Consumer good search: theory and evidence

Saman Darougheh[*]

This version: March 26, 2021

Abstract

I develop a model of the consumer good market where the individual's search decision is consistent with balanced-growth preferences. Here, optimal search is independent of income but increases with the time endowment. I characterize the potentially multiple equilibria and test whether the model can replicate differences in observed shopping behavior between employed and unemployed individuals. I use the American Time Use Survey to show that unemployed individuals have an almost 50% larger time endowment available for leisure and shopping and spend 27% more time shopping than the employed. In the calibrated model, however, unemployed households will spend around twice as much time shopping as employed households. I argue that micro-founded good search models are not yet ready for business cycle analysis, and discuss ways of reconciling the model with the data. (JEL D11, E21, E32, L11)

[*]Danmarks Nationalbank. Views are mine only. I thank Mark Aguiar, Karl Harmenberg, Per Krusell, Hannes Malmberg, Kurt Mitman, and Erik Oberg for valuable feedback, and Tim Maurer for excellent research assistance.

The presence of search frictions impedes the consumers' ability to compare prices. As a result, firms may charge higher markups. A corollary of this well-known result is that cyclical search may render firms' markups cyclical.

In recent years, macroeconomists have used this insight to integrate consumer good search into their analysis. For example, Qiu and Ríos-Rull (2019) show that search frictions in the good markets will deliver pro-cyclical markups in an otherwise standard New Keynesian framework. Integrating good search frictions into a Diamond-Mortensen-Pissarides framework will greatly increase the sensitivity of firm profits to the unemployment rate: Kaplan and Menzio (2016) assume different good search intensities for the employed and the unemployed, and show that this even permits multiple steady states in their calibration.

Both papers are representative for a larger literature that has taken the important first step of demonstrating the relevance of consumer good search in explaining aggregate phenomena. So far, there have been two predominant approaches to the integration of the household's search decision. Some papers assume that the search intensity is governed by an exogenous parameter (Alessandria, 2009; Kaplan and Menzio, 2016; Head et al., 2012; Menzio and Trachter, 2018). Others model the search decision in a reduced form by placing the search intensity inside the utility function (Bai, Rios-Rull, and Storesletten, 2011; Qiu and Ríos-Rull, 2019; Petrosky-Nadeau, Wasmer, and Zeng, 2014).

In this paper, I consider the possibility that, in searching for goods, the household has to trade off time spent searching against time enjoying leisure. This trade-off allows me to discipline the model by matching the search intensities of households with varying marginal utility of leisure. With such a setting, the time allocation to search and leisure can be responsive to economic policy as well as to changes in the macroeconomic environment. In fact, the time allocation has undergone significant swings over the past 15 years for both the employed and the unemployed, as illustrated in Figure 1 (see also Aguiar, Hurst, and Karabarbounis, 2013).

I propose a version of the Burdett and Judd (1983) framework that microfounds the household's good search decision while preserving most of the features of the original paper. The households are endowed with a fixed unit of time and income. They have to trade off spending time on leisure against shopping for consumer goods. When households spend time shopping, they randomly draw prices over time: spending more time shopping will on

