s, \varsigma_d\}$ in our specification of household preferences.

From the steady-state first-order conditions, we get

$$\left(\frac{\tilde{\phi}}{1-\tilde{\phi}}\right)^{\zeta} = \frac{1}{1-\beta(1-\delta)}\left(\frac{\phi}{\frac{1}{\delta}(1-\phi)}\right)^{\zeta} \tag{A.1}$$

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

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

$$\varsigma_c \equiv \frac{1+\frac{\zeta-\gamma}{1-\zeta}}{1-\frac{\zeta-\gamma}{1-\zeta}\frac{\frac{1}{\delta}(1-\phi)\beta(1-\delta)(1-\rho_b)}{\phi+(1-\beta(1-\delta))\frac{1}{\delta}(1-\phi)}} \tag{A.3}$$

Note that, in the separable case $\gamma = \zeta$ considered in much of this paper, these expressions simplify to $\alpha = \frac{1-\beta(1-\delta)\rho_b}{1-\beta(1-\delta)}$ and $\varsigma_c = 1$. Next, we set $\{\varsigma_s, \varsigma_d\}$ to ensure that a combination of sectoral shocks $b_t^s = b_t^d$ is isomorphic to an aggregate demand shock b_t^c of the same size:

$$\varsigma_s \equiv \varsigma_c \phi \tag{A.4}$$

$$\varsigma_d \equiv \varsigma_c(1-\phi) \tag{A.5}$$

We note that this choice of $\{\varsigma_s, \varsigma_d\}$ also implies 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.

For future reference, it will be useful to let

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

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

$$\begin{aligned} \widehat{c}_t &= \frac{\phi}{\phi + [1-\beta(1-\delta)]\frac{1}{\delta}(1-\phi)} \left(\widehat{s}_t + \frac{1}{1-\zeta} \left(b_t^c + b_t^s \right) \right) \\ &\quad + \frac{[1-\beta(1-\delta)]\frac{1}{\delta}(1-\phi)}{\phi + [1-\beta(1-\delta)]\frac{1}{\delta}(1-\phi)} \left(\widehat{d}_t + \frac{\alpha}{1-\zeta} \left(b_t^c + b_t^d \right) \right) \end{aligned} \tag{A.6}$$

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

$$\widehat{\lambda}_t = \widehat{r}_t^n - \mathbb{E}_t \left[\widehat{\pi}_{t+1} \right] + \mathbb{E}_t \left[\widehat{\lambda}_{t+1} \right] \tag{A.7}$$

Given the scaling factors define above, we can write the first-order conditions for services and durables as

$$(\zeta - \gamma)\widehat{c}_t - \zeta \widehat{s}_t + (b_t^c + b_t^s) = \widehat{\lambda}_t \tag{A.8}$$

$$\begin{aligned} (\zeta - \gamma)\widehat{c}_t - \zeta \widehat{d}_t + \alpha(b_t^c + b_t^d) &= \frac{1}{1-\beta(1-\delta)} \left[\widehat{\lambda}_t + \kappa(\widehat{d}_t - \widehat{d}_{t-1}) \right] \\ &\quad - \frac{\beta(1-\delta)}{1-\beta(1-\delta)} \mathbb{E}_t \left[\widehat{\lambda}_{t+1} + \frac{\kappa}{1-\delta} (\widehat{d}_{t+1} - \widehat{d}_t) \right] \end{aligned} \tag{A.9}$$

Note that, in our baseline case of $\gamma = \zeta$, we can re-write those conditions as

$$-\gamma \widehat{s}_t = \widehat{\lambda}_t - (b_t^c + b_t^s), \tag{A.10}$$

$$\begin{aligned} -\gamma \widehat{d}_t &= \frac{1}{1-\beta(1-\delta)} \left[\widehat{\lambda}_t - (b_t^c + b_t^d) + \kappa(\widehat{d}_t - \widehat{d}_{t-1}) \right] \\ &\quad - \frac{\beta(1-\delta)}{1-\beta(1-\delta)} \mathbb{E}_t \left[\widehat{\lambda}_{t+1} - (b_{t+1}^c + b_{t+1}^d) + \frac{\kappa}{1-\delta} (\widehat{d}_{t+1} - \widehat{d}_t) \right] \end{aligned} \tag{A.11}$$

where we have used that $\mathbb{E}_t(b_{t+1}^c) = \rho_b b_t^c$ 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 common aggregate shock can be perfectly offset by movements in real interest rates (via (A.7)), ensuring the desirable neutrality property discussed in Footnote 4.

