Demand Shocks as Technology Shocks*

Yan Bai
University of Rochester and NBER

José-Víctor Ríos-Rull
University of Pennsylvania, CAERP, CEPR, and NBER

Kjetil Storesletten
University of Oslo and CEPR

May 2017

Abstract

We provide a macroeconomic theory where *demand* for goods has a productive role. A search friction prevents perfect matching between producers and potential customers. Larger demand induces more search, which in turn increases GDP. When viewed through the lens of a standard neoclassical aggregate production function which ignores search effort as an input factor, an increase in demand appears as an increase in the Solow residual. We embed the product market friction in an otherwise standard neoclassical model with flexible prices and shocks to preferences and technology. We interpret preference shocks as shocks to demand. The model is estimated using standard Bayesian techniques. Demand shocks account for a large share of the fluctuations in consumption, GDP, and the Solow residual.We provide a macroeconomic theory where *demand* for goods has a productive role. A search friction difficults matching between producers and customers. Larger demand induces more search, increasing GDP, which when viewed through the lens of the standard aggregate production function appears as an increase in the Solow residual. We embed the friction in an otherwise standard neoclassical model with shocks to preferences (interpreted as shocks to demand) and technology. Demand shocks account for a large share of consumption, GDP, and the Solow residual fluctuations when estimated by Bayesian methods. The dynamics of the search effort implied by the model matches the empirical fall in shopping time in the U.S. during the Great Recession. The dynamics of the search effort implied by the model matches the empirical fall in shopping time in the U.S. during the Great Recession.

*We thank the editor and three anonymous referees, as well as seminar attendants at the University of Pennsylvania, New York University, the Bank of Canada, European Central Bank, FRB New York, FRB Minneapolis, Institute for Fiscal Studies, University of Toulouse, and CREI-CEPR for useful comments. We are especially thankful to Robert Gordon, Gary Hansen, Ellen McGrattan, Larry Christiano, and Zhen Huo. Ríos-Rull thanks the National Science Foundation for Grant SES-1156228.

1 Introduction

In the standard neoclassical model, output is a function of inputs such as labor and capital. There is no explicit role for demand because potential consumers are always available and Walrasian prices adjust so that all produced goods become used.In reality, customers and producers must meet in order for the produced good to be consumed, so value added depends on how well they are matched. As an example, consider a restaurant. According to the neoclassical view, the value added of a restaurant should be a function of its inputs (employees, tables, etc.), irrespective of the number of patrons and how hungry they are. Moreover, the restaurant owner would set prices so that all tables were in use. However, actual production takes place only when customers show up. The more customers demand the restaurant's meals, the larger the value added will be. The idea that the demand for goods plays a direct role extends to many forms of production: dentists need patients, car dealers need shoppers, all producers need buyers.

This paper provides a theory where search for goods – which we with some abuse of terminology refer to as *demand* – has a productive role. The starting point is that customers search for producers, and a standard matching friction prevents Walrasian market clearing, in the sense that all potential productive capacity necessarily translates into actual value added.[1]

Allowing an explicit role for demand has implications for business cycle analysis, especially for our understanding of the driving factors for business cycles. In our model changes in search effort affect output even if conventional inputs, and the intensity with which they are used, remain constant. A key purpose of this paper is to investigate how our notion of demand shocks works in terms of driving the Solow residual and to quantify how important this mechanism is for aggregate fluctuations, compared to more standard shocks in the business cycle literature.

To clarify our main argument, we start by analyzing a simple Lucas-tree version of our economy with search for goods and show that preference shocks generate fluctuations in measured TFP in this economy. Technically, we resolve the search friction by building on the competitive search model (Moen (1997)). To find goods for consumption, consumers have to exert costly search

[1] The National Income and Product Accounts (NIPA) measures consumption expenditure as value added only when goods are purchased by consumers. Indeed, a lost opportunity, such as an empty seat on a bus, is not value added. In the neoclassical model the true value of consumption – in money units – is the price times the quantity of consumption goods, consistent with the way value added is measured in NIPA. Note, however, that in the presence of a goods market search friction, a notion of value added that captures the benefits for the consumer would be different because obtaining goods requires a search cost. Therefore, price times quantity will underestimate such notion of value added.

effort. Firms post prices and customers trade off good prices versus congestion when searching for goods. In equilibrium firms charge higher prices in markets where they have a lower chance of realizing a sale and it is easier for customers to find goods. Search effort is determined by a familiar condition: the marginal rate of substitution between search costs and consumption (i.e., the ratio of marginal disutility of search to the marginal utility of consumption) must equal the marginal transformation of search, i.e., the marginal increase in the probability that a match will occur when search effort is increased marginally.

We then embed the search friction for consumption goods in an otherwise standard stochastic neoclassical growth model. To study the role of shocks to consumption demand – modeled as shopping effort preference shocks – we allow for the possibility that business cycles can be driven also by other shocks that are standard in the business cycle literature. In particular, we consider shocks that affect the marginal rate of substitution (MRS) between consumption and leisure, a standard economy-wide technology shock, and an investment-specific productivity shock (se e.g. Justiniano et al. (2010) and references therein).

Ceteris paribus, a preference shock that increases shopping effort will increase both consumption and the Solow residual. Firms respond in a similar way to a preference shock and to a true technology shock. Thus, if one ignores changes in search effort, consumption demand appears as if it were a positive externality in the economy. This mechanism reverses the direction of causation relative to the neoclassical model: there, a positive true technology shock would increase TFP and, in turn, consumption. Interestingly, aggregate data for the United States suggest that it is factors influencing consumption that drives TFP, rather than TFP driving consumption. In particular, a simple Granger test reveals that aggregate consumption turns out to Granger cause TFP, while TFP does not Granger cause aggregate consumption.[2]

The model is estimated using Bayesian estimation techniques targeting time series for output, labor productivity, investment, and the relative price of investment goods. This allows us to quantify the contribution of each shock to the variance of aggregate variables. The Bayesian estimation chooses shocks so as to maximize the probability of observing the realized aggregate time series. One of our main findings is that it is necessary to attribute a large role to demand shocks in order to match the data: demand shocks to consumption account for about half the variance of the Solow residual, 69 percent of the variance of consumption, and 29 percent of the variance of

[2] The Granger test was done with HP-filtered quarterly U.S. data, Q1:1967 to Q4:2013, with the Solow residual calculated by assuming a Cobb-Douglas production function with a labor share of 2/3. The p-value of TFP causing consumption is 0.34, while the p-value for consumption causing TFP is 0.002.

output. Moreover, the consumption demand shock accounts for 42 percent of the variance of the relative price of investment goods, which is substantially larger than the role of the investment-specific productivity shock for this variable. Thus, what may appear as technology shocks from the perspective of a standard neoclassical growth model do – according to our estimated model – to a large extent reflect fluctuations in resource utilization arising from consumers varying their shopping efforts for goods.

Our paper provides a model of endogenous shopping effort, and quantify the role of shocks to the disutility of shopping. Does our model make sense? One way to verify that the mechanism of our model is reasonable is to consider the model's predictions about shopping effort. As it turns out, the implied shopping effort in our estimated model is correlated with the *time* U.S. households devote to shopping. In particular, both the shopping effort in our model and the average empirical shopping time for U.S. households fell sharply from 2006 and onward (cf., the American Time Use Survey). In our model the fall in shopping effort after 2006 is mainly driven by the large and sustained fall in consumption and the weak development of the relative investment price during the Great Recession.

Our benchmark model features just-in-time production, in line with the restaurant example above. We pursue two important extensions to this framework in order to show that our findings are robust. First, note that if firms could store produced goods as inventories then demand shocks might – ceteris paribus – be expected to matter less for business cycles. However, one countervailing force is that households may also store these goods. To address this issue we extend the baseline model to incorporate storable goods in the form of firm inventories and a stock of consumer durables. Our main results hold up: demand shocks continue to be a major driving force for consumption, TFP, and GDP even when goods can be stored.[3] Moreover, our model replicates the empirical facts that purchases of durables are much more volatile than services and non-durable consumption goods and that the ratio of inventories-to-sales is countercyclical. These findings hinge on the assumption that goods can be stored by both households and firms.

The second extension pertains to capacity utilization of capital. Our exercise documenting how preference shocks can influence aggregate measured productivity is related to the well known fact that any mismeasurement of factor inputs will manifest itself as movements in the Solow residual. Indeed, in quantitative business cycle analysis it is standard to allow for varying capacity utilization

[3]When firms can store consumer durables, demand shocks account for one third of the variance of the Solow residual, half of the variance of consumption, and 20 percent of the variance of output.

of capital and labor.[4] For example, Wen (2004) and Nakajima (2005) argue that with variable capacity utilization, preference shocks that change the desired timing of consumption will cause changes in the utilization of input factors and, hence, generate fluctuations and cause changes in the measured Solow residual. This begs the question whether the significant role our estimation assigns to demand shocks through shopping utility preference shocks is simply a stand in for varying capacity utilization. To address this issue we extend our benchmark production model to allow for variable capacity utilization of capital along the lines of King and Rebelo (1999). We find that the quantitative role of the shopping preference shocks is robust to extending our model to allow for varying capacity utilization of capital in the sense that the demand shocks we consider continue to account for a significant share of the aggregate fluctuations.[5] Thus, we conclude that the mechanism for generating endogenous fluctuations in the variable capacity utilization literature is fundamentally different from that of the search-based theory we explore in this paper.

Our theory of matching frictions in product markets has far-reaching implications over and above the novel mechanism for how preference shocks generate productivity movements and propagate to the aggregate economy. In particular, our model provides a new theory of the relative prices of investment and consumption goods that is not based exclusively on exogenous technology shocks.[6] The model also embeds a theory of endogenous capacity utilization, based on lost goods and omitted factor inputs where the notion of utilization is the share of potential goods that are found. This is different from the earlier capacity utilization literature, where inputs are not being used at full capacity. We also provide an efficiency result: in our competitive search model where consumers search for firms and productions is endogenous, the equilibrium requires indexing markets not only by price and market tightness but also by the quantity of the good traded. Finally, our model provides a mechanism through which movements in search intensity affects asset prices.

Summing up, we see our paper as an implementation of Keynes' central idea that shocks to demand can give rise to business cycle fluctuations. This is done by extending a standard neoclassical framework with a search friction for goods. The role for demand is intrinsic to the process of production and is not arbitrarily imposed: markets clear and no agent has incentives to

[4]For example, Greenwood et al. (1988), Kydland and Prescott (1988), Bils and Cho (1994), Basu (1996), King and Rebelo (1999), Licandro and Puch (2000), and Francis and Ramey (2005) study variable capital utilization, and Burnside et al. (1993) introduce variable worker utilization in the form of labor hoarding during periods of low aggregate activity. In periods during which productivity and/or profits are high, firms will use the input factors more intensively, causing movements in the Solow residual that do not reflect true technology shocks.

[5]In the presence of variable capacity utilization we find that demand shocks account for half of the variance of the Solow residual, about two thirds of the variance of consumption, and 21 percent of the variance of output.

[6]See, for example, Greenwood et al. (1997), Krusell et al. (2000) and Fisher (2006) for papers that rely on exogenous technological shocks to derive changes in the relative price of investment.

deviate. There is a long tradition of attributing a role for demand in business cycle analysis, starting with Keynes' seminal contributions. However, in none of the earlier approaches has demand had a direct productive role.[7] We pursue neither the fixed-price tradition of the New Keynesian literature nor the coordination-problem tradition that sees a recession as a bad outcome within environments susceptible to multiple equilibria. Instead, we follow a tradition where a fall in demand generates a recession via infrautilization of productive capacity.

Our paper is also related to Petrosky-Nadeau and Wasmer (2015), which is developed independently from our paper. They also model costly search for goods in final goods markets although their focus is on how this search interacts with search in the labor market and influences the business cycle properties of the model. In particular, they do not focus on aggregate demand effects. Rudanko and Gourio (2014) study a business cycle model with a search friction in the market for consumption goods. Firms form long-lasting relationships with customers, and the authors focus on the role of customers as capital. Our contribution is also related to several papers emphasizing the effects of search frictions in shaping TFP (Alessandria (2005), Faig and Jerez (2005), and Lagos (2006)), although none of these focus on business cycles.

Some papers examine, as we do, how demand changes affect productivity and capacity utilization, although through very different mechanisms. In Fagnart et al. (1999), monopolistic firms with putty-clay technology are subject to idiosyncratic demand shocks, which causes fluctuations in capacity utilization. In Floetotto and Jaimovich (2008), changes in the number of firms cause changes in markups and, hence, changes in the measured Solow residual. Swanson (2006) shows that government expenditure shocks can increase aggregate output, consumption, and investment in a model with heterogeneous sectors. Finally, a number of papers build directly on our approach of consumption demand shocks in models with a search friction in product markets, including Huo and Ríos-Rull (2013), Duras (2015, 2016), and Bai and Ríos-Rull (2015).

The paper is organized as follows. Section 2 lays out the main mechanism in a simple Lucas-tree version of the economy where we show how preference shocks that increases demand are partially accommodated by an increase in productivity via more search. The full production economy is presented in Section 3. We then map the model to data in Section 4. In Section 5 we estimate

[7] Michaillat and Saez (2015) also study a model where aggregate demand can increase output and employment through a search friction in product markets. Their model relies on predetermined (i.e., exogenous) prices and wages as their equilibrium abstracts from a mechanism that can determine search intensity, Moreover, their model is static and the equilibrium is inefficient. In contrast, our dynamic model is embedded in a standard business cycle framework, it assumes flexible prices, and poses a theory of determination of search effort that yields a unique and efficient equilibrium as in Moen (1997).

the model and document that movements in the shopping effort implied by the model is consistent with dynamics of empirical shopping time in the U.S. since 2003. In Section 6 we extend the model to allow for varying capacity utilization of capital, as well as storable goods and firm inventories. Section 7 concludes. Appendices A-G provide proofs, computational and data details, and some additional tables and figures.

2 Competitive search for goods in a Lucas-tree model

We start by illustrating the workings of the model in a simple search model where output is produced by trees instead of capital and labor. We show that the search process has an impact on aggregate output in a way that appears as a level effect on the Solow residual. In an example we show how shocks to preferences are accommodated partly by increases in quantities.

2.1 Technology and preferences

The economy has a measure one of identical, infinitely lived households. Preferences are

$$E\left\{\sum_t \beta^t\, u\left(c_t, d_t, \theta_t\right)\right\}, \tag{1}$$

where c_t is consumption, d_t is shopping effort, and θ_t is a Markovian preference shock.

There is a continuum of trees (i.e., suppliers) with measure $T = 1$. Each tree yields Z pieces of fruit – the consumption good – every period. Fruit lasts for only one period and cannot be stored. A standard search friction makes it difficult for consumers to find trees. To overcome this friction, the household makes search effort $d \in [0, 1]$ to find fruit.

Trees are held by a mutual fund which in turn is owned by households. This is the only asset in the economy. Let a denote the household's ownership of shares in the mutual fund. Since shares add up to one and households are identical, the aggregate state of the economy is just $\Lambda \equiv (\theta, Z)$, whereas the individual state also includes individual wealth in terms of a.[8]

Following Moen (1997), we assume a competitive search protocol where households search in specific locations indexed by price and market tightness. Tightness q is defined as shopping effort per tree, $q = D/T$, where D and T are shopping effort and measure of trees at the location.[9]

[8] Throughout the paper, we exploit the perfect correlation between the idiosyncratic and aggregate shocks to preferences, and we write only one of them as a state variable.

[9] Here, shoppers are searching for trees. In related work (Duras (2015, 2016)) matches occur between household

The number of matches between trees and consumers at the location is given by

$$M = D^{\varphi}\ T^{1-\varphi}, \tag{2}$$

where φ is a matching technology parameter.[10] The rate at which shopping effort finds trees and, hence, fruit for sale, is then $\Psi_d(q) = q^{\varphi-1} = D^{\varphi-1}/T^{\varphi-1}$ and the probability that a tree is matched is $\Psi_T(q) = q \cdot \Psi_d(q) = q^{\varphi}$. Once a match is formed, fruit is traded at the posted price p per unit. The total amount of fruit purchased by the household is $d\Psi_d(q)Z$.

The budget constraint for a household shopping at a location offering a pair (p, q) is given by

$$P_a\, a' = (P_a + \pi)\ a\ -\ p\, d\Psi_d(q)Z, \tag{3}$$

where P_a is the price per unit of share in the mutual fund and π is the dividend per tree.

Let $V(\Lambda, a)$ be the value function that has already incorporated the optimal market choice. Note that when agents are identical, there will only be one active market.[11] A household that were to go to a market with price $\hat{p}$ and tightness $\hat{q}$ would get

$$\widehat{V}(\Lambda, a, \hat{p}, \hat{q}) = \max_d\ u\left(d\Psi_d(\hat{q})Z, d, \theta\right)\ +\ \beta\left\{V\left(\Lambda', \left[(P_a + \pi)\ a - \hat{p}\ d\Psi_d(\hat{q})Z\right]/P_a\right)|\Lambda\right\}, \tag{4}$$

where the expression for a' is derived from rewriting equation (3). V is then determined by the best market; $V(\Lambda, a) = \max_{(p,q)\in\Phi} \widehat{V}(\Lambda, a, p, q)$ where Φ is the set of available markets, yet to be determined.