Figure 1: shopping time over the recent business cycle

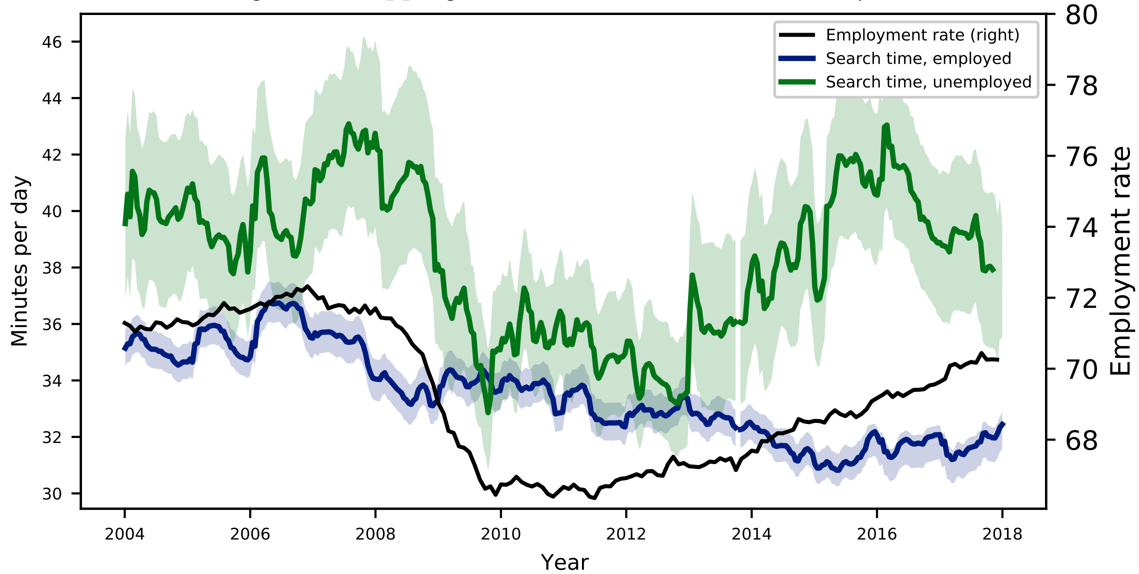

Shaded areas denote 95% confidence bands. Source: own computations based on the American Time Use Survey. Details in Appendix A. In the appendix, I show that this cyclicality is even stronger when controlling for demographics.

average lead to a larger basket of prices to compare. Households buy at the lowest observed price, implying that consumption weakly increases with time spent searching. Similar to the leisure-labor choice, the leisure-shopping decision is hence trading off leisure against consumption. On the other side of the market, firms are modeled as in Burdett and Judd (1983): they take as given the search intensities of the households and decide what prices to set. Setting a higher price leads to a higher revenue per customer, but a smaller set of customers to whom to sell. This trade-off can lead to a non-degenerate price distribution.

One advantage of this approach is that the environment is close to the original Burdett and Judd (1983) setup. As a result, many of the original theorems will still hold here, and the environment should be familiar to researchers. For example, the model always features an equilibrium without search, and the price distribution has a closed-form representation. As with similar setups, the model features zero, one, or two search equilibria with non-degenerate price distributions.

The household's search decision sets the model apart from previous search literature.

The model is a static counterpart to a dynamic model that could be used to analyze business cycles. Keeping the model static allows me to test the calibration that a dynamic model would also have to satisfy. Moreover, the only additional channel in a standard dynamic model would be the consumption-savings decision, which does not affect the search-leisure trade off that I am studying here.

With that in mind, I choose the King, Plosser, and Rebelo (1988) preferences that are consistent with balanced growth and widely used in business cycle analysis. First, I show that under these preferences, the household's search decision is independent of their income. This mirrors the well-known result that the household's labor-leisure decision under these preferences is also independent of their income – a key outcome that delivers balanced growth consistent with the Kaldor facts. The second result is that the search decision depends on the price distribution. The household's search intensity increases when the average price is lower, and when the dispersion of prices is higher. Households enjoy a higher price dispersion since they buy at the lowest price – which decreases when the variance of prices rises.

The model ties the search decision against the marginal utility of leisure, and so I attempt to validate the model by matching the search behavior of agents with different time endowments. For this exercise, I choose to match the search behavior of the employed and the unemployed. This is particularly interesting since differential search behavior by employment status has been found to be a potential amplifier of business cycle fluctuations (Kaplan and Menzio, 2016). I measure the search intensities and the time constraints in the American Time Use Survey and show that the unemployed spend 27% more time searching than the employed. I calibrate the price distribution to empirical moments. In particular, the model matches the empirical observation that the unemployed pay on average 2% less than the employed for a comparable consumption bundle (Kaplan and Menzio, 2016). I then test whether I can replicate the empirically observed differences in search behavior between the employed and the unemployed. The model qualitatively matches the empirical differences in search between the employed and the unemployed. However, the model has difficulties in quantitatively matching the levels of search intensity of the employed and the unemployed: in the model, the unemployed households spend between 100% and 200% more time searching than the employed - far more than the targeted 27%.