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 $\{b_t^c, b_t^s, b_t^d\}$:

$$\widehat{\pi}_t^w = \frac{(1-\beta\phi_w)(1-\phi_w)}{\phi_w(\frac{\varepsilon_w}{\varphi}+1)}\left[\frac{1}{\varphi}\widehat{\ell}_t - \left(\widehat{w}_t + \widehat{\lambda}_t - (\varsigma_c b_t^c + \varsigma_s b_t^s + \varsigma_d b_t^d)\right)\right] + \beta\mathbb{E}_t\left[\widehat{\pi}_{t+1}^w\right] \qquad \text{(A.12)}$$

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 $\widehat{\ell}_t$:

$$\widehat{\pi}_t = \zeta_p\left(\widehat{w}_t - \frac{y''(\ell)\ell}{y'(\ell)}\widehat{\ell}_t\right) + \beta\mathbb{E}_t\left[\widehat{\pi}_{t+1}\right] \qquad \text{(A.13)}$$

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.13) 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 2 we set $\gamma = \zeta = 1$, $\beta = 0.99$, $\delta = 0.021$, $\rho_b = 0$, $\kappa = 0$ and $\underline{\phi} = 0.9$. We then choose $\bar{\phi}$ to construct a recession with equal shares for durables and services spending. For the policy exercise in Figure 6, we largely keep the same parameters, but set $\kappa = 0.25$ and $\rho_b = 0.65$.

A.2 Further Extensions

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 Appendix B.2. All other model blocks are unaffected by the presence of hand-to-mouth households.

Many sectors. Household preferences over consumption bundles are now given as

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

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

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

and finally the linearized output market-clearing condition is

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

All other model equations are unchanged.

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 $\{z_t^c, z_t^s, z_t^d\}$ 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

$$\widehat{p}_t^s = -(z_t^c + z_t^s) \tag{A.14}$$
$$\widehat{p}_t^d = -(z_t^c + z_t^d) \tag{A.15}$$

The output market-clearing condition then becomes

$$\widehat{y}_t = \left[z_t^c + \phi z_t^s + (1-\phi) z_t^d\right] + \widehat{\ell}_t = \phi \widehat{s}_t + (1-\phi)\widehat{e}_t \tag{A.16}$$

All other model equations are unchanged.

Alternative specification for b_t^c. A natural alternative specification for household consumption preferences is

$$u(s,d;b) = \frac{e^{b^c}\left[e^{b^s}\tilde{\phi}^{\zeta} s^{1-\zeta} + e^{b^d}(1-\tilde{\phi})^{\zeta} d^{1-\zeta}\right]^{\frac{1-\gamma}{1-\zeta}} - 1}{1-\gamma} \tag{A.17}$$

Here, the common shock b^c affects the valuation of the total consumption bundle in one period relative to the next. With this specification, (A.7) still applies without change, but the first-order conditions for services and durables become

$$(\zeta-\gamma)\widehat{c}_t - \zeta \widehat{s}_t + (b_t^c + b_t^s) = \widehat{\lambda}_t \tag{A.18}$$
$$\begin{aligned}(\zeta-\gamma)\widehat{c}_t - \zeta \widehat{d}_t + (b_t^c + b_t^d) &= \frac{1}{1-\beta(1-\delta)}\left[\widehat{\lambda}_t + \kappa(\widehat{d}_t - \widehat{d}_{t-1})\right] \\ &\quad - \frac{\beta(1-\delta)}{1-\beta(1-\delta)}\mathbb{E}_t\left[\widehat{\lambda}_{t+1} + \frac{\kappa}{1-\delta}(\widehat{d}_{t+1} - \widehat{d}_t)\right]\end{aligned} \tag{A.19}$$