A firm, or a tree, chooses what market bundle $(\hat{p}, \hat{q})$ to offer. To attract shoppers, it has to offer bundles no worse than the most attractive one in the market. The firm therefore maximizes its profit $\pi(\hat{p}, \hat{q}) = \hat{p}\Psi_T(\hat{q})Z$ subject to the participation constraint of shoppers

$$\widehat{V}(\Lambda, a, \hat{p}, \hat{q}) \geq V(\Lambda, a). \tag{5}$$

search units and goods. There, more fruits per tree would increase the tightness in the market.

[10]We are implicitly assuming that preferences and technology imply that the equilibrium choices imply $D < T$ so that $M < T$, i.e. the number of matches does not exceed the number of trees.

[11]In Section 4.2.1 we analyze the case of cross-sectional heterogeneity in consumption expenditures, which we will use for calibration purposes.

2.2 Equilibrium

For convenience, we choose the consumption good as the numéraire good, so the price in the active market is $p = 1$.

A *competitive search equilibrium* is defined by value functions $V(\Lambda, a)$ and $\hat{V}(\Lambda, a, p, q)$, a set of individual decision rules, $c(\Lambda, a)$, $d(\Lambda, a)$, and $a'(\Lambda, a)$, aggregate allocations $D(\Lambda)$, $q(\Lambda)$, and $C(\Lambda)$, tree price $P_a(\Lambda)$, and tree dividend $\pi(\Lambda)$ so that

1. The individual value functions $V(\Lambda, a)$ and $V(\Lambda, a, p, q)$ and decision rules $d(\Lambda, a)$ solve the household problem (4) with $V(\Lambda, a) = \hat{V}(\Lambda, a, 1, q(\Lambda))$. Consumption is given by $c(\Lambda, a) = d(\Lambda, a)\Psi_d(q(\Lambda))Z$.

2. $q(\Lambda)$ and $p = 1$ maximizes firms' profit subject to condition (5). The optimal profit and dividends are $\pi(\Lambda) = q(\Lambda)^{\varphi} Z$.

3. The individual decision rules are consistent with the aggregate functions, $C(\Lambda) = c(\Lambda, 1)$, $D(\Lambda) = d(\Lambda, 1)$, and $a'(\Lambda, 1) = 1$.

4. Market tightness is $q(\Lambda) = D(\Lambda)$.

5. The goods market clears:

$$C(\Lambda) = D(\Lambda)^{\varphi} Z. \tag{6}$$

The competitive search equilibrium and its efficiency properties can then be characterized by the following proposition.

Proposition 1. *1. Aggregate search D is determined by the functional equation*

$$-u_d\left[D(\Lambda)^{\varphi} Z, D(\Lambda), \theta\right] = \varphi D(\Lambda)^{\varphi-1} Z \; u_c\left[D(\Lambda)^{\varphi} Z, D(\Lambda), \theta\right]. \tag{7}$$

The functions C, q, and π then follow directly from the equilibrium conditions.

2. The equilibrium tree price P_a is defined by the functional equation

$$u_c\left[C(\Lambda), D(\Lambda), \theta\right] = \beta \, \mathrm{E}\left\{ \frac{P_a(\Lambda') + \pi(\Lambda')}{P_a(\Lambda)} u_c\left[C(\Lambda'), D(\Lambda'), \theta'\right] \,\middle|\, \Lambda \right\}. \tag{8}$$

3. *The competitive equilibrium is efficient.*

The proof of the first two items is straightforward: simply derive the first-order conditions of households and combine them with the competitive search conditions (see Appendix A for details). For efficiency, we consider a planner solving $\max_{C,D} \{u(C, D, \theta)\}$ subject to the aggregate resource constraint $C = D^{\varphi} Z$. The solution to equation (7) solves this planner problem, which establishes efficiency. Note that the Euler equation (8) is the same as the one in the standard Lucas tree model.

We can rewrite the intra-temporal condition (7) as

$$-\frac{u_d\left[C\left(\Lambda\right)^{\varphi}, D(\Lambda), \theta\right]}{u_c\left[C\left(\Lambda\right)^{\varphi}, D(\Lambda), \theta\right]} = \frac{\partial M(q)}{\partial q} Z, \tag{9}$$

where M is the number of matches. The condition states that the marginal rate of substitution between shopping and consumption equals the marginal rate of transformation of shopping, i.e., the marginal increase in the probability of matching times the amount of fruit per tree.

In this economy, potential output is always Z. However, actual output is smaller than Z. Some units are not found, and hence some potential output is lost due to insufficient search. Nevertheless, the economy is efficient: finding additional goods is not worth the additional search effort. Total factor productivity is a function of the search effort, $TFP = D(\Lambda)^{\varphi} Z$, and fluctuates in response to preference shocks.

Note that despite our distinction between potential and real output we measure consumption and value added consistently with the way it is measured in NIPA: a lost opportunity does not contribute to output and it is only when it is purchased that potential production translates into consumption.[12] In our model search effort can be interpreted as an input to production. Note that when constructing measures for factor inputs in NIPA, the national statistical offices do not include search effort as an input. The reason why fluctuations in search can give rise to fluctuations in measured TFP is precisely because we (mis-)measure inputs in the same way as in NIPA. Indeed, if search effort were properly accounted for as an input factor, then TFP would be just Z.

[12]An exception from this principle is when goods that are not found can be stored as inventories by the producer and offered for sale later. We return to this issue in Section 6.2 where we extend the model to allow inventories.

2.3 An example

An analytical example allows us to illustrate how preference shocks that increase the desire to consume are partly accommodated by an increase in search effort that squeezes more output out of the economy, thereby making it more productive. Let the period utility function be

$$u(c, d, \theta_d) = \frac{1}{1-\gamma}\left(c - \theta_d \frac{d^{1+1/\eta}}{1+1/\eta}\right)^{1-\gamma}, \tag{10}$$

where the preference shock θ_d to the disutility of search effort is independently and identically distributed with $\mathrm{E}\{\theta_d\} = 1$ and $\theta_d \geq \varphi Z$. Given these preferences, the equilibrium conditions (6)-(7) yield the following equilibrium allocations:

$$D(\Lambda) = (\varphi Z/\theta_d)^{\frac{1}{1/\eta+1-\varphi}}, \qquad C(\Lambda) = (\varphi/\theta_d)^{\frac{\varphi}{1/\eta+1-\varphi}}(Z)^{\frac{1/\eta+1}{1/\eta+1-\varphi}}.$$

A decrease in the disutility of search (or, equivalently, a crave for current consumption) translates into an increase in search effort and, hence, increased consumption. The elasticity of consumption to the shocks is increasing in the share of search in the matching function, φ. As φ goes to zero, the shopping economy converges to the standard Lucas tree model, in which case aggregate consumption is invariant to the demand shock, and the preference shock affects prices only.

3 The stochastic growth model version of the economy

We now embed the matching frictions in product markets into an otherwise standard growth model suitable for quantitative business cycle analysis. We focus on matching frictions in consumption good markets and abstract from matching frictions in firms' purchases, including the investment goods market. While matching frictions must exist also for firm-to-firm transactions, these are presumably less severe than for households since firms buy most of their output from related parties and in larger quantities.[13] Consequently, there are two sectors: one producing consumption goods and the other investment goods. Production uses capital and labor as inputs. We start by describing technology and preferences. We then analyze the problems faced by households and firms, and study price determination in the presence of the search friction.

[13]This assumption mitigates the role of shopping effort for business cycle fluctuations. In a previous version of this paper we assumed a search friction also for investment goods.

3.1 Technology and Markets

There are two types of firms, a unit measure of firms producing consumption and a unit measure of firms producing investment goods. Each firm resides in a location, equivalent to a tree in Section 2. Consumption firms have a technology for transforming capital k and labor n into the consumption good via a standard production function $zf(k, n)$, where f is differentiable and strictly concave. Investment firms are subject to an additional aggregate investment shock z_i, so their production function is $z_i zf(k, n)$.

As in Section 2, we assume a competitive search protocol for consumption goods. Note that since production is endogenous, the equilibrium requires indexing markets not only by price and market tightness (p, q) but also by the quantity of the good offered for sale by the firm, F. The quantity F is in turn a function of the firm's labor and pre-installed capital stock.[14] In Appendix B we show that with endogenous production optimality requires that markets are indexed by the triplet (p, q, F) of price, market tightness, and the quantity of the good produced. We conjecture that this property applies to all competitive search models where consumers search for firms and firms can decide the magnitude of the available goods.

The aggregate state vector is $\Lambda \equiv (\theta, Z, K)$, comprising a vector of preference shocks θ, a vector of technology shocks $Z = (z, z_i)$, and a vector of capital stocks $K = (K_c, K_i)$ installed in consumption and investment firms, respectively. As in Section 2, the matching technology for shoppers and firms is given by equation (2).

To simplify the description of equilibrium we formulate the problem so that all dynamic choices are made by the household: it chooses how much capital to accumulate of each type. Moreover, we omit the possibility to trade shares in a mutual fund owning the firms. Instead, households receive dividends as a lump sum transfer. It would be straightforward to extend the model to explicitly incorporate trade in a mutual fund and let the firms carry out dynamic choices. The equilibrium allocation would not be affected by such extension.

3.2 Households

There is a measure one of households who have preferences over consumption c, shopping effort d, and labor supply n. Preferences are subject to a shock θ, perfectly correlated across households.

[14] Recall that the quantity produced was not part of the variables indexing markets in the simple Lucas-tree model. There, the amount of fruit per tree was exogenous and identical across trees by assumption.

The utility function is $u(c, d, n, \theta)$. The household owns k_c and k_i units of capital installed in consumption- and investment-firms, respectively and is entitled to the dividends of firms. The household's state vector includes the state of the economy Λ and individual holdings of (k_c, k_i).

Households take aggregate variables as given, including the price p, market tightness q, and quantity F in the active consumption-good market, the rental rates on capital invested in consumption- and investment-producing firms, R_c and R_i respectively, aggregate dividends π, the price of new capital P_i, the wage W, and the laws of motion of aggregate capital $K_c' = G_c(\Lambda)$ and $K_i' = G_i(\Lambda)$. These aggregate equilibrium objects are functions of the state vector Λ.

As in Section 2, we denote by $V(\Lambda, k_c, k_i)$ the value of a representative household and by $\widehat{V}\left(\Lambda, k_c, k_i, \widehat{p}, \widehat{q}, \widehat{F}\right)$ the value that it would obtain if it went shopping in a $(\widehat{p}, \widehat{q}, \widehat{F})$ market. Specifically, households in market $(\widehat{p}, \widehat{q}, \widehat{F})$ choose consumption c, shopping effort d, labor supply n, and future capital k_c' and k_i' so as to maximize

$$\widehat{V}\left(\Lambda, k_c, k_i, \widehat{p}, \widehat{q}, \widehat{F}\right) = \max_{d,n,c,k_c',k_i'} u(c, d, n, \theta) + \beta \; \mathrm{E}\{V(\Lambda', k_c', k_i') \mid \Lambda\} \tag{11}$$

subject to the shopping constraint, the budget constraint, the capital accumulation, and the aggregate laws of motion of capital,

$$c = d \; \Psi_d[\widehat{q}] \; \widehat{F}, \tag{12}$$

$$\widehat{p} \; c + P_i \; i = \pi + k_c \; R_c + k_i \; R_i + n \; W, \tag{13}$$

$$i = k_c' + k_i' - (1-\delta) \; (k_i + k_c), \tag{14}$$

$$K_c' = G_c(\Lambda), \qquad K_i' = G_i(\Lambda). \tag{15}$$

The value function V is determined by the best market;
$V(\Lambda, k_c, k_i) = \max_{(p,q,F)\in\Phi} \widehat{V}(\Lambda, k_c, k_i, p, q, F)$, where Φ is the set of available markets.

3.3 Firms

Representative firms rent capital and hire labor in spot markets. Given the state of the economy Λ, consumption firms make two choices: how much labor and consumption-specific capital to rent for producing output, and what market bundle (p, q, F) to offer. The problem for the consumption firms is

$$\begin{aligned}
\pi_c(\Lambda) = \max_{k,n,\widehat{p},\widehat{q},\widehat{F}} \ \widehat{p}\,\widehat{F}\,\Psi_T(\widehat{q}) \quad &- \quad W(\Lambda)\ n \ - \ R_c(\Lambda)\ k, && (16)\\
\text{subject to} \qquad \widehat{V}\left(\Lambda, K, \widehat{p}, \widehat{q}, \widehat{F}\right) \ &\geq \ V(\Lambda, K),\\
z\,f(k,n) \ &\geq \ \widehat{F}.
\end{aligned}$$

Note that the firm's problem (16) is static in the sense that current decisions do not influence the future value of the firm, and future dividends and prices do not matter for the firm's current problem. The problem determines the dividends π_c, the factor demands n_c and k_c, and the triplet (p, q, F), all expressed as functions of Λ.

The corresponding static problem for the investment firms is standard and is given by

$$\pi_i(\Lambda) = \max_{k,n} \ P_i\ z_i\ z\ f(k,n) \ - \ W\ n \ - \ R_i\ k, \tag{17}$$

which determines π_i, n_i, k_i, and aggregate investment.

3.4 Equilibrium

The competitive search equilibrium of the production economy consists of value functions and decision rules $\{V, \widehat{V}, c, d, n, k_c', k_i'\}$, all expressed as functions of the individual state (Λ, k), and a set of aggregate allocations $\{C, D, I, N, G_c, G_i, N_c, N_i, K_c', K_i'\}$, prices $\{p, W, R_c, R_i, P_i\}$, market tightness q, consumption capacity F, dividends π, and profits π_i, π_c, where all aggregate variables are expressed as functions of Λ, such that

1. The consumption good is the numéraire, $p = 1$.

2. The decision rules $c(\Lambda, k_c, k_i)$, $d(\Lambda, k_c, k_i)$, $n(\Lambda, k_c, k_i)$, $k_c'(\Lambda, k_c, k_i)$, $k_i'(\Lambda, k_c, k_i)$, and the associated value functions $\widehat{V}(\Lambda, k_c, k_i, p, q, f)$ and $V(\Lambda, k_c, k_i)$ solve the household problem (11-15), taking as given prices $\{W, R_c, R_i, P_i\}$ and dividend income π. Moreover, $V(\Lambda, K_c, K_i) = \widehat{V}(\Lambda, K_c, K_i, 1, q(\Lambda), F(\Lambda))$.

3. The triplet $p = 1$, $q(\Lambda)$, and $F(\Lambda)$, factor demands $N_c(\Lambda)$ and K_c, and profit $\pi_c(\Lambda)$ solve the problem of the producers of the consumption good (16), taking as given $\{W, R_c, V, \widehat{V}\}$, where $F(\Lambda) = zf(K_c, N_c(\Lambda))$.

4. Factor demands $N_i(\Lambda)$ and K_i, and profit $\pi_i(\Lambda)$ solve the investment producers' problem (17), taking as given $\{W, R_i, P_i\}$.

5. Individual decision rules are consistent with the aggregate functions; $C(\Lambda) = c(\Lambda, K)$, $D(\Lambda) = d(\Lambda, K)$, and $N(\Lambda) = n(\Lambda, K)$.

6. Market clearing conditions are satisfied,

$$
\begin{aligned}
C(\Lambda) &= \Psi_T(q(\Lambda))\, F(\Lambda), \qquad\qquad q(\Lambda) = D(\Lambda),\\
I(\Lambda) &= z_i\, z\, f(K_i, N_i(\Lambda)) = G_c(\Lambda) + G_i(\Lambda) - (1-\delta)(K_c + K_i),\\
N(\Lambda) &= N_i(\Lambda) + N_c(\Lambda), \qquad\qquad \pi(\Lambda) = \pi_i(\Lambda) + \pi_c(\Lambda).
\end{aligned}
$$

7. Aggregate laws of motion of capital are consistent with individual behavior, $K_c' = G_c(\Lambda) = k_c'(\Lambda, K)$ and $K_i' = G_i(\Lambda) = k_i'(\Lambda, K)$.