I discuss the different assumptions leading up to this result, and the degree to which relaxing them could reconcile the model with the data. One potent explanation is that unemployed and employed households are not sampling prices from the same distribution: firms may strategically set prices to discriminate between households with a high and low marginal evaluation of time. The unemployed spend 27% more time shopping than the employed and pay 2% less on the same consumption bundle. When assuming that the employed and the unemployed are drawing from the same price distribution, this can be used to discipline price dispersion and the returns to search. If firms can actually discriminate by employment status, such a calibration strategy would be misguided. Discrimination by employment status may be implemented with time-varying prices (Klenow and Malin, 2010; Kaplan and Menzio, 2015; Menzio and Trachter, 2018): firms might vary prices hourly or daily to prevent the employed – who on average are more constrained in their times of shopping – from finding the same prices as the unemployed.

Alternative extensions to match the data involve different preferences or technology assumptions. It is likely that the choice of microfoundation that one uses to fix this search gap will affect the behavior of the macroeconomic model that is built around it. For example, if stores are indeed able to discriminate by employment type, changes in the unemployment rate will not affect the price distributions, and an important amplification mechanism in Kaplan and Menzio (2016) becomes moot. I conclude that quantitatively matching the search intensities is a challenge. Going forward, more research on the source of this search gap is needed in order to microfound the consumer good search decision in macroeconomic models.

1 Model

I now build a model that microfounds households' shopping behavior in a way that paves the road for the analysis of business cycles. The model is supposed to satisfy several requirements. First, the model may be static, but should be easily extended to incorporate dynamics. Second, models that analyze business cycles are typically disciplined by targeting balanced growth. The model should be in line with that. Third, the model should be in line with previous work on consumer good search, such that well-known results from the literature will reappear in a similar fashion and researchers will be familiar with the environment.

The model that we study here is geared towards these requirements. It describes households and firms in a frictional goods market. The model is static, but permits closed-form solutions and could be extended to support a dynamic setup. The households have preferences over leisure and consumption that are standard in the business cycle literature and consistent with balanced growth. They have a fixed time and income endowment and want to spend all of their income on buying consumer goods. They take as given the distribution of prices that firms are charging in the economy and have to decide on how much of their time to spend on shopping. When they search for prices, they randomly meet with firms and draw prices from the distribution. Households that allocate more of their time to shopping will on average draw more prices. Households have perfect recall and will buy at the cheapest price that they find. Therefore, spending more time shopping leads on average to finding lower prices and consuming more. The firms are modeled exactly as in Burdett and Judd (1983): they take the search behavior of the households as given and set prices to maximize profits. Higher prices lead to a higher profit margin per unit sold, but fewer sales. This trade-off has the potential to lead to a non-degenerate price distribution: firms will charge different prices for the same good.

The fixed number of firms are modeled following Burdett and Judd (1983): they produce output with a constant marginal cost and set prices in order to maximize profits. Higher prices lead to a higher profit margin per unit sold, but fewer transactions. This trade-off has the potential to render firms indifferent in terms of charging various prices. Consequently, the firms may charge different prices in equilibrium.

In the following, I will first describe each side of the market in detail. Then, I study the various equilibria in section 1.3.

1.1 Households

The households have preferences over leisure ℓ and consumption ζ that can be represented by the utility function $u(\zeta, \ell)$. To consume, they need to spend their income endowment y on consumption units that they buy from firms. They have a total time endowment T, which they spend on shopping, t, and on leisure, ℓ. The search process will yield a final price p that they use to buy consumption goods. Households spend their entire income on consumption: $\zeta = y/p$. We denote by $U(p, t; y, T)$ the utility associated with spending t time for searching and finding the price p:

$$U(p, t; y, T) = u\left(\frac{y}{p}, T - t\right)$$

At the core of the model is the transformation between time spent searching and the resulting price draws. Following Burdett and Judd (1983), I assume non-sequential search. Here, households commit time to be spent on searching for prices and receive a Poisson draw of prices from the price distribution $F(p)$. For any time t, the number of prices drawn will be Poisson-distributed with mean $\lambda(t)$. I denote the probability of s draws given the arrival rate λ as $P(s, \lambda)$. We assume $\lambda(t) = aSt$, where S is the measure of firms operating in the economy and a is a search-efficiency parameter.[1]