Note that the two sectoral shocks $\{b_t^s, b_t^d\}$ enter exactly as in our baseline system (up to scale). The common aggregate demand shock b_t^c 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 extend the analysis in Section 2.4 on equilibrium characterization in the baseline model. In Appendix B.2 we then state formal results for the model extensions in Section 2.5. Finally, in Appendix B.3, we provide several further robustness checks for the quantitative model-based exercises in Section 4.3.

B.1 Baseline Model

We present two additional results: on the parameter θ_d in the model with adjustment costs and on equilibrium characterization in a model with partially sticky prices, a monetary policy rule as in (2) and arbitrary non-separability in preferences.

θ_d vs. $1-\delta$. Our quantitative explorations in 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.1. *Consider the model of Section 2.4. If $\theta_d = 1-\delta$, then*

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

The durables share after a common demand shock, $\frac{(1-\phi)\widehat{e}_0^c}{\phi\widehat{s}_0^c+(1-\phi)\widehat{e}_0^c}$, 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.[20]

Imperfectly sticky prices and non-separability. First, we consider the generalization of Proposition 2 to an economy with non-separability in household preferences as well

[20]The relative unconditional volatility documented in the bottom panel of Figure 3 suffices as evidence under the assumption that business cycles are largely driven by common aggregate fluctuations. A stronger test looks at relative volatilities *conditional* on a particular aggregate shock, e.g. to monetary policy. It is well-known that, even conditional on such common shocks, durables spending is much more volatile than non-durables spending (e.g. Christiano et al., 1999).

as imperfectly sticky prices (and a monetary policy rule which does not fix the real rate), but with wages being perfectly flexible.

Proposition B.2. *Consider the model of Section 2.1 with arbitrary degrees of price stickiness and non-separability (governed by ζ_p and $\{\zeta, \gamma\}$) but with flexible nominal wages ($\phi_w = 0$). Let $\{b_0^c, b_0^s, b_0^d\}$ be a vector of time-0 shocks resulting in services share ω. Then the normalized cumulative impulse response of aggregate output satisfies*

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

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

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

Second, we have the following generalization of Theorem 1

Theorem B.1. *Let $\widehat{\mathbf{s}}^c$ and $\widehat{\mathbf{e}}^c$ denote the normalized cumulative impulse responses of services and durables expenditure to a recessionary common demand shock $b_0^c < 0$, defined as in (8), resulting in services share ω^c.*

Then, the normalized cumulative impulse response of aggregate output $\widehat{\mathbf{y}}$ in Proposition B.2 is increasing in the services share ω if and only if

$$\widehat{\mathbf{s}}^c \quad > \quad \widehat{\mathbf{e}}^c\left(1+\frac{\phi}{1-\phi}\frac{1}{\omega^c}\theta_s\right)$$

Intuitively, in the case of net substitutability, we expect that $\theta_s < 0$ because an elevated lagged durables stock should depress current services spending, ensuring that (14) remains sufficient. With a conventional monetary rule and sticky prices, we expect endogenous persistence and co-movement, so that $\theta_s > 0$, implying (14) is necessary but not sufficient. Our numerical simulations in Section 4.3 confirm these intuitions, but reveal θ_s to be small.

B.2 Further Extensions

We here state formal propositions for the various model extensions discussed in Section 2.5 and for an alternative specification of the aggregate demand shock b_t^c.

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.3. *Consider the extended model with hand-to-mouth consumers, and suppose that that the monetary authority fixes the real rate of interest and that $\gamma = \zeta$. 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 2 and Theorem 1 apply without change.