As we show in Appendix B, the welfare theorems apply so the competitive equilibrium is efficient. It is therefore convenient to solve for the equilibrium using a standard social planner approach where the planner faces the same search friction in the consumption goods market. Given an aggregate state variable Λ, the social planner's problem is

$$
W(\Lambda) = \max_{D, N_c, N_i, K_c', K_c'} \left\{ u\left(D\, \Psi_d[D]\, zf(K_c, N_c), D, N_c + N_i, \theta\right) + \beta\ \mathrm{E}\{W(\Lambda') \mid \Lambda\} \right\} \tag{18}
$$

$$
\text{subject to} \qquad z_i z f(K_i, N_i) = K_c' + K_i' - (1-\delta)(K_c + K_i) \ \text{ and } \ \Lambda' = (\theta', Z', K_c', K_i').
$$

The first-order conditions yield four optimality conditions (see Appendix B for derivations). First, an intra-temporal first-order condition equating the marginal cost and the marginal gain of consumption search (literally, equation (19) equates the marginal rate of substitution between search and consumption to the marginal rate of transformation, i.e., the marginal increase in aggregate consumption from searching slightly harder),

$$
-\frac{u_d}{u_c} = \frac{\partial M[D]}{\partial D}\, z\, f(K_c, N_c). \tag{19}
$$

This corresponds to the intra-temporal condition in equation (9) in the search version of the Lucas model with the only difference that capacity, $z\, f$, is now endogenous. The second optimality condition is a standard intra-temporal first-order condition for labor supply, equating the marginal utility cost of working one additional hour to the marginal utility gain of the increased consumption production due to the additional hour worked,

$$-\frac{u_n}{u_c} = \Psi_T\left[D\right]\ z\frac{\partial f\left(K_c, N_c\right)}{\partial N_c}. \tag{20}$$

The third optimality condition equalizes the marginal gain of investing in each sector,

$$0 = \mathrm{E}\left\{ u_{n'}\ \left(\frac{\partial f\left(K_i', N_i'\right)/\partial K_i'}{\partial f\left(K_i', N_i'\right)/\partial N_i'} - \frac{\partial f\left(K_c', N_c'\right)/\partial K_c'}{\partial f\left(K_c', N_c'\right)/\partial N_c'}\right) \middle|\ \Lambda \right\}. \tag{21}$$

Finally, the planner program implies an Euler equation,

$$\frac{u_n}{z_i\ z\ \frac{\partial f\left(K_i,N_i\right)}{\partial N_i}} = \beta\ \mathrm{E}\left\{ \frac{u_n'}{z' z_i'\ \frac{\partial f\left(K_i',N_i'\right)}{\partial N_i'}} \left(z_i'\ z'\ \frac{\partial f\left(K_i', N_i'\right)}{\partial K_i'} + 1 - \delta\right) \middle|\ \Lambda \right\}. \tag{22}$$

The left-hand side is the utility cost of producing one more unit of the investment good by working $1/\left(z_i\ z\ \partial f\left(K_i, N_i\right)/\partial N_i\right)$ more hours. The right-hand side is the equivalent next-period utility gain of a marginal increase in K_i'. Note that if consumption and investment were the same good, as in the standard growth model, the term on the left-hand side of the equation would be just the marginal utility of consumption u_c and equation (22) would simplify to the familiar Euler equation.

Given the equilibrium allocations it is straightforward to back out the prices that would obtain in a decentralized competitive equilibrium. It is instructive to consider the implied relative price of investment goods (in terms of consumption goods), P_i,

$$P_i \quad = \quad \frac{1}{1-\varphi}\frac{\Psi_T(D)}{z_i}\frac{\partial f\left(K_c, N_c\right)/\partial N_c}{\partial f\left(K_i, N_i\right)/\partial N_i}. \tag{23}$$

The investment price is declining in the investment technology shock z_i since a higher investment productivity increases investment production and lowers the relative value of investment goods. Similarly, P_i is increasing in consumption market tightness $\Psi_T(D)$ and, hence, search effort. The reason is that with higher search effort the relative price of consumption falls.

4 Mapping the Model to Data

We now choose the functional forms for preferences and technology. We then follow the business cycle literature to calibrate the standard parameters. The parameters associated with shopping frictions are identified by using cross-sectional household data.

4.1 Functional forms

We assume a so-called Greenwood-Hercowitz-Huffman formulation for preferences over consumption and shopping search effort. This specification rules out wealth effects in search effort. As we shall see below, such a property is necessary to account for the cross-sectional aspects of shopping effort. This feature will also ensure that the model generates a procyclical shopping effort, in line with the empirical evidence on shopping time documented by Petrosky-Nadeau et al. (2016). We assume that preferences are additively separable between labor and the composite of consumption and search effort. The period utility function is then,

$$u(c, n, d) = \frac{1}{1-\gamma}\left(c - \chi_d \frac{d^{1+1/\eta}}{1+1/\eta}\right)^{1-\gamma} - \chi_n \frac{n^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}}. \tag{24}$$

The parameters are the discount rate β, the inverse of the intertemporal elasticity of substitution, γ, the Frisch elasticity of labor supply ν, and the parameters that determine average hours worked and average search effort, χ_n and χ_d, respectively. The parameter η determines the elasticity of shopping effort with respect to variation in the return to search.

Firms have decreasing returns to scale. This is a natural assumption in a model with frictions in the matching between consumers and firms.[15] The production function is Cobb-Douglas,

$$f(k, n) = k^{\alpha_k}\, n^{\alpha_n}. \tag{25}$$

4.2 Calibration

As far as possible, the calibration targets for the steady state are standard in business cycle research. Table 1 reports the targets and the parameters most closely associated with each target. The targets are defined in yearly terms even though the model period is a quarter.

[15] Note that if production were constant return to scale in capital and labor then the search friction would become irrelevant as one firm would produce all goods and all search would focus on finding the firm with available goods.

Table 1: Calibration Targets, Implied Aggregates, and (Quarterly) Parameter Values

Targets	Value	Parameter	Value
First Group: Parameters Set Exogenously			
Risk aversion	2	γ	2
Real interest rate	4%	β	0.99
Frisch elasticity	0.72	$\frac{1}{\nu}$	0.72
Second Group: Standard Targets			
Investment share of output	0.20	δ	1.78%
Physical capital to output ratio	2.75	α_k	0.25
Labor share of output	0.67	α_n	0.54
Third Group: Normalizations			
Steady-state output	1	$E(z)$	1.20
Relative price of investment	1	$E(z_i)$	0.74
Fraction of time spent working	30%	χ_n	14.91
Capacity utilization of consumption sector	81%	χ_d	1905
Fourth Group: Targets Specific to This Economy			
Cross-sectional stdev. of cons. good prices	9%	φ	0.23
Shopping time expenditure elasticity	7.5%	η	0.11

The first group of parameters does not need the computation of the steady state to be determined: the intertemporal elasticity of substitution is 0.5, and the rate of return is 4 percent. The Frisch elasticity is more controversial. We choose a value of 0.72 based on Heathcote et al. (2010b), who take into account the response of hours worked for both men and women in a model that explicitly incorporates households with husbands and wives.

The second group of parameters are determined simultaneously through specified steady-state targets. The physical capital to output ratio, the investment share of output, and the labor share allow us to identify the depreciation rate and the shares of capital and labor, α_k and α_n. In particular, the quarterly capital depreciation rate δ is set to 1.78% so as to generate an annual investment-output ratio of 20%.[16] A labor share of output of 0.67 implies a labor share parameter $\alpha_n = 0.54$. A capital-output ratio of 2.75 generates a capital share parameter $\alpha_k = 0.25$ (see Appendix C for derivations). The third group of parameters are normalizations that are immaterial to the cyclical behavior of the economy. Nevertheless we target values that have a direct interpretation. The average adult works about 1600 hours per year, so 30% of available time is spent working. Moreover, the share of potential goods that are found is set equal to the average capital utilization in the U.S. Capacity Utilization series published by the Board of Governors (81%).

4.2.1 Identifying Shopping Parameters with Cross-sectional Shopping Data

The last group of parameters, specific to our shopping theory, are the elasticity of shopping effort η and the weight of shopping in the matching function, φ.

Our model is formulated with a representative agent. However, to calibrate the parameters φ and η it is useful to study the cross-sectional implications of our shopping model. To this end, consider a version of our economy where J types of households differ in their household wealth. Let y_j and μ_j denote the consumption expenditure of households of type $j \in \{1, ..., J\}$ and the measure of type j households, respectively. The equilibrium can be characterized as follows (proof in Appendix D.)

Proposition 2. *The competitive equilibrium in the stochastic growth model with cross-sectional heterogeneity is a vector of prices and allocations where all firms offer the same capacity F and*

[16] Note that investment in our model is strictly business investment. Consumption durables are excluded since the markets for business investments are frictionless while consumer goods by assumption are subject to search frictions.

the posted price of the consumption good in the market targeting type j households is

$$p_j = \frac{C^{1-\varphi}}{F} \cdot (y_j)^{\varphi} (D_j)^{-\varphi}, \tag{26}$$

where C is aggregate consumption (in physical units). The measure of firms selling to type j households is equal to the households' expenditure share, $T_j = \mu_j y_j / C$, the real consumption of type j households is $c_j = D_j^{\varphi} (y_j/C)^{1-\varphi} F$, and the search effort is given by the solution to the first-order condition

$$-u_d = \varphi (D_j)^{\varphi-1} (y_j)^{1-\varphi} u_c F. \tag{27}$$

Consider now the utility function in equation (24). From equation (27) the shopping effort for type j households is given by $D_j = (\varphi F / \chi_d)^{\eta/[\eta(1-\varphi)+1]} y_j^{\eta(1-\varphi)/[\eta(1-\varphi)+1]}$. Taking the difference of logs across households of different types yields

$$\Delta \log (D_j) = \frac{\eta (1-\varphi)}{\eta (1-\varphi)+1} \Delta \log (y_j). \tag{28}$$

Consider now the price equation (26) in log form and take the standard deviation on each side of the equation. This yields

$$std (\log (p_j)) = \frac{\varphi}{\eta (1-\varphi)+1} std (\log (y_j)). \tag{29}$$

Kaplan and Menzio (2016) report that the cross-sectional standard deviation of household price indexes is 9% when using standardized bar codes to identify goods. With a standard deviation of log consumption expenditures on services and non-durables of 42.4% (cf. Heathcote et al. (2010a)), we get $std (\log (p_j)) / std (\log (y_j)) = 0.212$. Petrosky-Nadeau et al. (2016) report that in the cross section of individual households, shopping time is *increasing* in household income when controlling for observable household characteristics. For example, households with income between \$100,000 and \$150,000 spend 3.6 minutes more on shopping per day than households with income between \$25,000 and \$50,000. With an average of 42 minutes shopping per day, this difference amounts to $3.6/42 = 8.6\%$ of average shopping time. The difference in income between these groups is roughly a factor of three. This implies an elasticity $\Delta \log (D_j) / \Delta \log (y_j) \approx 0.075$. Applying these moments in equations (28)-(29) imply $\varphi = 0.23$ and $\eta = 0.11$.

5 Quantitative results

We now estimate the model with Bayesian methods to assess the role of demand shocks in accounting for aggregate fluctuations in the U.S. (Section 5.1). We then explore the implications of our model for the Great Recession which, as we shall see, was characterized by a sharp decline in shopping time. Our model imputes a decline in shopping effort so as to account for the sustained fall in consumption and the weak response of the price of investment (Section 5.2).

5.1 Estimating the model with U.S. data

We estimate the process for the shocks that affect the model using U.S. quarterly data (1967I-2013IV). We consider four types of shocks, two preference shocks and two technology shocks. The preference shocks are the disutility to shopping θ_d and the disutility to work θ_n, so preferences are

$$u(c, n, d; \theta_d, \theta_n) = \frac{1}{1-\gamma}\left(c - \theta_d\, \chi_d \frac{d^{1+1/\eta}}{1+1/\eta}\right)^{1-\gamma} - \theta_n\, \chi_n \frac{n^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}}.$$

There are two types of technology shocks: a shock to total factor productivity z and a shock to investment specific technology z_i. We assume that all four shocks follow AR(1) processes with respective persistence $\{\rho_d, \rho_n, \rho_z, \rho_i\}$. The innovations to the shocks are independent and normally distributed with standard deviations denoted by $\{\sigma_d, \sigma_n, \sigma_z, \sigma_i\}$.

We estimate jointly the autocorrelations $\{\rho_d,\ \rho_n,\ \rho_z,\ \rho_i\}$ and the standard deviation of innovations $\{\sigma_d,\ \sigma_n,\ \sigma_z,\ \sigma_i\}$ using Bayesian methods. We use four logged and linearly detrended data series in the estimation: output, investment, labor productivity and the relative price of investment to consumption. All data are real and from the National Income and Product Accounts for the period of 1967 to 2013 at a quarterly frequency. Since our model is a closed economy without a government sector, we measure investment as private investment and construct output from the sum of private consumption and private investment. Labor productivity is the ratio of our measure of output to total working hours.

The upper panel of Table 2 shows the priors and posteriors for all shock parameters. We assume that autocorrelations follow a Beta distribution and that standard deviation of innovations follow an inverse Gamma distribution. We assume an initial prior of equal autocorrelation for all shocks and let the estimation tell us which one is more persistent. The estimation has a log-likelihood of 2172. The 90th percentile intervals show that our estimation is tight. The estimated autocorrelations of

Table 2: Bayesian Estimation: Benchmark Model

Priors and Posteriors for the Shock Parameters likelihood = 2172.07					
Parameter	Density	Para(1)	Para(2)	Mean	90% Intv.
ρ_d	Beta	0.930	0.05	0.9995	[0.9989, 0.9999]
ρ_z	Beta	0.930	0.05	0.7894	[0.7782, 0.7994]
ρ_i	Beta	0.930	0.05	0.9849	[0.9807, 0.9883]
ρ_n	Beta	0.930	0.05	0.9986	[0.9974, 0.9999]
σ_d	Inv Gamma	0.150	Inf	0.6099	[0.5542, 0.6749]
σ_z	Inv Gamma	0.002	Inf	0.0120	[0.0107, 0.0131]
σ_i	Inv Gamma	0.002	Inf	0.0164	[0.0148, 0.0179]
σ_n	Inv Gamma	0.010	Inf	0.0193	[0.0176, 0.0207]

	Data			Model			Var Decomp			
	var(x)	$\frac{var(x)}{var(Y)}$	Corr	var(x)	$\frac{var(x)}{var(Y)}$	Corr	θ_d	z	z_i	θ_n
Y	4.68	1.00	1.00	4.10	1.00	1.00	28.61	54.03	11.31	6.05
Y/N	1.47	0.31	0.50	5.61	1.37	0.87	53.15	41.60	2.54	2.71
Investment	48.48	10.36	0.95	37.38	9.11	0.72	0.78	47.24	50.31	1.67
P_i	0.70	0.15	-0.01	7.06	1.72	0.58	41.81	32.32	25.58	0.30
Labor	3.52	0.75	0.83	1.33	0.32	-0.04	31.25	1.55	7.76	59.45
Consumption	1.58	0.34	0.93	3.21	0.78	0.80	69.00	23.34	1.95	5.71
TFP	2.05	0.44	0.79	4.85	1.18	0.94	47.07	48.02	4.74	0.17
Capacity	19.15	4.09	0.89	3.23	0.79	0.54	99.98	0.01	0.00	0.00
Shopping				7.06	1.72	0.58	99.98	0.01	0.00	0.00
$corr(C,I)$		0.76			0.16					

The table shows the Bayesian estimation and the business cycle properties for the benchmark model. All variables are HP-filtered logarithms of the original series. The estimation targets the data series of GDP, labor productivity, investment, and investment price. The sample is quarterly U.S. data, from Q1 in 1967 to Q4 in 2013. Note that the variance $var(x)$ is multiplied by a factor of 10,000.

the preference shocks θ_d and θ_n are close to unity.

The variance decompositions of the major aggregate variables are reported in the lower panel of Table 2. The preference shock to the disutility of shopping, θ_d accounts for about 30 percent of the variance of output and labor.[17] It is also the main driving force for consumption (69%), labor productivity (53%), and the relative price of investment (42%). Combined, the two shocks θ_d and the neutral productivity shock z account for almost the entire variation in consumption, labor productivity, and the Solow residual. The two shocks account for an equal share of the variance of TFP. The shock to z accounts for roughly half of the fluctuations in output and investment. The investment-specific technology shock accounts for a significant share of the investment and investment price dynamics – 50 and 26 percent respectively – but has a negligible role for other variables. The shock to θ_n matters mainly for labor dynamics and contributes little to other variables. Finally, while the previous business cycle literature has usually treated fluctuations in the relative investment price P_i as reflecting investment-specific technology shocks (cf. Greenwood et al. (1997)), we find that θ_d accounts for a larger share of the variance in the relative price of investment goods than does z_i – 42% versus 26%, respectively. The Bayesian estimation procedure chooses shocks so as to maximize the probability of the observed aggregate time series. We conclude that to match the data – especially the consumption dynamics – it is necessary to attribute a significant role to θ_d.

The model generates output fluctuations close to the data, with a standard deviation of 2.02% versus 2.16%. Labor productivity and the investment price are more volatile than in the data, while the variance of consumption, investment, and labor relative to that of output are close to their empirical counterparts. All empirical variables except for the relative investment price are procyclical. The model generates similar patterns. Both the model and the data have positively correlated consumption and investment. Labor in the model is, however, acylical.

5.2 Verifying the model over the Great Recession

The key departure of our model from the existing literature is the role of shoppers in shaping productivity via their endogenous choice of shopping effort. One way to verify that the mechanism

[17] As aa robustness exercise we explore alternative formulations for the technology shocks on goods and investment relative to that of the benchmark economy. This includes (i) assuming independent shocks for consumption and investment production, and (ii) assuming a joint shock z and a relative shock z_x so TFP for consumption and investment production are zz_x and z/z_x, respectively. When estimated, these alternative formulations imply a larger role for θ_d than in Table 2. The benchmark formulation yields a somewhat larger likelihood than the alternative models (2172 versus 2130 or less). See Table 6 for details.

of our model is reasonable is to investigate the model's predictions about the shopping effort. Unfortunately, shopping effort in the model does not have an immediate empirical counterpart. However, it seems reasonable that the *time* households devote to shopping should be correlated with the shopping effort.

This section compares the empirical time spent shopping in the U.S. over the Great Recession with the behavior of shopping effort implied by our model over the same period (Section 5.2.1). We then investigate the mechanism responsible for the shopping effort over that period (Section 5.2.2).