In this class of model, we need to anchor the price distribution exogenously. Burdett and Judd (1983) assume that there is a reservation price above which households would never buy. I follow Kaplan and Menzio (2016) by specifying that households always have the option of transforming their income to consumption units at the reservation price r. Therefore, we can simplify the household's problem by assuming that all prices in the distribution $F(p)$ are below or equal to r. Households that drew multiple prices will purchase at the lowest price that they found. For a given number of price draws s, I denote the CDF and the PDF

[1] Appendix C shows that the search process can be microfounded by a model in which households search for firms that are uniformly distributed on a unit circle.

of the minimum-price distribution as $H(p;s)$ and $h(p;s)$. They can be characterized as

$$H(p;s) = Prob(\min \leq p) = 1 - \prod_{x=1}^{s}(1 - F(p))$$
$$h(p;s) \equiv \frac{\partial H(p;s)}{\partial p} = s(1 - F(p))^{s-1}F'(p).$$

The household's objective is to trade off leisure against consumption by choosing the amount of search that maximizes its objective function. The objective function is denoted as $K(t)$:

$$K(t;y,T) = P(0;\lambda(t))U(r,t;y,T) + \sum_{s=1}^{\infty} P(s;\lambda(t)) \int U(p,t;y,T)h(p,s,F)dp$$

Households receive zero price draws with probability $P(0;\lambda(t)))$. In that case, they transform at the reservation rate r. Otherwise, they receive $s > 0$ draws and purchase at the lowest price.

Preferences We now specify the functional form of the utility function $u(\zeta,\ell)$. The model is a precursor to a dynamic analysis of business cycles, and so we will follow that literature by assuming preferences that are consistent with balanced growth. In the United States, hours worked per capita have been stable over the past 80 years while real wages have more or less steadily increased (Prescott, 1986). In models that are used for the analysis of growth and business cycles, this stylized fact is matched by specifying preferences following King, Plosser, and Rebelo (1988). In these preferences, the income effect and the substitution effect of an increase in real wages exactly cancel out. We follow that literature by assuming

$$u(\zeta,\ell) = \frac{1}{1-\sigma}\left[\zeta^{\gamma} \cdot \ell^{1-\gamma}\right]^{1-\sigma}.$$

Under these preferences, the following lemma holds.

Lemma 1. *The objective function has the following compact representation:*

$$K(t;y,T) = e^{-\lambda(t)}\frac{1}{1-\sigma}\left[\left(\frac{y}{r}\right)^{\gamma}(T-t))^{1-\gamma}\right]^{1-\sigma} + \lambda(t)\int e^{-\lambda(t)F(p)}f(p)\frac{1}{1-\sigma}\left[\left(\frac{y}{p}\right)^{\gamma}(T-t)^{1-\gamma}\right]^{1-\sigma}dp \tag{1}$$

Proof. In the appendix. □

The household chooses t to maximize the objective function (2), and I denote the policy function as $g(y,t,F)$.

$$g(y,T,F) = \arg\max_{t} K(t;y,T,F) \tag{2}$$

How does the solution vary with the individual's endowments and the given distribution?

Proposition 1. *The objective function is homogeneous of degree one in income, and optimal search is independent of income:*

$$g_y(y,T,F(p,t')) = 0$$

Proof. Note that $K(t;y,T,F) = y^{\gamma(1-\sigma)}\cdot K(t;1,T,F)$. Income only scales the objective function and hence does not affect the optimal search intensity. □

Household income does not affect the search decision. This result relies heavily on the specified preferences which are chosen such that a change in wages would not affect hours worked in a standard business cycle framework. This necessarily implies that a change in income does not affect the search decision in this framework. To see that, note that the two problems are isomorphic: an increase in the wage rate proportionally affects the return to work, and affects the leisure-consumption trade-off. In this setup, an increase in income proportionally affects the return to search since the household spends its entire income on consumption. When consumption and leisure are log-additive in preferences, the income and substitution effect will therefore cancel out in both scenarios. This is not a technical point, but an economically meaningful result from the preceding microfoundation: if preferences are indeed fundamental in the sense that an individual's choices in different environments can be rationalizable with the same preference structure, then the income and substitution effects will operate in the shopping decision in the same way as they do in the work decision. This result can still hold when households enjoy shopping, as long as consumption, leisure, and shopping all continue to enter multiplicatively in the production function – for example as $\zeta^{\gamma}\ell^{\beta}t^{1-\gamma-\beta}$. Furthermore, this result holds also when including a labor choice. Consider for example an environment where households have a fixed wage rate w and their income is a result of their choice of working hours, $y = w \cdot h$. The time spent searching is independent of the wage rate. Formally,