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

Proposition B.4. *Consider the extended model with N sectors, and suppose that that the monetary authority fixes the real rate of interest and that $\gamma = \zeta$. Consider an arbitrary shock mix $\{b_0^c, \{b_0^i\}_{i=1}^N\}$ with sectoral spending shares ω_i. Then*

$$\widehat{\mathbf{y}} \quad = \quad -\sum_{i=1}^{N} \omega_i \frac{\delta_i}{1-\theta_d^i} \quad = \quad -\sum_{i=1}^{N} \omega_i \mathbf{e}_i^c \tag{B.1}$$

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

GENERAL ADJUSTMENT COSTS. Our results on the equivalence between CIR rankings and the effects of spending composition on recovery dynamics go through without change in a model with arbitrary adjustment costs.

Proposition B.5. *Consider the extended model with arbitrary adjustment costs. Suppose that the monetary authority fixes the real rate of interest and that $\gamma = \zeta$, and consider an arbitrary shock vector $\{b_0^c, b_0^s, b_0^d\}$ with services share ω. Then the normalized cumulative impulse response of aggregate output satisfies*

$$\widehat{\mathbf{y}} \quad = \quad \omega \mathbf{s}^c \;+\; (1-\omega)\mathbf{e}^c$$

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

$$\mathbf{s}^c \quad > \quad \mathbf{e}^c$$

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

Alternative specification for b_t^c. Many — but not all — of our main results extend to the alternative preference specification (A.17). For simplicity, we restrict our discussion here to the separable case ($\zeta = \gamma$).

Consider first the equilibrium characterizations and recession recovery results in Sections 2.2 to 2.4. Comparing (A.18)-(A.19) to (A.10)-(A.11), 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 $\widehat{s}_t$ and $\widehat{e}_t$ to any shock tuple $\{b_0^c, b_0^d, b_0^s\}$ are unaffected *up to scale*. Intuitively, the new preference specification is isomorphic to a rescaling of the shocks b_t^c and b_t^d in (A.11); since Propositions 1 and 2 condition on sectoral spending declines, they (as well as Theorem 1) are entirely unaffected by the change in preferences.[21] 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.17). 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.17), a monetary shock m_t with persistence ρ_m is now equivalent (up to scale) to a mixture of sectoral demand shocks

$$b_t^s + \frac{1 - \beta(1-\delta)\rho_b}{1 - \beta(1-\delta)} b_t^d \tag{B.2}$$

The weight on b_t^d is simply equal to our preference scaling parameter α, exactly as expected.[22]

[21] The analogous expressions for Lemmas 1 and 2 change, however, as a common shock b_t^c is now a different weighted average of pure sectoral shocks.

[22] In the more general case of partially flexible prices, the sectoral labor disutility parameters $\{\varsigma_s, \varsigma_d\}$ would need to additionally satisfy $\varsigma_c = \phi$, $\varsigma_d = \frac{1}{\alpha}(1-\phi)$ to ensure equivalence.

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.2), 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.

B.3 Quantitative Analysis

In Section 4.3 we study the effect of pent-up demand on expected recovery dynamics as a function of the strength of adjustment costs and price stickiness, simply because these are the model features mostly likely to neutralize the pent-up demand logic. We here complement these findings by providing an analogous plot for output CIRs as a function of shock persistence (ρ_b) and the monetary authority's policy rule (ϕ_π). All other parameters are fixed as in Table 4.1, and we further set $\zeta = 0.02$ and $\kappa = 0.25$, in line with the discussion in Section 4.3.

Results are displayed in Figure B.1. The key take-away is that pent-up demand effects remain quite strong throughout: for heterogeneous shocks, the counterfactual causal effect is robustly large, while for a common shock in heterogeneous economies the effect only vanishes for a very aggressive monetary rule in conjunction with highly transitory shocks.