5.2.1 The Behavior of Shopping Effort

To measure shopping time we use the American Time Use Survey (ATUS), which provide annual data from 2003 and onward. We measure shopping time as the average time spent shopping for goods, professional and personal care services, and household services, plus the travel time associated with this shopping.

Panel A of Figure 1 shows the percentage change in average shopping time for various groups, relative to 2003. Shopping time for the working age population (25-59) fell by about 12.5 percent over the sample period. Most of the fall in shopping time occurred between 2006 and 2010, and shopping time is roughly stationary before 2006 and after 2010. This fall in shopping time during the Great Recession is consistent with the evidence documented by Aguiar et al. (2013) and Petrosky-Nadeau et al. (2016). As is clear from the figure, shopping time fell by similar magnitudes also for various subgroups of the population, including the full-time employed and those not full-time employed (11.8 percent and 14.4 percent, respectively, between 2006 and 2010). Aguiar et al. (2013) interpret the decline in shopping time during the Great Recession as reflecting an aggregate trend toward less shopping time. Petrosky-Nadeau et al. (2016) use cross-state and individual regressions analysis and argue that the recent decline in aggregate shopping time reflects that consumer search in the goods market is procyclical.[18]

[18]Note that even when allowing for a linear trend in shopping time between 2003 and 2014, there is a sharp decline in aggregate shopping time during the Great Recession, with the trough in shopping time (relative to the linear trend) occurring during 2009-2010. Aguiar et al. (2013) use data only up until 2010 and show that if a linear trend were applied only for the 2003-2010 period, then the 2009-2010 period would appear as years with shopping time above trend (cf. Table 1 in their paper). Petrosky-Nadeau et al. (2016) argue that the fact that shopping time declined during the Great Recession – a period during which unemployment increased sharply – runs counter to the common perception that shopping time tends to increase during recessions. Since people who are less than full time employed spend about 12 percent more time on shopping than do the full-time employed, one might expect that a reduction in employment should be a force for increasing the average shopping time during recessions (cf. Kaplan and Menzio (2016)). However, this compositional effect has only a minor impact on aggregate shopping time: if shopping time for all groups were held constant, then a fall in employment rate for full time workers of six

We now compare the dynamics of the empirical shopping time with the theoretical shopping effort implied by our estimated model. Given the empirical facts documented above – including the procyclical nature of shopping time (cf. Petrosky-Nadeau et al. (2016)) – it seems natural to require that a theory of shopping effort be consistent with a large empirical drop in shopping time since 2003. To infer the shopping effort predicted by our theory, we first back out the shock innovations using the four observed data series, output, labor productivity, investment, and the relative investment price for the period of 2003 to 2013. The implied annualized shopping effort is shown by the solid line in Panel B of Figure 1. The key observation is that our model generates dynamics of shopping effort qualitatively similar to those of our measure of the empirical shopping time. In particular, the estimated model implies a sharp decline in shopping effort during the sample period, including a decline during the Great Recession.[19]

One additional dimension of the data that may shed light on the model's mechanism is the distinction between goods that are purchased often such as for example groceries, for which the search frictions presumably are smaller than for other goods. If search frictions were to matter for business cycles then one would expect that shopping time for goods with more frictions – i.e., goods other than groceries – should be tighter correlated to the business cycle than is shopping time for goods with less matching frictions. To investigate this issue we split the total shopping time in Figure 1 into shopping for groceries and shopping for everything else. The main finding is that there is no clear trend in shopping time for groceries – it is essentially constant between 2006 and 2013, while shopping time for other goods essentially follows the same pattern as in Figure 1 (cf. Figure 3 in Appendix E). Moreover, while shopping time for other goods is procyclical, the shopping time for groceries is acyclical.[20]

percentage points – equivalent to the change over the 2006-2010 period – should increase average shopping time by merely one half of a percent. This compositional effect is dwarfed by the massive empirical declines in shopping time across groups over this period, documented in Figure 1.

[19]An alternative way to verify our model's predictions of lower search effort after 2006 is to consider the frequency of stockouts of goods. Our model abstracts from the possibility of stockouts. Nevertheless, we conjecture that if our model were extended to incorporate such inventory stockouts, then these would be more likely during periods with surprisingly large search effort. Thus, an increase in search effort should manifest itself as an increase in the probability of observing episodes of stockout. Bils (2016) documents stockouts using micro-CPI data. He finds that the occurrence of stockouts tend to fall during recessions, and especially so during the Great Recession (cf. his Figure 1). This is broadly consistent with Figure 1.

[20]The correlation of grocery shopping time with HP-filtered annual GDP is -0.056, while the correlation of shopping time for other goods with GDP over the 2003-2014 period is 0.53. To calculate these correlations, we HP-filter yearly GDP from 1967 to 2014 and use the cyclical values from 2003 to 2014.

Figure 1: Shopping

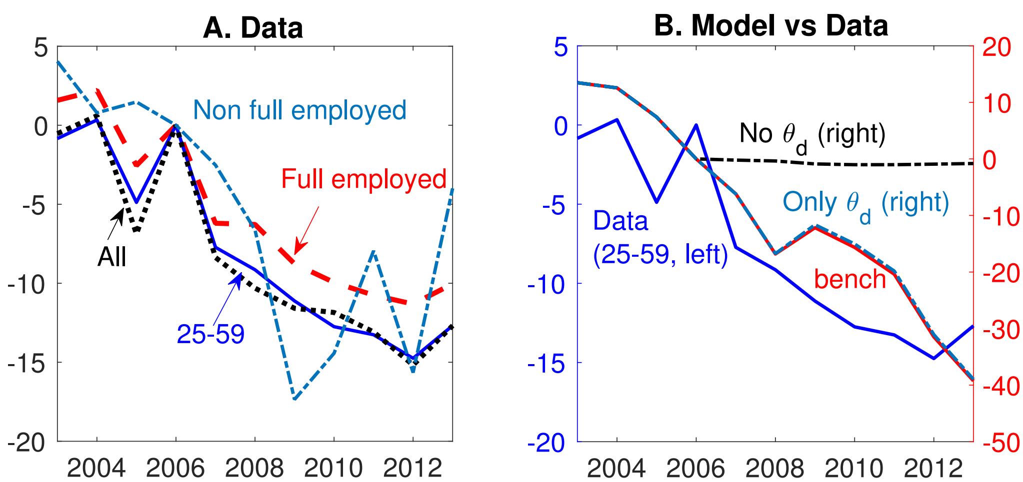

Panel A shows the percentage change in shopping time between 2003 and 2013 for various groups. Shopping time is measured as the average time spent shopping for goods, professional and personal care services, and household services, plus the travel time associated with this shopping. Data source: American Time Use Survey (ATUS).

Panel B shows the evolution of average shopping time for the group 25-59 years old (blue solid line) plotted against the shopping effort implied by our estimated benchmark model (red solid line) over the ATUS sample period. The panel also displays shopping effort under two alternative scenarios – one where the θ_d shock follows its estimated path while the remaining shocks (θ_n, z, z_i) are fixed after 2006 (green dashed line), and one where θ_d is constant after 2006 while the remaining shocks (θ_n, z, z_i) follow their estimated paths (black dash-dot line).

5.2.2 Inspecting the mechanism

We now investigate why the estimated model implies a large decline in shopping effort after 2006. We start by documenting which of the four shocks of the estimated model $(\theta_d, \theta_n, z, z_i)$ is the main driver of the implied endogenous shopping effort. To this end we show how the shopping effort predicted by the model would have evolved from 2006 and onward under two counterfactual scenarios: one where the θ_d shock follows its estimated path after 2006 while the remaining shocks (θ_n, z, z_i) are held constant at their 2006 levels, and one where the shocks (θ_n, z, z_i) follow their estimated paths after 2006 while the θ_d shock is held constant at its 2006 levels. The implied counterfactual paths for shopping effort are displayed in Panel B of Figure 1. As is clear from the figure, the shopping effort in the model is entirely driven by the dynamics of θ_d: the scenario where only θ_d moves is visually indistinguishable from the benchmark scenario involving all the shocks.

Figure 2: Counterfactual Series for Aggregate Variables under Fixed θ_d.

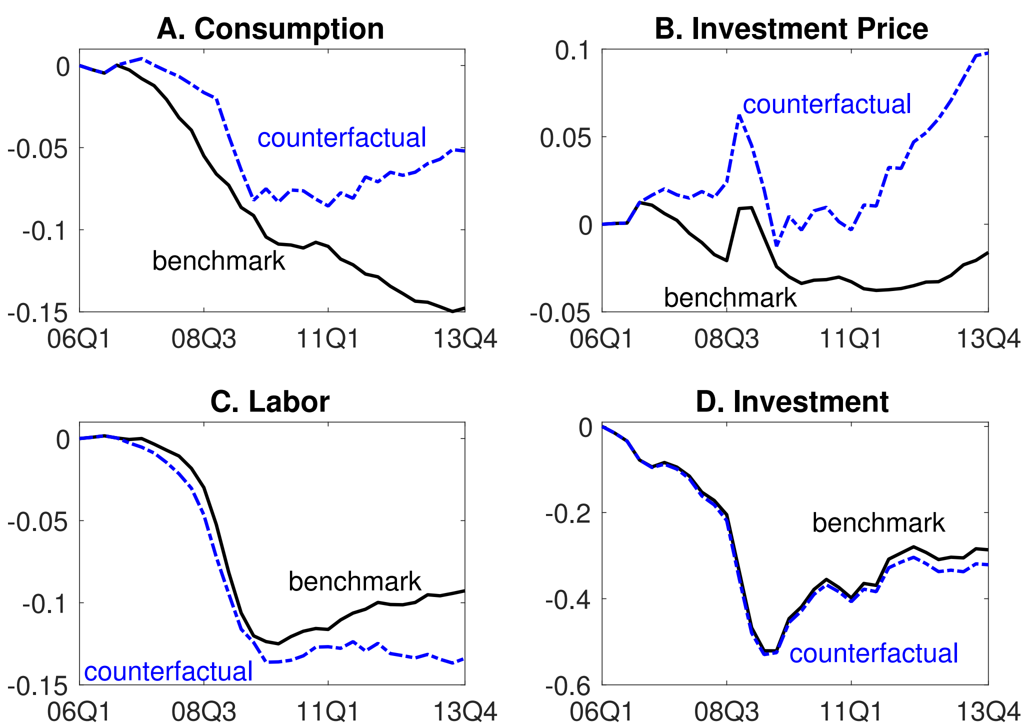

The four panels display the aggregate time series of C_t, P_{it}, N_t, and I_t after 2006, plotted against their respective counterfactual series predicted by the estimated model under the assumption that the *search shock* θ_d remains constant after 2006. Panel A displays consumption C_t, Panel B displays the relative investment good price P_{it}, Panel C displays labor supply N_t, and Panel D displays investment I_t.

Given the insight that shopping effort is driven by the *search shock* θ_d, we would like to know why our Bayesian estimation predicts such a large fall in θ_d after 2006. Recall that the model is estimated by targeting the aggregate time series of C_t, I_t, N_t, and P_{it}. We now investigate what aspects of these empirical observations identify the implied dynamics of θ_d. To this end we perform a series of counterfactual analyses, asking what the evolution of each of these aggregate variables would have been – according to the model – if the shopping shock θ_d had remained fixed from 2006 and onward.

Panel A of Figure 2 shows that the large and sustained fall in aggregate consumption (relative to trend) during the Great Recession is an important reason why the model implies a sustained fall in shopping effort. As can be seen from the figure, aggregate consumption would – in the absence of the decline in shopping effort – have risen by ten percent relative to its actual path over the 2007-2014 period. This effect is especially large after 2009. The larger disutility of search effort makes the households search less for consumption goods – consistent with the observed fall in shopping time. This in turn makes consumption more expensive. The behavioral response in the model is to reduce consumption and, to a lesser extent, increase labor supply (we discuss this in detail below).

Panel B of Figure 2 shows that the fall in the relative price of investment goods is also a factor pushing for a decline in the implied shopping effort, especially after 2009. The empirical investment price fell slightly during the Great Recession and remained low after 2009. In the absence of the decline in shopping effort the investment price should have risen by twelve percent relative to its actual path. To understand this point, note that given the assumed Cobb-Douglas production function we can rewrite the investment price expression in equation (23) as

$$P_i = \frac{1}{1-\varphi}\frac{D^{\varphi}}{z_i}\left(\frac{k_c}{k_i}\right)^{\alpha_k}\left(\frac{n_c}{n_i}\right)^{\alpha_n-1}. \tag{30}$$

Clearly, a fall in search effort D should be associated with a lower investment price P_i.[21]

Panel C of Figure 2 illustrates how aggregate labor supply would have evolved – according to the model – if shopping effort had not fallen after 2006. Between 2006 and 2009 the effect is small. After 2009, labor supply would have fallen more if θ_d had remained constant. Ceteris paribus, a lower search effort lowers the tightness in the goods market and, hence the probability that a firm

[21]Note that the ratio $(k_c/k_i)^{\alpha_k}(n_c/n_i)^{\alpha_n-1}$ is approximately constant. This ratio would be exactly constant under perfect foresight or if capital could be reallocated between K_i and K_c within the period. It follows that the dynamics of P_i is mainly driven by fluctuations in θ_d and the technology shocks.

will be matched, $\Psi_T[D]$. From the intra-temporal first-order condition for labor supply (20), a lower $\Psi_T[D]$ will in turn lower the marginal disutility of labor supply, inducing a lower n. Finally, Panel D shows that aggregate investment is hardly affected by the dynamics of θ_d.

We conclude that according to our model, the unusually large and sustained fall in consumption (relative to output) and the weak development of the relative investment price observed during and after the Great Recession point toward a large and sustained fall in shopping effort.

6 Extensions: Variable capacity utilization and Storable Goods

We now extend the model in two directions. The first is to allow variable capacity utilization. The aim of this extension is to separate our mechanism from that of variable capacity utilization and mismeasured inputs (Section 6.1). The second studies environments where goods can be stored by both households and firms (Section 6.2).

6.1 Varying capacity utilization

In our model demand shocks – in the form of for example a crave for consumption or a change in the disutility of search effort – induce an endogenous change in search effort and, hence, a change in the share of the consumption capacity that is realized as consumption. Thus, shopping effort is a factor of production. However, since this factor is unmeasured, fluctuations in search influence the measured Solow residual, as illustrated in Section 2.3.

This point is related to the well known measurement concern that any mismeasurement of factor inputs will show up as movements in the Solow residual. It is a common practice in DSGE models to allow for varying capacity utilization, i.e., changes in the intensity of factors of production, which in turn can induce movements in the Solow residual, possibly independently from the true movements in technology (see for instance King and Rebelo (1999) and Francis and Ramey (2005)).

The purpose of this section is to demonstrate that the significant role our estimation attributes to (shopping utility) preference shocks – see Table 2 – is robust to allowing for varying capacity utilization of capital and shocks to this utilization. To this end we extend our benchmark model to allow for varying capacity utilization of capital along the lines of Christiano et al. (2005).[22] In particular, we assume that to utilize a fraction h of pre-installed capital, households have to pay

[22] Since wages are not very pro-cyclical, we ignore the possibility of varying capacity utilization of labor.

the variable cost $\psi(h)$, given by

$$\psi(h) \ = \ \xi \, \frac{h^{1+\sigma_a} - 1}{1+\sigma_a}, \tag{31}$$

where $1/\sigma_a$ captures the elasticity of depreciation of capital with respect to how intensively it is used. Households can choose a separate capacity utilization for each capital stock, denoted h_c and h_x for utilization of capital for consumption and investment sector, respectively. The budget constraint – the equivalent to equation (13) with consumption as the numéraire good – becomes

$$c + P_i \ i \ = \pi + (h_c k_c) \ R_c + (h_i k_i) \ R_i + n \ W,$$

where the expression for investment i – the equivalent to equation (14) – now incorporates the additional cost of using capital more intensively,

$$i = k_c' \ + \ k_i' \ - \ (1 - \delta - \psi(h_c)) \ k_c \ - \ (1 - \delta - \psi(h_i)) \ k_i.$$

We first estimate this version of model with the same shocks and data series as in the benchmark model, but allowing the value for σ_a to be estimated. We assume that σ_a follows an inverse Gamma distribution. Christiano et al. (2005) calibrate σ_a to be 0.1. We therefore use this value as the prior for σ_a. Table 3 presents the estimated results. In line with the findings of Christiano et al. (2005) the estimated value for σ_a implies a large elasticity.

Capacity utilization on capital gives the model an additional channel for labor productivity and Solow residual movements. This in turn generates smaller estimated fluctuations for the demand shock θ_d and productivity shock z. The standard deviation of θ_d drops from 0.6 to 0.4 and the standard deviation of z drops from 0.012 to 0.008. The variance decomposition, however, shows that even when allowing for a standard mechanism of capital capacity utilization, our shopping mechanism is still significantly important (cf. Table 3). In particular, the demand shock θ_d accounts for about 19 percent of GDP variability, about 50 percent of TFP variability, and almost 2/3 of consumption variability (versus 29, 47, and 69 percent, respectively, in the benchmark economy).