Lemma 2. *Let*

$$\tilde{K}(t,\ell,h) = e^{-\lambda(t)}\frac{1}{1-\sigma}\left[\left(\frac{y}{r}\right)^{\gamma}\ell^{1-\gamma}\right]^{1-\sigma} + \lambda(t)\int\frac{1}{1-\sigma}\left[\left(\frac{y}{p}\right)^{\gamma}\ell^{1-\gamma}\right]^{1-\sigma}e^{-\lambda(t)F(p)}f(p)dp$$

and

$$\begin{aligned}(t^*(w),\ell^*(w),h^*(w)) = \arg\max_{t,\ell,h} K(t,\ell,h)\\ s.t.\\ y = wh\\ T \geq t+\ell+h\end{aligned}$$

The optimal search decision is independent of the wage rate:

$$\frac{\partial t^*(w)}{\partial w} = 0.$$

Proof. In the appendix. □

With respect to the second endowment, T, households' search decisions appear to be weakly increasing in the amount of time that they have available: households that are at a corner solution and spend no time searching might not respond to an increase in T, but those that are spending some time shopping will increase their search intensity when provided with more time.

Finally, we look at how $g(y, T, F)$ varies with the distribution F. It is difficult to make general statements: instead, we will analyze how individual moments of the price distribution affect the search decision. Instead of using the price distribution that results from the firms' optimal behavior, I will instead assume that prices are drawn from the truncated normal distribution with mean μ and variance σ. I do this for two reasons. First, it is more pedagogical to use a well-known distribution than the equilibrium price distribution that I have not introduced at this point. Second, the truncated normal distribution allows me to change both the mean and the variance independently. The equilibrium price distribution would instead respond to any parameter changes with a simulatenous change in the mean, the variance, and further higher-order moments.

Figure 2 displays how g varies with μ and σ. In the left-hand panel, the standard deviation is tiny and constant: all draws from the distribution will be very similar. Here, the main motivation for search is to have at least one draw from F, and the value of additional draws is negligible. As $\mu \to r$, the gains from search decrease, and optimal search decreases. Since search is random and costly, households already exert zero effort when μ is close – but not equal – to r.

In the right-hand panel, we fix $\mu = r$: the expected value of each price drawn is equal to the outside option. Now, the primary motivation from search comes from the dispersion of prices: the distribution F has a positive support for prices below r, and the household searches to find those. A higher variance also entails prices with a positive support above r, but since the household cares about the minimum price drawn and can always fall back to its outside option, a higher variance is always beneficial to it.

Due to the positive variance, there is a value in additional searches – and the higher the variance, the higher the value of searching more: the households want to search more when the variance is higher.

Notice that there is a discontinuity when varying σ, but not when varying μ. When the distribution is degenerate, a single draw is always sufficient, and effectively t is chosen to trade off the probability of 0 vs 1 draw. In the right-hand panel, the motivation for search is variance: the distribution warrants no searches at all when the variance is small. When the

Figure 2: Optimal search decreases in mean and increases in variance

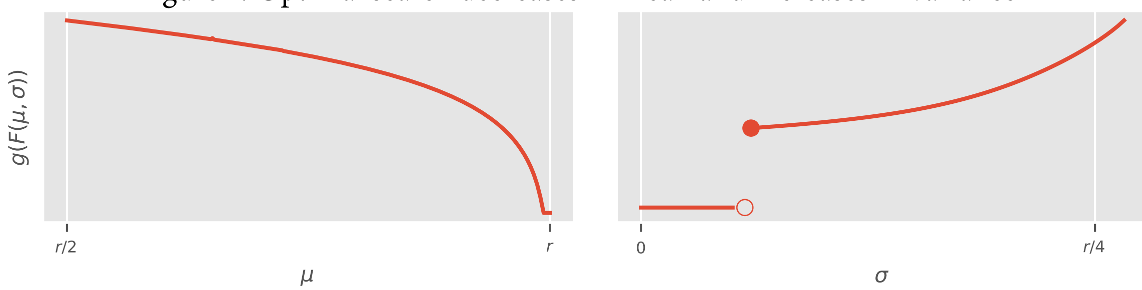

Optimal search under normally distributed prices with mean μ and standard deviation σ. The left-hand panel varies μ, holds $\sigma = 0.001$ constant. The right-hand panel fixes $\mu = r$, and varies σ.