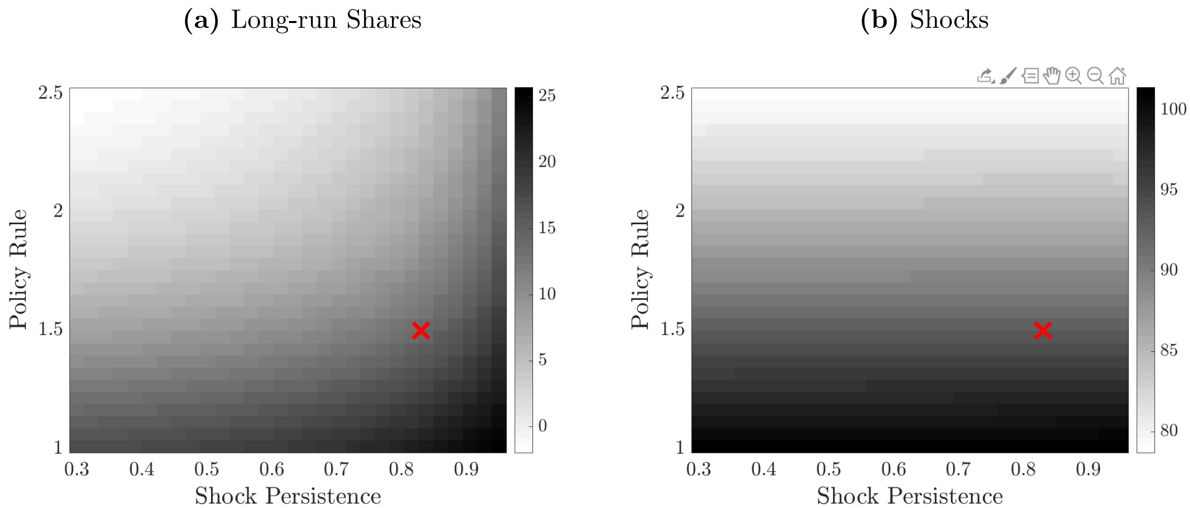

Figure B.1: Left panel: Percentage gap between the CIR to a common demand shock b_0^c in an economy with the U.S. vs. Canada long-run expenditure shares, as a function of shock persistence (x-axis) and the Taylor rule coefficient (y-axis). Right panel: Percentage gap between the CIR to demand shocks (b_0^c, 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 shock persistence (x-axis) and the Taylor rule coefficient (y-axis). The red cross in both figures indicates our preferred parameterization.

C Empirical Appendix

This appendix provides further details for the empirical exercises in Section 3.

C.1 Monetary Policy

We estimate a recursive VAR in a sectoral measure of consumption, aggregate consumption, aggregate GDP (all real), the GDP deflator, and the federal funds rate, in this order. We consider three specifications, changing the sectoral measure of consumption from durables to non-durables to services. All series are taken from the St. Louis Fed's FRED database.

Our three VARs are estimated on a quarterly sample from 1960:Q1 — 2007:Q4, with four lags, a constant and a linear time trend, and with a 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 CIR difference of durables- and services-led recession, we estimate a single VAR containing all consumption series.

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 (2020), 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.[23] 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.1 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.