Ideally we would like to measure the role of varying capacity utilization versus the varying shopping effort. Fortunately, there exists aggregate measurements of capital capacity utilization – for example the series Total Capacity Utilization in Manufacturing – that allows us to undertake such a comparison by incorporating a quarterly version of such data explicitly in the estimation of

Table 3: Bayesian Estimation: with Capacity Utilization on Capital

Priors and Posteriors for the Shock Parameters likelihood = 2274.84					
Parameter	Density	Para(1)	Para(2)	Mean	90% Intv.
ρ_d	Beta	0.930	0.05	0.9998	[0.9996, 0.9999]
ρ_z	Beta	0.930	0.05	0.8135	[0.8120, 0.8156]
ρ_i	Beta	0.930	0.05	0.9992	[0.9986, 0.9998]
ρ_n	Beta	0.930	0.05	0.9988	[0.9982, 0.9995]
σ_d	Inv Gamma	0.150	Inf	0.4122	[0.3862, 0.4469]
σ_z	Inv Gamma	0.002	Inf	0.0080	[0.0074, 0.0085]
σ_i	Inv Gamma	0.002	Inf	0.0076	[0.0068, 0.0084]
σ_n	Inv Gamma	0.010	Inf	0.0189	[0.0173, 0.0202]
σ_a	Inv Gamma	0.100	Inf	0.0390	[0.0226, 0.0558]

	Data			Model			Var Decomp			
	var(x)	$\frac{var(x)}{var(Y)}$	Corr	var(x)	$\frac{var(x)}{var(Y)}$	Corr	θ_d	z	z_i	θ_n
Y	4.68	1.00	1.00	2.54	1.00	1.00	19.29	57.54	8.91	14.26
Y/N	1.47	0.31	0.50	2.82	1.11	0.83	44.45	45.53	7.74	2.28
Investment	48.48	10.36	0.95	19.97	7.87	0.83	0.81	77.81	16.89	4.48
P_i	0.70	0.15	-0.01	2.92	1.15	0.63	49.58	31.11	18.72	0.60
Labor	3.52	0.75	0.83	0.91	0.36	0.21	19.28	0.65	0.01	80.06
Consumption	1.58	0.34	0.93	1.51	0.60	0.86	62.93	18.26	1.21	17.59
TFP	2.05	0.44	0.79	2.01	0.79	0.71	50.47	42.89	6.06	0.58
Capacity	19.15	4.09	0.89	1.54	0.61	0.45	99.98	0.01	0.00	0.01
Shopping				2.92	1.15	0.63	99.98	0.01	0.00	0.01
$corr(C, I)$		0.76			0.44					

The table shows the Bayesian estimation and the business cycle properties for the shopping model with capacity utilization on capital. All variables are HP-filtered logarithms of the original series. The estimation targets the data series of GDP, labor productivity, investment, and investment price. The sample is quarterly U.S. data, from Q1 in 1967 to Q4 in 2013. The inverse of σ_a is the elasticity of capital depreciation rate to capital usage. Note that the variance $var(x)$ is multiplied by a factor of 10,000.

Table 4: Bayesian Estimation: with Capacity Utilization Shock

Priors and Posteriors for the Shock Parameters likelihood = 2843.28

Parameter	Density	Para(1)	Para(2)	Mean	90% Intv.
ρ_d	Beta	0.930	0.05	0.9997	[0.9994, 1.0000]
ρ_z	Beta	0.930	0.05	0.8325	[0.8136, 0.8513]
ρ_i	Beta	0.930	0.05	0.9985	[0.9973, 0.9996]
ρ_n	Beta	0.930	0.05	0.9983	[0.9971, 0.9997]
ρ_h	Beta	0.930	0.05	0.9091	[0.8985, 0.9220]
σ_d	Inv Gamma	0.150	Inf	0.4497	[0.4071, 0.4804]
σ_z	Inv Gamma	0.002	Inf	0.0076	[0.0071, 0.0082]
σ_i	Inv Gamma	0.002	Inf	0.0082	[0.0074, 0.0092]
σ_n	Inv Gamma	0.010	Inf	0.0194	[0.0181, 0.0208]
σ_h	Inv Gamma	0.005	Inf	0.0120	[0.0108, 0.0132]
σ_a	Inv Gamma	0.100	Inf	0.0743	[0.0314, 0.1165]

	Data			Model			Var Decomp				
	var(x)	$\frac{var(x)}{var(Y)}$	Corr	var(x)	$\frac{var(x)}{var(Y)}$	Corr	θ_d	z	z_i	θ_n	θ_h
Y	4.68	1.00	1.00	2.74	1.00	1.00	20.70	49.59	9.79	13.31	6.60
Y/N	1.47	0.31	0.50	3.20	1.17	0.84	45.57	37.85	7.58	2.15	6.86
Investment	48.48	10.36	0.95	20.80	7.60	0.82	0.94	68.12	20.75	4.79	5.40
P_i	0.70	0.15	-0.01	3.18	1.16	0.62	52.42	26.37	18.90	0.71	1.61
Labor	3.52	0.75	0.83	0.96	0.35	0.16	21.40	0.47	0.07	77.86	0.20
Consumption	1.58	0.34	0.93	1.71	0.63	0.87	64.49	15.53	0.95	14.86	4.16
TFP	2.05	0.44	0.79	2.22	0.81	0.72	53.42	34.95	5.93	0.75	4.94
Capacity	19.15	4.09	0.89	1.78	0.65	0.47	99.98	0.01	0.00	0.01	0.00
Shopping				3.18	1.16	0.62	99.98	0.01	0.00	0.01	0.00
$corr(C, I)$		0.76			0.42						

The table shows the Bayesian estimation and the business cycle properties for the shopping model with capacity utilization on capital. All variables are HP-filtered logarithms of the original series. The estimation targets the data series of GDP, labor productivity, investment, investment price, and total capacity utilization. The sample is quarterly U.S. data, from Q1 in 1967 to Q4 in 2013. The inverse of σ_a is the elasticity of capital depreciation to capital usage. ρ_h and σ_h are the autocorrelation and standard deviation of innovation of the cost for capacity utilization of capital. Note that the variance $var(x)$ is multiplied by a factor of 10,000.

the model. Since we target an additional data series we must introduce an additional shock. In a way parallel to the shock to the disutility of shopping effort we introduce a shock θ_h to the cost of varying capacity utilization, affecting both capital stocks. The cost of capacity utilization h then becomes

$$\psi(h, \theta_h) \; = \; \theta_h \, \xi \, \frac{h^{1+\sigma_a} - 1}{1 + \sigma_a}, \tag{32}$$

We assume that the shock θ_h follows an AR(1) process with an autocorrelation of ρ_h and a standard deviation of innovation of σ_h. We estimate this shock together with the four shocks of the benchmark model using the same data series as before plus the aggregate Total Capacity Utilization series from the Federal Reserve Bank of St. Louis. Table 4 provides the results. The estimation, variance decomposition, and the business cycle statistics are similar to the previous estimation without the extra shock. The main finding is that the role of the shopping preference shock θ_d in the capacity utilization model is unaffected by the existence of the new shock. It still contributes to over 20 percent of the variance of output, to 64 percent of consumption, and over half of the variation of TFP. Moreover, the capacity utilization shock plays a small role in accounting for the aggregate variables, including the Solow residual – less than a third of the contribution of the shopping shock.

We conclude that varying capital capacity utilization is a mechanism fundamentally different from our shopping friction mechanism, and that our shopping effort mechanism remains important regardless of whether or not we also allow for varying capacity utilization and shocks to capacity utilization costs in the model.

6.2 Storable Goods

In our benchmark model newly produced consumption goods can either be found by a shopper (and consumed instantly) or not found, in which case the goods are permanently lost. This property seems a good description of services, which are consumed the moment they are produced. However, for a large number of goods – especially for durables – producers have a third alternative: goods that are not instantly found can be stored as inventories and offered for sale again later. As it turns out, inventories are large. For example, the ratio of inventories to annual sales for total business (i.e., manufacturing, retail, and wholesale) is 11.5% (source: Bureau of Labor Statistics (BLS)).

The possibility to store goods could in principle change the mechanism for how shocks to shopping utility propagates to the aggregate economy, relative to the mechanism in our benchmark model. To investigate this possibility, we extend our shopping environment to model explicitly the distinction between services and consumer goods, allowing goods to be stored as inventories for firms and as a stock of consumer goods for households. Since a good can be stored by firms, it is natural to allow it to be stored also by households. As we shall see, this property is important. Investment goods and services are modeled exactly as in the benchmark model in Section 3.

6.2.1 A Shopping Model with Storable Goods

There are search frictions for both consumption goods and services but not for investment. An unmatched service location produces zero output. Goods, however, can be stored by both households and firms.[23] Consumer goods are produced in two stages. First, there is a manufacturing stage where a physical object is produced using capital and labor. This object is then added to the firm's stock of inventories and these goods become available in a location that may or may not be matched with a shopper. If there is a match, a sale is produced that adds to output. If there is no match, the goods become inventories that may in turn be traded across firms in a secondary frictionless market.

Output is the sum of the sales of goods, services, and investment, plus the change in inventories. This procedure for measuring output is consistent with the way GDP is measured in NIPA. The aggregate state of the economy is $\Lambda = (\theta, G, X, K)$ where G denotes households' stocks of consumer goods, X denotes firms' inventories and K is a vector of the capital stocks, $K = (K_g, K_s, K_i)$, with K_g for goods sector, K_s for service sector, and K_i for investment sector.

Markets are differentiated by what is traded, and by price, tightness, and quantity. The triplets offered in consumer goods markets and in service markets are (p_g, q_g, F_g) and (p_s, q_s, F_s), respectively, where p_j is the price, $q_j = D_j/T_j$ is market tightness, and F_j is the capacity in market $j \in \{s, g\}$. The price of consumer good inventories is given by p_x. The measure of service firms is normalized to $T_s = 1$.

Households Let g be the stock of consumer goods the household carries over from the previous period and let c be the new purchases. The stored goods depreciate at rate δ_h. The current stock of consumer goods is therefore $c + (1 - \delta_h)g$. Households have preferences over services from their stock of consumer goods, services s, total shopping effort on goods d_g and services d_s, and hours worked n. Current utility is given by $u(c + (1 - \delta_h)g, s, d_g + d_s, n; \theta)$ where $\theta = (\theta_d, \theta_n)$ is the

[23] Since there is no search friction for investment goods, we abstract from the possibility that investment producers may store investment goods.

same vector of preference shocks as in the benchmark model.

Each period, households start with an initial holding of consumer goods g, and capital in each sector $k = (k_g, k_s, k_i)$. Households choose consumption goods, services, shopping efforts, hours worked, and investments in each sector $i = (i_g, i_s, i_i)$, taking as given the aggregate state Λ, the wage W, and rental returns $R = (R_g, R_s, R_i)$, to maximize

$$V(\Lambda, g, k) = \max_{c,s,d_g,d_s,n,i,k'} u(c + (1-\delta_h)g, s, d_g + d_s, n; \theta) + \beta\, E\{V(\Lambda', c + (1-\delta_h)g, k')\}$$

subject to the budget constraint

$$\begin{aligned} c\, p_g(\Lambda) + & s\, p_s(\Lambda) + p_i(\Lambda)(i_g + i_s + i_i) \\ & \leq \pi(\Lambda) + n\, w(\Lambda) + R_g(\Lambda)K_g + R_s(\Lambda)k_s + R_i(\Lambda)k_i, \end{aligned}$$

the shopping constraints for goods and services, $c = d_g\ \Psi_d\,[q_g(\Lambda)]\ F_g(\Lambda)$ and $s = d_s\ \Psi_d\,[q_s(\Lambda)]\ F_s(\Lambda)$, and capital accumulations for each sector $m = \{g, s, i\}$, $k'_m = (1-\delta)k_m + i_m$.

Let the temporary value functions $\widehat{V^g}(\Lambda, g, k, \hat{p}_g, \hat{q}_g, \hat{F}_g)$, $\widehat{V^s}(\Lambda, g, k, \hat{p}_s, \hat{q}_s, \hat{F}_s)$ denote the utility attained by the household if the market for consumption goods and services were given by the triplets $(\hat{p}_g, \hat{q}_g, \hat{F}_g)$ and $(\hat{p}_s, \hat{q}_s, \hat{F}_s)$ as opposed to the equilibrium triplets $(p_g(\Lambda),\ q_g(\Lambda),\ F_g(\Lambda))$ and $(p_s(\Lambda),\ q_s(\Lambda),\ F_s(\Lambda))$, respectively.

Firms As in the benchmark model, each sector has a fixed measure of firms and a production function f as in equation (25). Moreover, service production is subject to a neutral technology shock z, while consumption goods production and investment production have sector-specific TFP shocks given by $z_g * z$ and $z_i * z$, respectively, where z, z_g, and z_i are independent.[24] Firms in the services sector and in the investment good sector solve the same problem as in the benchmark model (cf. equations (16)-(17)) with the participation constraints in the service sector given by $\widehat{V^s}(\Lambda, G, K, \hat{p}_s, \hat{q}_s, \hat{F}_s) \geq V(\Lambda, G, K)$.

Consider now the problem of consumer goods producers. Firms carry inventories x from last period that depreciate at rate δ_f. Firms then choose inputs of labor n_g and capital k_g for production

[24] We also experimented with a range of different specifications for the productivity and preference shocks, including (i) a joint productivity shock for consumption goods and services plus a relative productivity shock that increases productvity for goods relative to services, (ii) independent sector-specific TFP shocks z_s, z_g, and z_i, and (iii) a model where the search effort preference shock is allowed to differ for services and for consumption goods, while TFP for goods and services are identical. Our benchmark specification is the one that yields the highest likelihood in the Byesian estimation. See Table 6 in Appendix G for details.

and add the current production $z_g\, z\, f(k_g, n_g)$ to their net inventories $(1-\delta_f)x$. This sum yields current capacity F_g. Firms also choose which triplet (p_g, q_g, F_g) to offer to customers. After the matching process, firms choose how much extra investment of inventory to purchase or to sell in the frictionless market for inventories. The unmatched firms purchase $i_{x,u}$ units of inventory and the matched firms purchase $i_{x,m}$. This in turn gives the next period's inventories x'_n and x'_m for the unmatched and matched firms, respectively,

$$\begin{aligned} x'_n &= z_g\, z\, f(n_g, k_g) + (1-\delta_f)x + i_{x,u}, \quad (33)\\ x'_m &= z_g\, z\, f(n_g, k_g) + (1-\delta_f)x - F_g + i_{x,m}. \quad (34)\end{aligned}$$

To summarize, a consumer goods producer chooses $\{n_g, k_g, p_g, q_g, F_g, i_{x,u}, i_{x,m}, x'_m, x'_n\}$ to solve the following recursive problem, where $\Psi_{T,g}(q_g)$ is the probability that the firm will be matched with a shopper, $\Omega(\Lambda, x)$ is the value of the firm, and $\Upsilon(\Lambda, \Lambda')$ is the stochastic discount factor.

$$\begin{aligned}\Omega(\Lambda, x) = \max_{n_g, k_g, i_{x,m}, i_{x,u}, p_g, q_g, F_g} \quad & -W(\Lambda)n_g - R_g(\Lambda)k_g \\ & + \Psi_{T,g}(q_g)\left\{p_g F_g - p_x(\Lambda)i_{x,m} + E\left[\Upsilon(\Lambda,\Lambda')\,\Omega(\Lambda', x'_m)|\Lambda\right]\right\} \\ & \quad + (1-\Psi_{T,g}(q_g))\left\{-p_x(\Lambda)i_{x,u} + E\left[\Upsilon(\Lambda,\Lambda')\,\Omega(\Lambda', x'_n)|\Lambda\right]\right\}\end{aligned}$$

subject to the capacity constraint

$$F_g \le zf(n_g, k_g) + (1-\delta_f)x,$$

the dynamics of inventories for non-matched (33) and matched state (34), and the participation constraint for shoppers $\widehat{V}_g(\Lambda, G, K, \hat{p}_g, \hat{q}_g, \hat{F}_g) \ge V(\Lambda, G, K)$.

Average dividends of consumption good producers are

$$\pi_g(\Lambda) = \Psi_{T,g}(q_g)p_g F_g - W(\Lambda)n_g - R_g(\Lambda)k_g - \Psi(q_g)p_x(\Lambda)i_{x,m} - (1-\Psi(q_g))p_x(\Lambda)i_{x,u}.$$

Finally, in Appendix F we show that firms with identical initial inventories choose the same inventory holdings for the following period, i.e., $x'_n(\Lambda, x) = x'_m(\Lambda, x)$..

Equilibrium Let the consumption good be the numéraire, so $p_g = 1$. The competitive equilibrium consists of allocations $\{S,\ C,\ D_g,\ D_s,\ I_g,\ I_s,\ I_i,\ N,\ N_g,\ N_s,\ N_i,\ G,\ K'_g,\ K'_s,\ K'_i,\ X',\ x'_m,\ x'_n,\ i_{x,u},$ $i_{x,m}\}$, dividend and profits $\{\pi, \pi_g, \pi_i, \pi_s\}$, values V and Ω, prices $\{W, R_g, R_s, R_i, p_x, p_i, \Upsilon\}$, market tightness and capacity $p_g, p_s, q_g, q_s, F_g, F_s$ such that

1. Households choose $\{C, S, D_g, D_s, I_g, I_s, I_i, N, \pi, G', K'_g, K'_s, K'_i, V, a'\}$ to solve their problem taking as given prices $\{W, R_g, R_s, R_i, p_i, P_a\}$ and dividends π. The stochastic discount factor Υ satisfies $\Upsilon(\Lambda, \Lambda') = \beta u_s(\Lambda')/u_s(\Lambda)$.