variance is large enough for the household to search, it immediately wants multiple draws. Therefore, the household either chooses a search intensity consistent with zero price draws or one that is likely to lead to multiple price draws, thus generating a discontinuous search profile.

1.2 Firms

There is a fixed measure S of stores in the economy. They make profits by selling to a measure H of households. Households that draw multiple prices purchase at the store that offers them the lowest price. Out of the customers that arrive at a particular store, the store will only sell to those that have not drawn a lower price elsewhere. I will refer to these customers as "captured". For any price p, the probability of capturing a customer conditional on contact is denoted as $\eta(p;\lambda,F)$. Captured customers spend their total income on consumption and will hence buy y/p units of consumption. Stores produce the consumption good at unit cost c and hence make per-unit profits of $p-c$. Chaining these components allows us to compute the profits $\pi(p;\lambda,F)$ as in (3). The appendix shows that $\eta(p)$ permits the compact formulation as in (4).

$$\pi(p;F,\lambda) = \underbrace{\frac{H\lambda}{S}}_{\#\text{ contacts}} \cdot \overbrace{\eta(p;\lambda,S)}^{\text{Share captured}} \cdot \underbrace{\frac{y}{p}}_{\#\text{ goods solds}} \cdot \overbrace{(p-c)}^{\text{profit per sale}} \tag{3}$$

I denote the lowest and the highest price observed in $F(p)$ as $\underline{p}, \overline{p}$:

$$\underline{p} \equiv \min\{p : f(p) > 0\}$$
$$\overline{p} \equiv \max\{p : f(p) > 0\}$$

Lemma 3. *The capturing probability $\eta(p; \lambda, F)$ is given by*

$$\eta(p; \lambda, F) = \frac{e^{-\lambda}}{1 - e^{-\lambda}} \frac{1}{1 - F(p)} \left[e^{\lambda(1-F(p))} - 1 \right] \tag{4}$$

It satisfies

$$\eta(\underline{p}; \lambda, F) = 1$$
$$\eta(\overline{p}; \lambda, F) = \frac{\lambda e^{-\lambda}}{1 - e^{-\lambda}}$$

Proof. In the appendix. □

Naturally, the firms that offer the lowest price in the economy capture all customers that contact them. The firms that offer the highest price only capture the customers that have a single price draw. The probability of meeting such customers – conditional on capture – is given by $\lambda e^{-\lambda}/(1 - e^{-\lambda})$.

We can now state a definition of a partial equilibrium. Informally, a partial equilibrium is such that any observed price maximizes the profits.

Definition 1 (Partial firm equilibrium)**.** *For any $\lambda \geq 0$, a partial firm-side equilibrium is given by a density of prices $\{f(p; \lambda)\}$ such that*

$$\pi(p; F, \lambda) \geq \pi(p'; F, \lambda) \qquad \forall p : f(p; \lambda) > 0 \,, \forall p'$$

Lemma 4. *In any equilibrium with a strictly positive search, profits are strictly positive. The offer price distribution $F(p)$ is continuous and connected. It satisfies $c < \underline{p} \leq r$ and $\overline{p} = r$.*

Proof. In the appendix. □

Proposition 2 (Equilibrium price distribution)**.** *Under positive search, the unique offer price distribution consistent with these characteristics is given by* (5)*.*

$$F(p;\lambda) = \frac{1}{\lambda}(z + LambertW(-ze^{-z})) + 1 \tag{5}$$
$$z = \frac{r}{p}\frac{p-c}{r-c}$$

The lower bound of the distribution satisfies

$$\underline{p}(\lambda) = \frac{rc}{r - \lambda\frac{e^{-\lambda}}{1-e^{-\lambda}}(r-c)}$$

Proof. In the appendix. □

Next, we discuss how the offer price distribution responds to changes in search intensity.