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

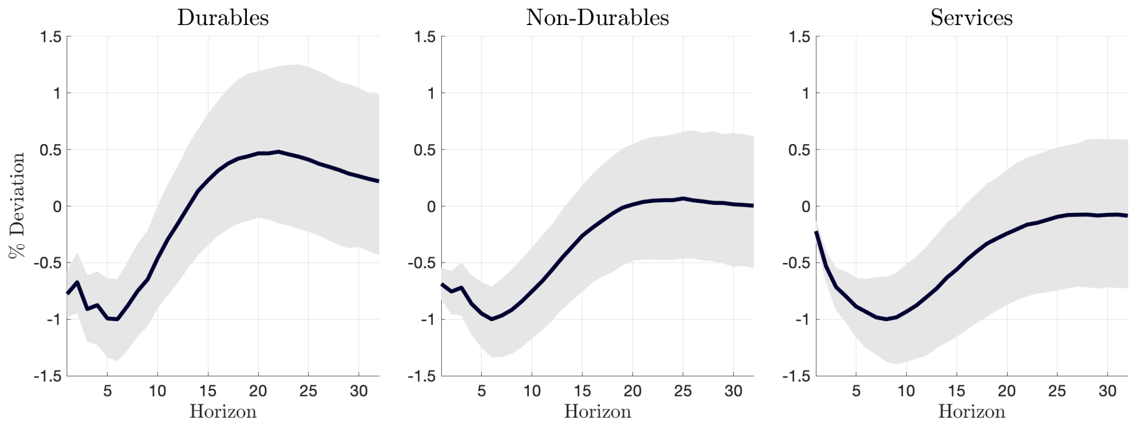

Figure C.1: Quarterly impulse responses to an uncertainty shock (à la Basu & Bundick (2017)) by consumption spending category, all normalized to drop by -1% at the trough. The solid blue line is the posterior mean, while the shaded areas indicate 16th and 84th percentiles of the posterior distribution,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

Figure C.2 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.

(a) Long-run Shares **(b)** Shocks

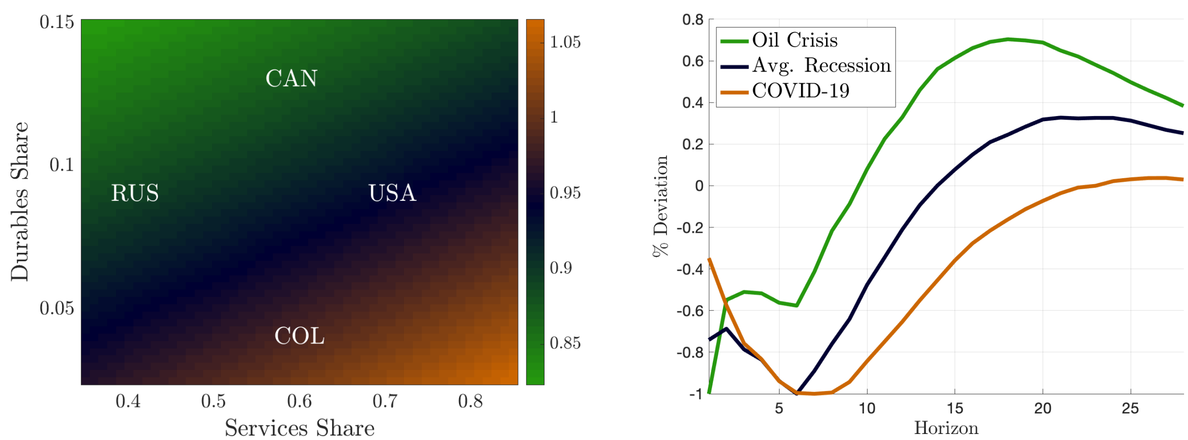

Figure C.2: Left panel: CIR to a common demand shock b_0^c as a function of long-run expenditure shares, with the U.S. CIR normalized to 1, computed using the posterior mode point estimates from Figure 1. Right panel: Impulse response of total consumption to sectoral demand shocks reproducing expenditure composition changes in (i) ordinary recessions, (ii) the 1973 oil crisis, and (iii) the COVID-19 recession, all normalized to lead to a peak-to-trough consumption contraction of -1 per cent and evaluated again using the posterior mode point estimates from Figure C.1.

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 include several granular measures of sectoral consumption. The results from a subset of our experiments are reported in Figure C.3. 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, clothes and financial services all decline in the initial recession, but then only recover gradually and without much of an overshoot.