2. The allocation $\{p_g, q_g, F_g, N_g/T_g, K_g/T_g\}$, profit π_g, inventory choices $X'/T_g = x'_m = x'_n$, $i_{x,u}$, $i_{x,m}$, and value function Ω solve the problem of the consumption goods producers, taking as given $\{\Upsilon, W, R_g, p_x\}$.

3. The allocation $\{N_i, K_i\}$ and the profit π_i solve the problem of the investment producers, taking as given $\{W, R_i, p_i\}$.

4. The allocation $\{p_s, q_s, F_s, N_s, K_s\}$ and the profit π_s solve the problem of the consumption service producers, taking as given $\{W, R_s\}$.

5. Market clearing conditions are satisfied

$$
\begin{aligned}
C &= T_g \Psi_{T,g}(q_g) F_g & S &= \Psi_{T,s}(q_s) z f(N_s, K_s) \\
z z_i f(N_i, K_i) &= I_g + I_s + I_i & N &= N_g + N_s + N_i \\
X' &= T_g [1 - \Psi_{T,g}(q_g)] F_g & \pi &= \pi_g + \pi_i + \pi_s \\
F_g &= (1-\delta) X/T_g + z z_g f(K_g/T_g, N_g/T_g). & &
\end{aligned}
$$

The equilibrium is optimal (we omit the proof since it is similar to the one in Section 3).

6.2.2 Calibration and Quantitative Analysis

Following Herrendorf et al. (2013) we assume that consumption is a constant-elasticity-of-substitution (CES) composite of services and consumption goods. Households' preferences are otherwise similar to those in the benchmark model and are given by

$$
u(c, s, d_c, d_s, n) = \frac{1}{1-\gamma} \left\{ \left[\mu c^{\frac{\omega-1}{\omega}} + (1-\mu) s^{\frac{\omega-1}{\omega}} \right]^{\frac{\omega}{\omega-1}} - \chi_d \frac{(d_c + d_s)^{1+1/\eta}}{1+1/\eta} \right\}^{1-\gamma} - \chi_n \frac{n^{1+1/\nu}}{1+1/\nu},
$$

where μ is the share of storable consumption goods and ω is the elasticity of substitution between services and the consumption good.

Most parameters are calibrated using the same moments as in the benchmark (the parameter values and targets are summarized in Table 1 in the Supplementary Appendix). We set μ so the model matches the empirical share of consumer durables in total consumption expenditures, i.e., 8 percent. BEA defines consumer durables as goods which are storable and last for at least a year. This corresponds well to our notion of storable consumption goods. Similarly, the empirical counterpart to services in the model is the sum of non-durable consumption and services. In the robustness analysis we investigate the effect of opting for a wider notion of storable goods. According to the BEA, the average annualized depreciation rate of durable consumption goods (held by consumers) is 18.5 percent. Richardson (1995) argues that the annual cost of holding inventories is between 25 percent and 55 percent of the value of the stored goods. Taking out interest rate costs, the cost of holding inventories is between 19 percent to 43 percent per year. Thus, the depreciation rates of goods does not seem to be higher for households than for firms. For simplicity we equate these rates and set them to $\delta_h = \delta_f = 19$ percent. We set the elasticity between goods and services to $\omega = 0.85$, as in Herrendorf et al. (2013).

The measure of service and investment producing firms are each normalized to unity. To calibrate the measure of goods producing firms T_g we consider the ratio of inventories to sales since the market tightness D_g/T_g pins down the relative size of inventory to sales. The average annual real inventories to sales ratio for total business (manufacturing, wholesale and retail trade) is 11.5 percent (Source: BLS). We also normalize the steady-state investment goods price to 1 as in the benchmark calibration and equalize the steady-state productivity of goods and service production. These normalizations pin down the steady state productivities z_s, z_g, and z_i.

As in the benchmark, we estimate the model using Bayesian methods. The model has five shocks: the disutility shock of shopping θ_d, the disutility shock of labor θ_n, and the technology shocks for three sectors, z, zz_g, and zz_i. In addition to the four data series in the benchmark, we introduce an extra time-series to do the Bayesian estimation, namely the ratio of consumption goods to consumption services. This data series helps us identify the relative productivity of goods versus services. The priors and distributions are the same as in the benchmark estimation.

Table 5 reports the estimated results together with business cycle statistics and variance decomposition. The variance decomposition shows that the demand shock θ_d still plays a significant role in accounting for business cycle fluctuations – about 20 percent of the variance of output, 35 percent of TFP, and more than half of the the variance of consumption. Note that the role of θ_d is sensitive to the specification of the technology and preference shocks. Estimating the model with different specifications of the technology shocks and preference shocks we found that the share

Table 5: Bayesian Estimation: Model with Storable Goods

Priors and Posteriors for the Shock Parameters (Likelihood = 2570.6)
Data used: Y, Y/N, Inv, P_i, C/S

Parameter	Density	Para(1)	Para(2)	Mean	90% Intv.
ρ_d	Beta	0.930	0.05	0.9958	[0.9918, 0.9995]
ρ_{z^g}	Beta	0.930	0.05	0.5634	[0.5327, 0.6001]
ρ_{z^i}	Beta	0.930	0.05	0.9946	[0.9909, 0.9988]
ρ_{z^s}	Beta	0.930	0.05	0.8799	[0.8569, 0.8996]
ρ_n	Beta	0.930	0.05	0.9825	[0.9742, 0.9899]
σ_d	Inv Gamma	0.150	Inf	0.5631	[0.5152, 0.6179]
σ_{z^g}	Inv Gamma	0.002	Inf	0.0274	[0.0250, 0.0302]
σ_{z^i}	Inv Gamma	0.002	Inf	0.0111	[0.0100, 0.0122]
σ_{z^s}	Inv Gamma	0.002	Inf	0.0125	[0.0114, 0.0137]
σ_n	Inv Gamma	0.020	Inf	0.0176	[0.0162, 0.0189]

	Data (1967Q1-2013Q4)			Model			Var Decomposition				
	var(x)	$\frac{var(x)}{var(Y)}$	Corr	var(x)	$\frac{var(x)}{var(Y)}$	Corr	θ_d	z^g	z^i	z	θ_n
Y	4.68	1.00	1.00	4.39	1.00	1.00	19.71	2.09	3.61	68.51	6.08
Y/N	1.47	0.31	0.50	5.22	1.19	0.88	41.26	0.46	1.22	54.45	2.62
Investment	48.48	10.36	0.95	29.25	6.66	0.77	1.43	2.05	23.13	69.45	3.93
P_i	0.70	0.15	-0.01	1.18	0.27	0.23	0.02	21.76	62.69	14.70	0.83
Goods/service	14.14	3.02	0.82	63.80	14.53	0.31	2.89	87.31	0.60	8.94	0.26
Labor	3.52	0.75	0.83	1.14	0.26	0.07	25.26	2.04	2.05	2.14	68.51
TFP	2.05	0.44	0.79	4.73	1.08	0.94	35.04	0.86	1.92	62.07	0.11
Consumption	1.58	0.34	0.93	3.27	0.74	0.88	51.67	7.32	0.89	35.61	4.51
Goods	19.13	4.09	0.87	66.73	15.20	0.48	2.15	81.06	0.81	15.17	0.81
Services	0.70	0.15	0.88	2.85	0.65	0.82	65.15	0.41	0.53	29.47	4.44
$\frac{Inventory}{sale}$	2.44	0.52	-0.42	57.64	13.13	-0.58	29.13	59.16	0.49	10.66	0.57
$\frac{P_{goods}}{P_{service}}$	1.47	0.31	-0.19	4.24	0.97	0.69	62.02	9.95	0.51	26.81	0.72

The table shows the Bayesian estimation and the business cycle properties for the shopping model with storable goods and services. The service share of consumption is 0.92. All variables are HP-filtered logarithms of the original series. The estimation targets the data series of output, labor productivity, investment, investment price, and the ratio of goods to services. *Goods* are measured as durable consumption goods expenditures in NIPA, and *Services* as consumption services plus non-durable consumption goods. The sample is quarterly U.S. data, from Q1 in 1967 to Q4 in 2013. Note that the variance $var(x)$ is multiplied by a factor of 10,000.

of output accounted for by search-related preference shocks range from 18 to 35 percent. See Table 6 in Appendix G for details.[25] Note also that the shopping friction is quantitatively more relevant for services than for durable consumption goods: according to the variance decomposition the "demand shock" θ_d accounts for 65 percent of the variance of services and just two percent of the variance of durables (cf. Table 5).

The extended model has richer predictions than the benchmark model. In particular, we study the dynamics of goods, services, and the inventories-to-sales ratio. In the data, aggregate purchases of durable goods are much more volatile than services and non-durable goods, and the model has similar predictions. Note that this empirical success of the model hinges on the assumption that both firms and households can store durables. To see this, compare the relative variance of goods to services in economy BM5 relative to the benchmark (BM1) in Table 6. If only firms could store goods (BM5), then the variance of goods relative to the variance of services in the reestimated model would be just one tenth of the ratio in the benchmark economy.[26] Moreover, the inventories-to-sales ratio in the model is countercyclical as in the data, although it is more volatile in the model than that in the data.

In our benchmark analysis we have used durable consumption goods as the empirical counterpart to our notion of consumption goods. However, many so-called non-durable goods can be stored a limited period of time and might be perceived as being part of storable consumption goods. As a sensitivity analysis we assume that the empirical counterpart to goods in the model is durable goods plus non-durable goods. The series used for consumption goods and services in the Bayesian estimation are redefined accordingly. In this case the share of storable goods in consumption expenditures is 40 percent. When reestimating the model the empirical time series for goods and services are modified accordingly. With a larger share of storable goods the preference shocks generally play a smaller role: across various specifications for the technology shocks and preference shocks, the preference shocks account for between 6.5 percent and 18 percent of output, see Table 6 in Appendix G for details.

[25]For example, if we were to assume independent productivity shocks for each good, then θ_d would account for a larger share of aggregate fluctuations – 34.5 and 44.2 percent of output and TFP, respectively.

[26]The low relative volatility of services in the benchmark model also depends on the way the technology shocks are specified. As it turns out, it is only in the benchmark specification – when the TFP shock on service production is the *neutral* technology shock – that the model generates a variance of durable good expenditure which is 20-25 times larger than the variance of service consumption, similar to the difference in the U.S. data. For example, if the neutral technology shock were to apply to goods production and TFP of service production were zz_s, then services become almost as volatile as consumption goods (see Economy BM3 in Table 5). In this case the estimation assigns a very volatile neutral technology shock and volatile θ_d shock in order to make consumption goods sufficiently volatile. This in turn makes services substantially more volatile.

Summing up, the main results from the benchmark model are robust to allowing firms to hold inventories of consumer durables and to let households store such goods: demand shocks are a major driver of aggregate consumption, TFP, and output. In addition, the extension with inventories is consistent with a number of stylized facts on cyclical properties of services, consumption goods, and inventories. In particular, the model replicates the correlation between the inventory to sales ratio and output, and services is substantially less volatile than consumption durables.

7 Conclusion

This paper provides a business cycle theory with an explicit productive role for the demand for goods. A search friction prevents perfect matching between producers and potential consumers. A larger consumer demand for consumption is associated with more intense search for goods and, hence, a larger utilization of the potential production. Shocks that cause changes in search effort will therefore induce changes in aggregate output even if standard factors such as capital and labor are fixed. Thus, when applying a neoclassical production function that ignores search effort as an input, changes in demand and search effort generate procyclical movements in measured TFP.

A competitive search protocol resolves the matching friction and the equilibrium outcome is efficient and unique. The framework is otherwise a standard neoclassical model with flexible prices.

Our main quantitative exercise is to estimate the model using standard Bayesian techniques. Business cycles are driven by preference shocks to search effort and labor supply, true technology shocks, and investment-specific shocks. Preference shocks affecting search effort account for a large share of the fluctuations in consumption, GDP, and the Solow residual. This finding is robust to extending the model to allow for varying capacity utilization of capital and to allow firms and households to store consumption goods as inventories and durables.

While our notion of search effort for consumption goods is unobservable, we note that the shopping effort implied by our estimated model turns out to be correlated with the measured *time* U.S. households devote to shopping. Interestingly, both the shopping effort in our model and the average empirical shopping time for U.S. households experienced a sharp decline during the Great Recession. According to our model, this fall is needed to account for the large and persistent fall in consumption as well as the suppressed relative price of investment after 2006 in the U.S.

Our paper is consistent with Keynes' idea that consumer demand can have real effects. We show that this holds true even in a neoclassical model with flexible prices, amended with a product market matching friction.

In future work, we plan to extend this environment to contexts where the demand shocks are generated by financial frictions, government expenditures, or foreign demand shocks. It would also be interesting to consider additional frictions that could break the efficient outcome of the competitive search model, such as coordination failures or additional labor market frictions. Finally, it is straightforward to embed our search friction for goods within the New Keynesian and Mortensen-Pissarides approaches to fluctuations, which assume frictions in either price setting or labor markets to generate large fluctuations in output and hours worked. Ultimately, these two traditions build on technology shocks as a major source of fluctuations, and our findings provide a rationale for substituting productivity shocks for demand shocks in these models.

References

AGUIAR, M., E. HURST, AND L. KARABARBOUNIS (2013): "Time Use During the Great Recession," *American Economic Review*, 103, 1664–1696.

ALESSANDRIA, G. (2005): "Consumer Search, Price Dispersion, and International Relative Price Volatility," Working Paper No. 05-9 Federal Reserve Bank of Philadelphia.

BAI, Y. AND J.-V. RÍOS-RULL (2015): "Demand Shocks and Open Economy Puzzles," *American Economic Review Papers and Proceedings*, 105, 644–649, american Economic Review Papers and Proceedings.

BASU, S. (1996): "Procyclical Productivity: Increasing Returns or Cyclical Utilization?" *Quarterly Journal of Economics*, 111, 719–751.

BILS, M. (2016): "Deducing markups from stockout behavior," *Research in Economics*, 70, 320–331.

BILS, M. AND J.-O. CHO (1994): "Cyclical Factor Utilization," *Journal of Monetary Economics*, 33, 319–354.

BURNSIDE, C., M. EICHENBAUM, AND S. REBELO (1993): "Labor Hoarding and the Business Cycle," *Journal of Political Economy*, 101, 245–273.

CHRISTIANO, L., M. EICHENBAUM, AND C. EVANS (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45.

DURAS, J. (2015): "Goods Market Search and the Labor Wedge," Manuscript, Texas Tech University.

——— (2016): "A Search Model of Unemployment and Inventories," Manuscript, Texas Tech University.

FAGNART, J. F., O. LICANDRO, AND F. PORTIER (1999): "Firm heterogeneity, capacity utilization, and the business cycle," *Review of Economic Dynamics*, 2, 433–455.

FAIG, M. AND B. JEREZ (2005): "A theory of commerce," *Journal of Economic Theory*, 122, 60–99.

Fisher, J. D. (2006): "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks," *Journal of Political Economy*, 114, 413–451.

Floetotto, M. and N. Jaimovich (2008): "Firm Dynamics, Markup Variations and the Business Cycle," *Journal of Monetary Economics*, 55, 1238–1252.

Francis, N. and V. A. Ramey (2005): "Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited," *Journal of Monetary Economics*, 52, 1379–1399.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988): "Investment, Capacity Utilization and the Real Business Cycle," *American Economic Review*, 78, 402–417.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997): "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 83, 342–362.

Heathcote, J., F. Perri, and G. Violante (2010a): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967- 2006," *Review of Economic Dynamics*, 1, 15–51.

Heathcote, J., K. Storesletten, and G. L. Violante (2010b): "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy*, 118, 681–722.

Herrendorf, B., R. Rogerson, and Á. Valentinyi (2013): "Two Perspectives on Preferences and Structural Transformation," *American Economic Review*, 103, 2752–2789.

Huo, Z. and J.-V. Ríos-Rull (2013): "Paradox of Thrift Recessions," NBER Working Paper No. 19443.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010): "Investment Shocks and Business Cycles," *Journal of Monetary Economics*, 57, 132?145.

Kaplan, G. and G. Menzio (2016): "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations," *Journal of Political Economy*, 124, 771–825.

King, R. G. and S. Rebelo (1999): "Resuscitating Real Business Cycles," in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, North Holland, chap. 14, 927–1007.

Krusell, P., L. E. Ohanian, J. V. Ríos-Rull, and G. L. Violante (2000): "Capital-skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 68(5), 1029–1054.

Kydland, F. E. and E. C. Prescott (1988): "The Workweek of Capital and its Cyclical Implications," *Journal of Monetary Economics*, 21, 343–360.

Lagos, R. (2006): "A Model of TFP," *Review of Economic Studies*, 73, 983–1007.

Licandro, O. and L. A. Puch (2000): "Capital utilization, maintenance costs and the business cycle," *Annales d'Economie et de Statistique*, 143–164.

Michaillat, P. and E. Saez (2015): "Aggregate Demand, Idle Time, and Unemployment," *Quarterly Journal of Economics*, 130, 507–569.

Moen, E. R. (1997): "Competitive Search Equilibrium," *Journal of Political Economy*, 105, 385–411.

Nakajima, T. (2005): "A Business Cycle Model with Variable Capacity Utilization and Demand Disturbances," *European Economic Review*, 49, 1331?1360.

Petrosky-Nadeau, N. and E. Wasmer (2015): "Macroeconomic Dynamics in a Model of Goods, Labor and Credit Market Frictions," *Journal of Monetary Economics*, 72, 97–113.

Petrosky-Nadeau, N., E. Wasmer, and S. Zeng (2016): "Shopping Time," *Economic Letters*, 143, 52–60.

Richardson, H. (1995): "Control your costs then cut them," *Transportation and Distribution*, 36, 94–96.

Rudanko, L. and F. Gourio (2014): "Customer Capital," *Review of Economic Studies*, 81, 1102–1136.

Swanson, E. T. (2006): "The relative price and relative productivity channels for aggregate fluctuations," *Contributions in Macroeconomics*, 6, 1–39.

Wen, Y. (2004): "What Does It Take to Explain Procyclical Productivity?" *B.E. Journal of Macroeconomics (Contributions)*, 4, Article 5.

Appendix

A Proof of Proposition 1

Part 1. We show that aggregate search satisfies the intra-temporal condition (7).

The firm's problem is static, and given by $\max_{\{p,q\}} \Pi(p,q)$ subject to (5). Let the Lagrangian multiplier on (5) be λ and the household's optimal shopping effort be d^*. The first-order conditions over q and p are given by

$$\begin{aligned} p\varphi q^{\varphi-1} Z &= (1-\varphi)\, q^{\varphi-2} Z d^* \lambda \left[u_c - \frac{p}{P_a} \beta\, \mathrm{E}\left\{ V_a\left(1 + \frac{\pi - p d^* \Psi_d(q) Z}{P_a}, \theta'\right) \,\middle|\, \theta \right\} \right] \\ q^{\varphi} Z &= \frac{d^* q^{\varphi-1} Z}{P_a} \lambda \beta\, \mathrm{E}\left\{ V_a\left(1 + \frac{\pi - p d^* \Psi_d(q) Z}{P_a}, \theta'\right) \,\middle|\, \theta \right\}. \end{aligned}$$

Combining the above two first-order conditions yields

$$(1-\varphi)\, u_c = \frac{p}{P_a} \beta\, \mathrm{E}\left\{ V_a\left(1 + \frac{\pi - p d^* \Psi_d(q) Z}{P_a}, \theta'\right) \,\middle|\, \theta \right\}. \tag{A.1}$$

We now find the optimal search effort by differentiate equation (4),

$$u_c + u_d \frac{1}{\Psi_d(q) Z} = \frac{p}{P_a} \beta\ \mathrm{E}\left\{ V_a\left(1 + \frac{\pi - p d^* \Psi_d(q) Z}{P_a}, \theta'\right) \,\middle|\, \theta \right\}. \tag{A.2}$$

Combining (A.1)-(A.2) yields the intra-temporal first-order condition (7) with $D = q$ in equilibrium.

Part 2. The Euler equation for a' in the household's problem is given by,

$$u_c \Psi_d(q) Z + u_d = \beta\, \mathrm{E}\left\{ \frac{P_a(\Lambda') + \pi(\Lambda')}{P_a(\Lambda)} \left[u_c(\Lambda') \Psi_d(q(\Lambda')) Z + u_d(\Lambda') \right] \right\}.$$

Plugging the intratemporal condition (7) into the above Euler equation, we can show the usual Euler equation (8) holds.

B Planner Problem, Competitive Equilibrium, and the First Welfare Theorem in the Production Economy

In the text we appeal to a version of the First Welfare Theorem and focus on the planner problem. This Appendix characterizes the planner problem and the competitive equilibrium, and proves that the allocation implied by the decentralized competitive equilibrium is equivalent to the planner allocation.

B.1 Derivations of optimality condition for the planner

The optimality conditions follow from the first-order conditions of the planner problem (18), rewritten as

$$\begin{aligned}\max_{D,N_c,N_i,K_c'} \; & u\left(D\,\Psi_d\left[D\right]\; zf\left(K_c,N_c\right),D,N_c+N_i,\Lambda\right)\\ & +\;\beta\;\mathrm{E}\{W\left(\theta',K_c',z_i\;z\;f\left(K_i,N_i\right)+\left(1-\delta\right)\left(K_c+K_i\right)-K_c'\right)\mid\Lambda\}.\end{aligned}$$

The first-order conditions over D, N_c, N_i, and K_c' are given by

$$u_c\,\frac{\partial\Psi_T\left[D\right]}{\partial D}\;zf\left(K_c,N_c\right)+u_d \;=\; 0 \tag{B.1}$$

$$u_n+u_c\;\Psi_T\left[D\right]\;z\frac{f\left(K_c,N_c\right)}{\partial N_c} \;=\; 0 \tag{B.2}$$

$$u_n+z_i\;z\;\frac{\partial f\left(K_i,N_i\right)}{\partial N_i}\beta\;\mathrm{E}\left\{\frac{\partial W\left(\Lambda'\right)}{\partial K_i'}\;\middle|\;\Lambda\right\} \;=\; 0 \tag{B.3}$$

$$\beta\;\mathrm{E}\left\{\frac{\partial W\left(\Lambda'\right)}{\partial K_c'}-\frac{\partial W\left(\Lambda'\right)}{\partial K_i'}\;\middle|\;\Lambda\right\} \;=\; 0. \tag{B.4}$$

And the envelope conditions are

$$\frac{\partial W\left(\Lambda\right)}{\partial K_c} \;=\; u_c\;\Psi_T\left[D\right]\;z\frac{\partial f\left(K_c,N_c\right)}{\partial K_c}-\frac{\left(1-\delta\right)u_n}{z_i\;z\;\frac{\partial f_n\left(K_i,N_i\right)}{\partial N_i}} \tag{B.5}$$

$$\frac{\partial W\left(\Lambda\right)}{\partial K_i} \;=\; -\left(z_i\;z\;\frac{\partial f\left(K_i,N_i\right)}{\partial K_i}+1-\delta\right)\;\frac{u_n}{z_i\;z\;\frac{\partial f\left(K_i,N_i\right)}{\partial N_i}}. \tag{B.6}$$

Optimality conditions (19)-(20) follow directly from equations (B.1)-(B.2). Combining equations (B.4)-(B.5)-(B.6) with (B.2) yields equation (21):

$$0 = \mathrm{E}\left\{ \frac{\partial W(\Lambda')}{\partial K_c'} - \frac{\partial W(\Lambda')}{\partial K_i'} \,\middle|\, \Lambda \right\} = \mathrm{E}\left\{ u_{n'} \left(\frac{\frac{\partial f\left(K_i', N_i'\right)}{\partial K_i'}}{\frac{\partial f\left(K_i', N_i'\right)}{\partial N_i'}} - \frac{\frac{\partial f(K_c', N_c')}{\partial K_c'}}{\frac{\partial f(K_c', N_c')}{\partial N_c'}} \right) \,\middle|\, \Lambda \right\}.$$

Euler equation (22) follows from plugging (B.6) into equation (B.3).

B.2 Competitive Equilibrium

Consider the household problem in the benchmark economy, equations (11)-(14) in the text. Let the multiplier on the household's budget constraint (13) be denoted λ. The first-order conditions for the households can then be expressed as,

$$\begin{aligned}
\lambda \hat{p} &= u_c + \frac{u_d}{\Psi_d(\hat{q})\hat{f}} \qquad \text{(B.7)}\\
\lambda p_i &= \beta E \lambda' [R_c(\Lambda') + p_i(\Lambda')(1-\delta) \mid \Lambda] \\
\lambda p_i &= \beta E \lambda' [R_i(\Lambda') + p_i(\Lambda')(1-\delta) \mid \Lambda] \\
\lambda W &= u_n \\
V_{\hat{p}} &= -\lambda d A \hat{q}^{\varphi-1} \hat{f} \\
V_{\hat{q}} &= (\varphi - 1) d A \hat{q}^{\varphi-2} \hat{f} (u_c - \lambda \hat{p}) \\
V_{\hat{f}} &= u_c \, d A \hat{q}^{\varphi-1} (u_c - \lambda \hat{p}).
\end{aligned}$$

Now let's consider the problem of a consumption-producing firm (16). Let the multiplier for the participation constraint be γ and for the production constraint be μ. We can write the firm's first-order conditions as

$$\begin{aligned}
A \hat{q}^{\varphi} \hat{f} + \gamma V_{\hat{p}} &= 0 \\
\varphi A \hat{q}^{\varphi-1} \hat{p} \hat{f} + \gamma V_{\hat{q}} &= 0 \\
\hat{p} A \hat{q}^{\varphi} - \mu + \gamma V_{\hat{f}} &= 0 \\
W &= \mu z \frac{\partial f(k,n)}{\partial n} \\
R_c &= \mu z \frac{\partial f(k,n)}{\partial k}
\end{aligned}$$

The first-order conditions for the investment-producing firms are standard and are given by

$$\begin{aligned} W &= p_i z z_i \frac{\partial f(k,n)}{\partial n} \\ R_i &= p_i z z_i \frac{\partial f(k,n)}{\partial k}. \end{aligned} \tag{B.8}$$

Equilibrium In equilibrium, $p = \hat{p} = 1$, $q = \hat{q}$, and $F = \hat{f}$. The competitive equilibrium consists of 15 functions, three functions from the households' problem, $\{C, D, N\}$, five functions from the firms' problems, $\{K_c', K_i', N_c, N_i, \pi\}$, one function from competitive search problem, $\{q\}$, and six price functions, $\{W, p, P_a, p_i, R_c, R_i\}$, which simultaneously satisfy the following functional equations.

First, the household problem yields the following functional equations,

$$u_c = \beta \, \mathrm{E}\left\{ \frac{R_c(\Lambda') + p_i(\Lambda')(1-\delta)}{p_i(\Lambda')} u_c' \mid \Lambda \right\} \tag{B.9}$$

$$u_c = \beta \, \mathrm{E}\left\{ \frac{R_i(\Lambda') + p_i(\Lambda')(1-\delta)}{p_i(\Lambda')} u_c' \mid \Lambda \right\} \tag{B.10}$$

$$(1-\varphi) u_c = \frac{u_n}{W} \tag{B.11}$$

$$-u_d = \varphi u_c D^{\varphi-1} f(K_c, N_c). \tag{B.12}$$

Second,the firms' problems yield four functional equations,

$$(1-\varphi) W = \Psi_T(q) z \frac{\partial f(K_c, N_c)}{\partial N_c} \tag{B.13}$$

$$\frac{W}{R_c} = \frac{\partial f(K_c, N_c)/\partial N_c}{\partial f(K_c, N_c)/\partial K_c} \tag{B.14}$$

$$\frac{W}{p_i} = z z_i \frac{\partial f(K_i, N_i)}{\partial N_i} \tag{B.15}$$

$$\frac{W}{R_i} = \frac{\partial f(K_i, N_i)/\partial N_i}{\partial f(K_i, N_i)/\partial K_i}. \tag{B.16}$$

Finally, market clearing yields

$$
\begin{aligned}
q &= D \\
C &= \Psi_T(q) z f(K_c, N_c) \\
z_i z f(K_i, N_i) &= K_c' + K_i' - (1-\delta)(K_c + K_i) \\
N &= N_c + N_i.
\end{aligned}
$$

Incidentally, by combining equations (B.13) and (B.15) we obtain the expression for the relative price of investment, equation (23) in the text.

B.3 A First Welfare Theorem in the Production Economy

We now show that the competitive equilibrium corresponds to the planner allocation. To this end, we show that the first-order conditions of the competitive equilibrium are the same as the ones in the planner problem.

First, equation (B.12) in the equilibrium is identical to equation (19) in the planner problem. Second, combining equation (B.11) and (B.13) yields the first-order condition on labor (20) in the planner problem. Finally, equation (21) in the planner problem holds in equilibrium. To see this, take the difference between (B.9) and (B.10) from the household's problem, and get

$$0 = \mathrm{E}\left\{ \left(R_i(\Lambda') - R_c(\Lambda')\right) \frac{u_c'}{p(\Lambda')} \mid \Lambda \right\}$$

Now we replace u_c with u_n/W using equation (B.11) and obtain

$$0 = \mathrm{E}\left\{ \left(R_i(\Lambda') - R_c(\Lambda')\right) \frac{u_n'}{W(\Lambda')} \mid \Lambda \right\}.$$

Then replace R/W with firms' first-order conditions (B.14) and (B.16) to obtain the planner's Euler equation.

C Details about the Calibration of the Production Economy

Depreciation δ. Let the investment share of output be $\kappa = p_I I/Y = 20\%$ and the size of the aggregate capital-output ratio $p_I K/Y = 2.75$. This allows us to pin down the depreciation rate from the steady-state relation $\delta K = I$, so

$$\delta = \frac{p_I I/Y}{p_I K/Y} = \frac{0.2}{2.75} = 7.3\%$$

Labor's share θ_n: for an investment firm the labor's share of output is simply θ_n (proof omitted). Consider now the share of output for a consumption-producing firm that is paid to labor. Note that the equilibrium price of the consumption good is given by

$$\widehat{p} = (1-\varphi)\frac{W}{\Psi_T(q) z \frac{\partial f(K_c, N_c)}{\partial N_c}}, \tag{C.1}$$

where $\widehat{p} \equiv 1$ since consumption is the numéraire good. The wage bill is WN_c. Total (measured) value added is the sales $\Psi_T(q)F$, where $F = zf(K_c, N_c)$, so labor's share of output in consumption firms is

$$\frac{WN_c}{C} = \frac{WN_c}{\Psi_T(q)zf(K_c, N_c)} \cdot \frac{\frac{\partial f}{\partial N_c}}{\frac{\partial f}{\partial N_c}} = \frac{WN_c}{(1-\varphi)W} \cdot \frac{\frac{\partial f}{\partial N_c}}{f(K_c, N_c)} = \frac{\theta_n}{1-\varphi},$$

where the second equality follows from equation (C.1) and the third equation from the assumption that f is Cobb-Douglas so $\theta_n f(K_c, N_c) = N_c \frac{\partial f}{\partial N_c}$. The aggregate labor share is then given by,

$$\frac{WN}{Y} = \theta_n \frac{1-\kappa\varphi}{1-\varphi}.$$

With an aggregate share of output that goes to labor $WN/Y = 0.67$ and $\kappa = 0.2$, we have $\theta_n = 0.54$.

Capital's share θ_k. The share of output that goes to reproducible capital is

$$\frac{RK}{Y} = \theta_k \frac{1-\kappa\varphi}{1-\varphi}.$$

From households' first order conditions, we knows that in steady state

$$\frac{R}{p_I} = \frac{1-\beta(1-\delta)}{\beta},$$

and $\beta = 1/(1+r)$. In addition, in steady state $R/W = (\partial f(N_c, K_c)/\partial K_c)/(\partial f(N_c, K_c)/\partial N_c)$. With $\varphi = 0.23$, $\delta = 7.3\%$, $r = 4\%$, $\kappa = 0.2$, and $p_I K/Y = 2.75$, this implies $\theta_k = 0.25$.

Weight on search effort χ_d. The capacity utilization in consumption is defined as the probability that a production unit gets matched, i.e., $\Psi_T[D] = D^{\varphi}$. A capacity utilization of 81% then implies $D = (0.81)^{1/0.23} = 0.4$. Given the expressions for u, Ψ_T and F, the intra-temporal first-order condition (19) imply

$$\chi_d = \varphi C\,(D)^{-\left(\frac{1}{\eta}+1\right)} = 0.23 * 0.8 * (0.4)^{-\left(\frac{1}{0.11}+1\right)} = 1905.$$

Mean TFP $Z = E(z)$ **and** $Z_I = E(z_i)$. To compute average TFPs, we first obtain the share of labor and capital in each sector. In steady state, the two sectors, consumption and investment, have the same ratio of wage to rental. Under the assumption of Cobb-Douglas of f, this implies $\frac{N_i}{K_i} = \frac{N_c}{K_c} = \frac{N}{K}$. In addition, we divided equation (B.13) by equation (B.15) and get

$$(1-\varphi)p_I = \frac{\Psi_T(q)z\partial f(K_c, N_c)/\partial N_c}{zz_i\partial f(K_i, N_i)/\partial N_i} = \frac{C}{I}\frac{N_i}{N_c},$$

which implies

$$\frac{N_i}{N_c} = \frac{(1-\varphi)\kappa}{1-\kappa}. \tag{C.2}$$

Note that we choose $p_I = 1$ in steady state, thus $I = \kappa$. Under our quarterly model, capital used for production $K = 4I/\delta = 11$. Taking the ratio of consumption over investment, we have

$$\frac{C}{I} = \frac{1-\kappa}{\kappa} = \Psi_T(q)\frac{1}{Z_I}\left(\frac{(1-\varphi)\kappa}{1-\kappa}\right)^{\alpha_k+\alpha_n}.$$

We can therefore calibrate Z_I,

$$Z_I = \frac{\kappa}{1-\kappa}\Psi_T(q)\left(\frac{(1-\varphi)\kappa}{1-\kappa}\right)^{-\alpha_k-\alpha_n} = \frac{0.2}{0.8} * 0.81 * \left(\frac{(1-0.23)*0.2}{0.8}\right)^{-0.25-0.54} = 0.74.$$

From equation (C.2) and the fraction of time spent on working in steady state $N = .3$, we know that $N_c = 0.25$ and $N_i = 0.05$. Similarly, $K_c = 0.83 * K$ and $K_i = 0.17 * K$. We can now back up the neutral technology in steady state

$$Z = \frac{C}{\Psi_T(D)f(K_c, N_c)} = \frac{0.8}{0.81K_c^{\alpha_k}N_c^{\alpha_n}} = 1.2.$$

D Proof of Proposition 2

Consider a static economy with J types of agents who differ in their consumption expenditure y_j, $j \in \{1, ..., J\}$. With a price level of $P = 1$, the aggregate expenditure must equal $\sum_j y_j \mu_j = C$, where μ_j is the population share of type j households, and C is aggregate consumption. There is a unit measure of firms, each of whom has a production function $f(k, n) = (k)^{\alpha_k} (n)^{\alpha_n}$. We start by showing that in equilibrium all firms supply the same capacity F.