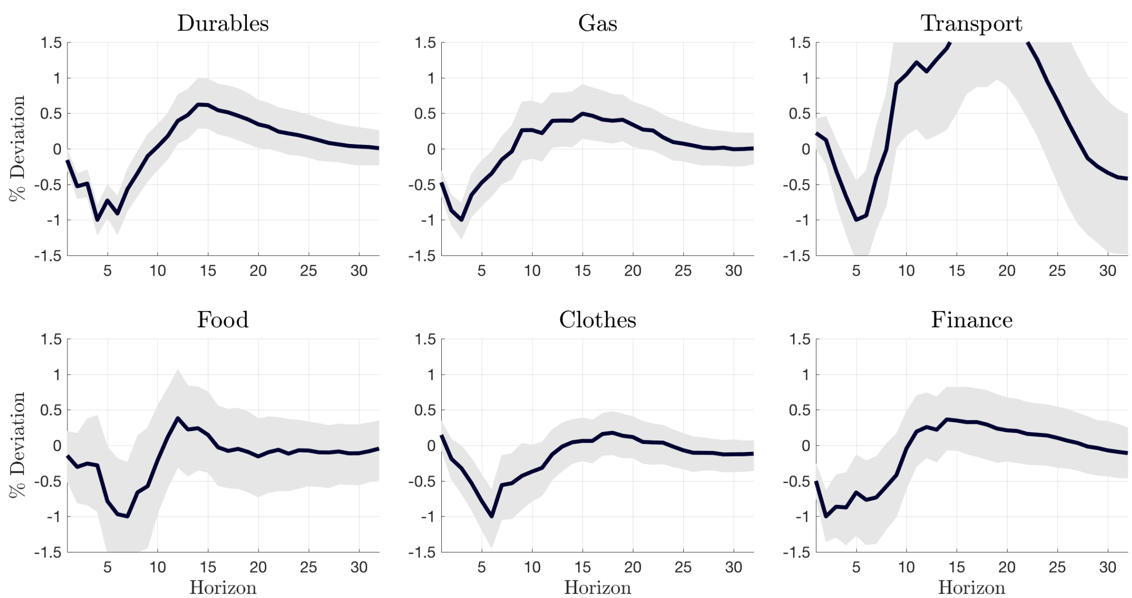

Figure C.3: Quarterly impulse responses to an oil shock (à la Hamilton (2003)) by consumption spending category, all normalized to drop by -1% at the trough. The solid blue line is the posterior mean, while the shaded areas indicate 16th and 84th percentiles of the posterior distribution, respectively.

C.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 CIRs of total consumption to each of the three reduced-form Wold innovations, with the impact consumption response normalized to 1.[24]

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

[24]For 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.

D Proofs

D.1 Proof of Lemmas 1 and 2

Under the simplifying assumption (iv) of fixed prices 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.6), (A.8) and (A.9).

Then, durable expenditures are given by

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

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

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

Under the additional restrictions that $\zeta = \gamma$ and all shocks having common persistence ρ_b, we use the method of undetermined coefficients to obtain the dynamic responses

$$\widehat{s}_t = \frac{1}{\gamma}(b_t^c + b_t^s) \tag{D.3}$$

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

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 $b_t^c + b_t^s = (b_0^c + b_0^s)\rho_b^t$, the above directly implies the dynamic response for $\widehat{s}_t$ in Lemma 2.

Moreover, iterating $\widehat{d}_t$ forward, we obtain

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

and thus we obtain the dynamic response for $\widehat{e}_t$ in Lemma 2

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

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

$\theta_b = \frac{1}{1-\beta(1-\delta)}$. □

D.2 Proof of Propositions 1 and 2

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

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

Using the dynamic responses in Lemma 2, we have that

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

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

Thus, replacing above, we obtain the expression for $\widehat{\mathbf{y}}$ in equation (13) of Proposition 2

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

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

D.3 Proof of Theorem 1

Firs, note that $\widehat{\mathbf{y}}$ in equation (13) of Proposition 2 is decreasing in ω if and only if $\delta < 1 - \theta_d$.