Lemma 1. *All firms supply the same capacity of consumption goods, given by $F = zf(K_c, N_c)$.*

Proof. Since there are no externalities or distortions, it is straightforward that the welfare theorems apply. Accordingly, we formulate the problem as a planner problem,

$$\max \sum_j \xi_j u(C_j, D_j)$$

$$\text{s.t. } K_C = \sum_j K_j \qquad N_C = \sum_j N_j \qquad 1 = \sum_j T_j, \qquad C_j = D_j^{\varphi} T_j^{1-\varphi} zf(K_j, N_j),$$

where ξ_j denotes the planner weight on households of type j and K_j, N_j, and T_j denote the capital, labor, and share of firms allocated to production intended for group j. D_j and T_j represent the aggregate search effort in market j and the aggregate measure of firms catering to market j. We rewrite the problem as a Lagrange problem

$$\Theta = \sum_j \xi_j u\left(D_j^{\varphi} (T_j)^{1-\varphi} zf(K_j, N_j), D_j\right) - \lambda_K \left(\sum_j K_j - K_C\right) - \lambda_N \left(\sum_j N_j - N_C\right) - \Gamma \left(\sum_j T_j - 1\right)$$

Taking the first-order conditions yields,

$$-u_{d,j} = u_{c,j} \cdot \varphi \left(\frac{D_j}{T_j}\right)^{\varphi-1} zf(K_j, N_j)$$

$$\frac{\Gamma}{\lambda_K} = \frac{\xi_j u_{c,j} \cdot (1-\varphi) \left(\frac{D_j}{T_j}\right)^{\varphi} f(K_j, N_j)}{\xi_j u_{c,j} \cdot \left(\frac{D_j}{T_j}\right)^{\varphi} T_j \frac{\partial f(K_j, N_j)}{\partial K_j}} = (1-\varphi) \frac{f(K_j, N_j)}{T_j \frac{\partial f(K_j, N_j)}{\partial K_j}}$$

$$\frac{\lambda_K}{\lambda_N} = \frac{\xi_j u_{c,j} \cdot \left(\frac{D_j}{T_j}\right)^{\varphi} T_j \frac{\partial f(K_j, N_j)}{\partial K_j}}{\xi_j u_{c,j} \cdot \left(\frac{D_j}{T_j}\right)^{\varphi} T_j \frac{\partial f(K_j, N_j)}{\partial N_j}} = \frac{\frac{\partial f(K_j, N_j)}{\partial K_j}}{\frac{\partial f(K_j, N_j)}{\partial N_j}}.$$

When f is Cobb-Douglas, the last two conditions imply

$$\frac{\Gamma}{\lambda_K} = (1-\varphi)\frac{(K_j)^{\alpha_k}(N_j)^{\alpha_n}}{T_j\alpha_k(K_j)^{\alpha_k-1}(N_j)^{\alpha_n}} = \frac{1-\varphi}{\alpha_k}\frac{K_j}{T_j}$$

$$\frac{\lambda_K}{\lambda_N} = \frac{\alpha_k(K_j)^{\alpha_k-1}(N_j)^{\alpha_n}}{\alpha_n(K_j)^{\alpha_k}(N_j)^{\alpha_n-1}} = \frac{\alpha_k}{\alpha_n}\frac{N_j}{K_j}.$$

It follows that all firms have the same capital-labor ratio, equal to K_c/N_c, and all firms have the same output per location, $F = zf(K_c, N_c)$. □

We now solve the decentralized problem. Conjecture that there will be J different markets open, which all provide F but differ in the offered pair (p_j, q_j). Let $\bar{\pi}$ represent the expected revenue for a firm operating in the most profitable market. Profit maximization then imposes the following arbitrage condition on any offered (p_j, q_j),

$$p_j = \frac{\bar{\pi}}{F}\cdot(q_j)^{-\varphi} = \frac{\bar{\pi}}{F}\cdot\left(\frac{D_j}{T_j}\right)^{-\varphi}.$$

The search technology implies that for households of type j, $c_j = D_j^{\varphi}(T_j)^{1-\varphi}F$. We can thus obtain income of a household $y_j = p_j c_j = \bar{\pi}T_j$. Thus, since $\int T_j dj = 1$, aggregate revenue equals aggregate expenditure, $\bar{\pi} = \int y_j dj = C$. It follows that the measure of firms catering to type j households is equal to j's expenditure share $T_j = \mu_j y_j/C$. This in turn implies equation (26). Finally, the intra-temporal first-order condition yields equation (27).

E Grocery and Non-Grocery Shopping

Figure 3 shows the shopping time between 2003 and 2013 for the working age group (25-59 years old). The data source is American Time Use Survey (ATUS). Panel A shows the level for all shopping time (black-solid line), grocery shopping time (blue-circled line), and others shopping time (red-dashed line) of the working-age group. Panel B shows the percentage deviation of these categories.

Figure 3: Shopping

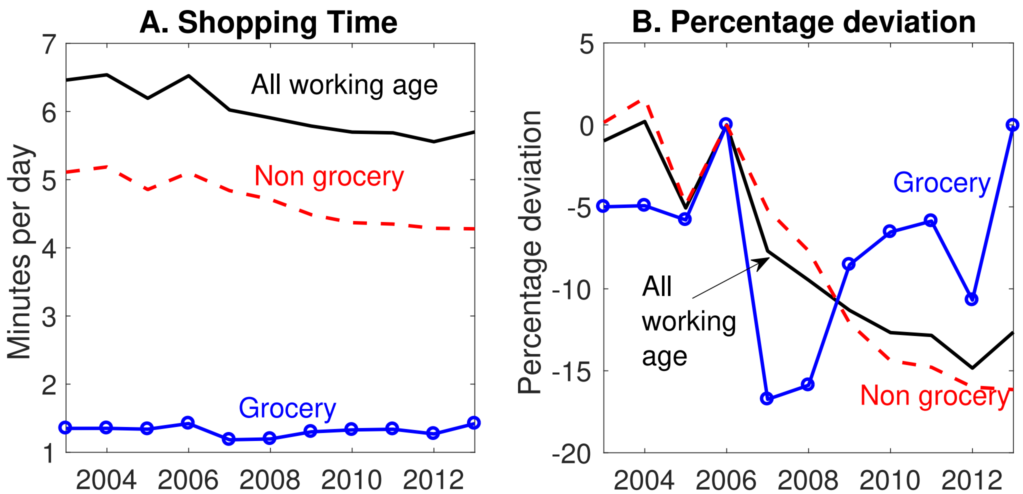

This figures shows the level and percentage deviation in shopping time between 2003 and 2013 for the working age group (25-59 years old). The data source is American Time Use Survey (ATUS). Panel A shows the level for all shopping time (black solid line), grocery shopping time (blue circled line), and other shopping time (red dashed line) of the working-age group. Panel B shows the percentage deviation of these categories.

F Identical inventory choice at the storable good model

In this subsection, we show that consumption goods producers with identical initial inventory choose the same holdings of next period's inventory, i.e.

$$x'_n(\Lambda, x) = x'_m(\Lambda, x). \tag{F.1}$$

A consumption good producer chooses $\{n_g, k_g, p_g, q_g, F_g, i_{x,u}, i_{x,m}, x'_m, x'_n\}$ to maximize

$$\begin{aligned}\Omega(\Lambda, x) = \max \quad & -W(\Lambda)n_g - R_g(\Lambda)k_g \\ & + \Psi_{T,g}(q_g)\left\{p_g F_g - p_x(\Lambda)i_{x,m} + E\left[M(\Lambda, \Lambda')\,\Omega(\Lambda', x'_m)|\Lambda\right]\right\} \\ & + (1 - \Psi_{T,g}(q_g))\left\{-p_x(\Lambda)i_{x,u} + E\left[M(\Lambda, \Lambda')\,\Omega(\Lambda', x'_n)|\Lambda\right]\right\}\end{aligned}$$

subject to the capacity constraint

$$F_g \leq zf(n_g, k_g) + (1 - \delta_f)x,$$

the inventory accumulation constraint for the unmatched state and for the matched state,

$$x'_n = zf(n_g, k_g) + (1-\delta_f)x + i_{x,u}, \quad \text{(F.2)}$$
$$x'_m = zf(n_g, k_g) + (1-\delta_f)x - F_g + i_{x,m}, \quad \text{(F.3)}$$

and the participation constraint of households. Let μ_n and μ_m be the multipliers on the unmatched constraint (F.2) and the matched inventory constraint (F.3) respectively. Taking the first order conditions on $i_{x,u}, i_{x,m}, x'_m, x'_n$, we have

$$\Psi_{T,g}(q_g)p_x(\Lambda) = \mu_m, \qquad (1-\Psi_{T,g}(q_g))p_x(\Lambda) = \mu_n$$
$$\Psi_{T,g}(q_g)E[M(\Lambda,\Lambda')\Omega_x(\Lambda', x'_m)] = \mu_m, \qquad (1-\Psi_{T,g}(q_g))E[M(\Lambda,\Lambda')\Omega_x(\Lambda', x'_n)] = \mu_n$$

Combining these four equations, we have

$$\frac{\Psi_{T,g}(q_g)E[M(\Lambda,\Lambda')\Omega_x(\Lambda', x'_m)]}{(1-\Psi_{T,g}(q_g))E[[M(\Lambda,\Lambda')\Omega_x(\Lambda', x'_n)]} = \frac{\mu_m}{\mu_n} = \frac{\Psi_{T,g}(q_g)p_x(\Lambda)}{(1-\Psi_{T,g}(q_g))p_x(\Lambda)}$$

which implies

$$\frac{E[[M(\Lambda,\Lambda')\Omega_x(\Lambda', x'_m)]}{E[[M(\Lambda,\Lambda')\Omega_x(\Lambda', x'_n)]} = 1.$$

It is therefore $x'_m = x'_n$.

G Robustness estimation on the storable good model

In this section, we report the robustness check over the estimation of the storable good model for alternative shock specifications and alternative definitions of storable goods. Table 6 presents the likelihood for the estimated models, the contribution of the θ_d shock on output, TFP, and consumption (according to the variance decomposition), as well as two moments unique to the storable good model. These moments are the relative variance of storable consumption goods to services and the correlation of the inventory-to-sale ratio with output. The Supplementary Appendix reports the details of the various estimated models, including the estimated parameter values, business cycle statistics, and the complete variance decompositions for all variables.

We consider four alternative specifications for the technology and preference shocks. Economy 1 has the benchmark specification where service production is subject to a neutral productivity shock (labeled z_s), consumption good production has a TFP given by the neutral TFP shock times a good-specific shock so that TFP is $z_s z_g$, and investment production has TFP $z_s z_i$ as in the benchmark model in Sections 3-5. Economy 2 has two types of shopping disutility shocks, one applying to search for storable goods, θ_d^g, and one applying to search for services, θ_d^s. With some abuse of notation we record the sum of variance contribution attributable to θ_d^g and θ_d^s in the column labeled θ_d. Moreover that economy has only two productivity shocks: a neutral TFP shock z_s applying to both consumption goods and services and the standard specification $z_s z_i$ for investment production. Economy 3 has z_g as the neutral technology shock, so TFP for service production becomes $z_s z_g$. Finally, Economy 4 assumes that TFP for the three production functions are independent, i.e., z_s on service production, z_g on goods production, and z_i on investment production.

The upper panel of Table 6, "Benchmark definition of storable goods: Durable goods", reports the statistics for the benchmark calibration where the empirical counterpart to storable goods is durable consumption goods and the consumption expenditure share of these goods is calibrated to eight percent. Our benchmark specification – Economy BM1 – has the highest likelihood and matches best the observed relative variance of goods to service. Across specifications the θ_d shock(s) account for 18 to 35 percent of output variation, at least one third of TFP variation, and more than 50 percent of consumption variation. All specifications feature countercyclical inventory to sales ratios.

Economy BM5 considers a version of the storage economy where the quarterly depreciation rate of consumption goods is set to 99 percent. The specification of shocks is the same shock specification as BM1. This economy is intended to capture a counterfactual economy where firms can store goods but households are effectively prevented from storing goods.

Economies BM6 and BM7 reestimates the storage model but with just fours shocks, similar to the benchmark model in Sections 3, and targets the same data series as in 5 (GDP, labor productivity, aggregate consumption and the relative price of investment to goods). Economy BM6 allows households to store the storable good (and with $\delta_h = 0.05$) while Economy BM7 effectively rules out storage (with $\delta_h = 0.99$)).

The lower panel of Table 6, "Alternative definition of storable goods: Durable plus non-durable goods", reports the statistics for an alternative calibration where the empirical counterpart to storable goods is non-durable plus durable consumption goods. In this case the expenditure share of these goods is calibrated to 40 percent, equal to the NIPA expenditure share on non-durable plus durable consumption goods. Note that the data series used in the estimation are adjusted accordingly. Therefore, the likelihood is not comparable across the upper and lower panels.

We conduct four different shock specifications also under this alternative calibration. Clearly, with a large expenditure share of storable goods, the θ_d shock plays a smaller role. The contribution varies from 7 percent to 18 percent for output, 16 percent to 33 percent for TFP, and 22 percent to 45 percent for consumption. With goods including nondurable consumption, the relative goods volatility reduces from 27 times to 7 times of service volatility. The benchmark shock specification still has the highest likelihood and a relative variance of consumption goods to services closest to the data.

Table 6: Robustness Estimation on Storable Good Model

			contribution of θ_d				
Economy	Shocks	Likelihood	Y	TFP	Con.	$\frac{Var(goods)}{Var(services)}$	Corr(Y, $\frac{inventories}{sales}$)
Benchmark definition of storable goods: Durable goods							
Data						27.33	-0.42
BM1	$z_s, z_s z_g, z_s z_i, \theta_d, \theta_n$	2570.57	19.71	35.04	51.67	23.41	-0.58
BM2	$z_s, z_s, z_s z_i, \theta_d^g, \theta_d^s, \theta_n$	2553.36	18.88	33.85	55.52	19.36	-0.13
BM3	$z_g z_s, z_g, z_g z_i, \theta_d, \theta_n$	2511.71	21.57	29.63	38.75	1.31	-0.62
BM4	$z_s, z_g, z_i, \theta_d, \theta_n$	2452.48	34.46	44.15	49.78	3.68	-0.49
BM5	$\delta_h = 0.99$, BM1 shocks	2023.96	18.31	33.02	61.37	2.50	-0.03
BM6	$z_s, z_s, z_i, \theta_n, \theta_d$	2335.44	18.34	33.96	52.39	4.50	-0.66
BM7	$\delta_h = 0.99$, BM6 shocks	2322.48	20.61	41.04	60.71	0.31	-0.19
Alternative definition of storable goods: Durable plus non-durable goods							
Data				7.12	-0.42		
ND1	$z_s, z_s z_g, z_s z_i, \theta_d, \theta_n$	2835.17	6.51	16.07	22.18	8.02	-0.63
ND2	$z_s, z_s, z_s z_i, \theta_d^g, \theta_d^s, \theta_n$	2820.89	7.11	17.15	45.02	5.84	-0.21
ND3	$z_g z_s, z_g, z_g z_i, \theta_d, \theta_n$	2680.90	7.97	19.28	31.88	1.42	-0.57
ND4	$z_s, z_g, z_i, \theta_d, \theta_n$	2699.52	18.03	32.89	38.04	1.80	-0.49

The table summarizes key statistics for the storable good model for different specifications of productivity and preference shocks and for different calibrations of the share of storable goods. The upper panel reports the results for the benchmark calibration where the empirical counterpart to storable goods is durable consumption goods. The lower panel, "Alternative definition of storable goods: Durable plus non-durable goods" reports the results for an alternative calibration where the empirical counterpart to the storable good is durable plus non-durable consumption goods. The specifications of the technology and preference shocks are reported in the following sequence: (1) TFP shock for service production, (2) TFP shock for storable consumption good production, (3) TFP shock for investment production, (4) shock(s) to the disutility of shopping, and (5) preference shock for leisure. Economies BM3 and ND3 have two shopping disutility shocks, θ_d^g and θ_d^s. In this case the column labeled θ_d records the sum of variance contribution attributable to θ_d^g and θ_d^s. Economy BM5 has the same shock specification as BM1 but prevents households from storing goods (with $\delta_h = 0.99$). Economy BM6 and BM7 has the same shock specification as our benchmark (non-storage) economy and estimates the model with the four data series used in the benchmark non-storage economy (Y, Y/N, P_i, and C).