Second, using the dynamic responses in Lemma 2, we have that

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

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

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

D.4 Proof of Proposition 3

Plugging the policy rule (2) into the bond FOC (A.7), we get

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

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

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

Similarly, the services and durables FOCs become

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

and the wage-NKPC becomes

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

Therefore, for shocks m_t with persistence ρ_m and size m_0, these equations are identical to those in a model with common demand shocks b^c_t with persistence $\rho_b = \rho_m$ and size $b^c_0 = -\frac{1}{\varsigma_c}\frac{1}{1-\rho_m} m_0$, completing the argument.[25] □

D.5 Proof of Proposition 4

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

1. We know from the expressions in Lemma 2 that the impulse responses $\widehat{s}^c_t$ and $\widehat{e}^c_t$ are independent of ϕ. The claim then follows immediately from the market-clearing con-

[25] In the separable case, with $\varsigma_c = 1$, the conclusion is immediate by comparing to (A.10) - (A.11). In the general case, our definitions of $\{\alpha, \varsigma_c\}$ are precisely so that we recover (A.8) - (A.9).

dition

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

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

2. Since services spending impulse responses to any shock vector (b_0^c, b_0^s, b_0^d) scale in $b_0^c + b_0^s$, while durables spending impulse responses scale in $b_0^c + b_0^d$, it follows that $\widehat{s}_t \propto \widehat{s}_t^c$ and $\widehat{e}_t \propto \widehat{e}_t^c$. 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

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

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

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

replicating the flexible-price allocation. $\square$

D.7 Proof of Proposition 6

In the flexible-price analogue of the economy of Section 2.1 with $\gamma = \zeta$, 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:

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

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. □

D.8 Proof of Proposition B.1

Set $\theta_d = 1 - \delta$ in the expressions for θ_b, θ_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 a common 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. $\square$

D.9 Proof of Proposition B.2 and Theorem B.1

Consider the equations characterizing the equilibrium in the baseline model in Appendix A.1 for arbitrary degrees of price stickiness and non-separability (governed by ξ and $\{\zeta, \gamma\}$) but with 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

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

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

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

Moreover, we have that

$$\mathbf{s}^c = \frac{1}{1-\rho_b} + \frac{\phi}{1-\phi}\frac{1-\omega^c}{\omega^c}\theta_s\mathbf{e}^c$$
$$\mathbf{e}^c = \frac{\delta}{1-\theta_d}\frac{1}{1-\rho_b}$$

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

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

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

D.10 Proof of Proposition B.3

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

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

Aggregating across constrained households H and unconstrained households R:

$$\widehat{s}_t \quad = \quad (1-\mu)\widehat{s}_t^R + \mu\widehat{s}_t^H$$
$$\widehat{e}_t \quad = \quad (1-\mu)\widehat{e}_t^R + \mu\widehat{e}_t^H$$

This set of equations completes the equilibrium characterization.

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 (2019), 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.11 Proof of Proposition B.4

We now for each good i get the optimality condition

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

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

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

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 2 and Proposition 2 can proceed exactly as in the baseline case. □

D.12 Proof of Proposition B.5

For the output CIR to *any* shock combination $\{b_0^c, b_0^s, b_0^d\}$ to satisfy

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

it suffices to show that (i) $\widehat{s}_t$ depends only on $b_0^c + b_0^s$ and that (ii) $\widehat{e}_t$ depends only on $b_0^c + b_0^d$. Since the real rate of interest is fixed, (i) follows from

$$-\gamma\widehat{s}_t = b_t^c + b_t^s$$

But since the FOC for d_t depends only on $b_t^c + b_t^d$, $b_{t+1}^c + b_{t+1}^d$ and $\{d_{t+\ell}\}_{\ell=-t}^{\infty}$, it follows that d_t and so e_t can similarly be obtained as a function of only $b_0^c + b_0^s$, so the result follows. □

D.13 Proof of Proposition B.6

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

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

$$-\frac{\beta(1-\delta)}{1-\beta(1-\delta)}\mathbb{E}_t\left[\widehat{\lambda}_{t+1}+\widehat{p}^d_{t+1}+\frac{\kappa}{1-\delta}(\widehat{d}_{t+1}-\widehat{d}_t)\right]$$

With $\widehat{\lambda}_t = 0$ from our assumptions on equilibrium selection, it follows immediately from (A.14)-(A.15) 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.16) to ensure overall output market clearing. It follows that all our main results apply without change to supply shocks. $\square$