Proposition 3 (Offer price distribution and search intensity)**.** *Distributions consistent with a lower λ first order stochastically dominate those with a higher λ:*

$$F(p;\lambda) \geq F(p;\lambda') \qquad \forall\lambda > \lambda', \forall p$$
$$F(p;\lambda) > F(p;\lambda') \qquad \forall\lambda > \lambda', \forall p \in [\underline{p}(\lambda'), r)$$
$$\underline{p}'(\lambda) = \frac{e^{\lambda}(1-\lambda)-1}{(e^{\lambda}-1)^2}\frac{r-c}{rc}\underline{p}(\lambda)^2 < 0$$

Proof. In the Appendix. □

Figure 3 draws the offer price distribution for varying search intensities to demonstrate this point. A corollary of Proposition 3 is that distributions consistent with a lower λ have higher mean prices.

The relationship between λ and the variance of the distribution is ambiguous. When λ is low, the distribution has most of its mass close to r: an increase in λ increases the dispersion. As the distribution spreads out, the lower bound of its support converges to c. An increase in λ leads to even more prices being close to c: when λ is high, an increase in λ leads to a

Figure 3: Offer price distribution and varying λ

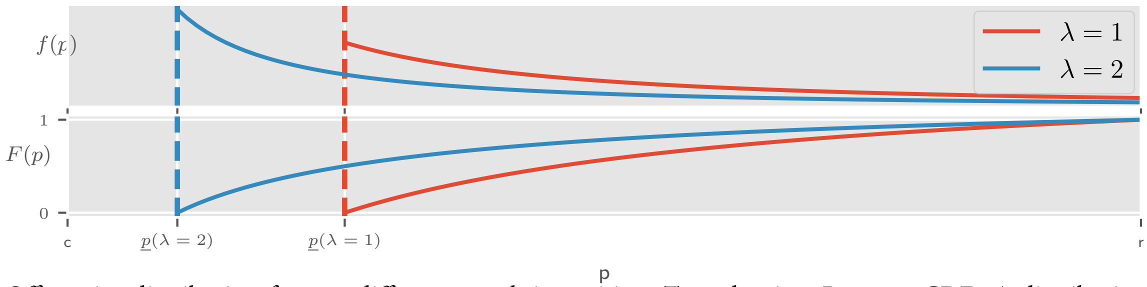

Offer price distribution for two different search intensities. Top: density. Bottom: CDF. A distribution consistent with a lower λ first order dominates the higher-λ distribution and has a smaller support.

Figure 4: Mean and variance under varying λ

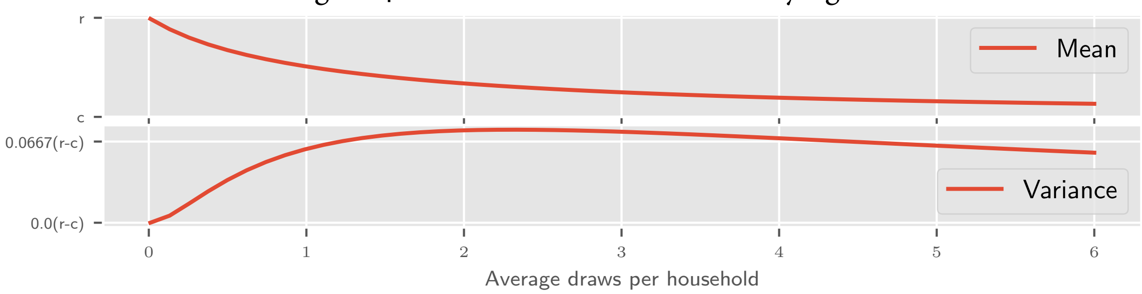

concentration of prices around $\underline{p}(\lambda)$ and a decrease in dispersion. Figure 4 demonstrates this by plotting the mean and variance of $F(p)$ against λ.

1.3 Equilibrium

We can now define an equilibrium for this economy. We will not consider firm entry and treat the measure of firms S as an exogenous parameter.

Definition 2 (Equilibrium)**.** *Given measures $\{S, H\}$ and endowments $\{y, T\}$, an equilibrium is described by a tuple $\{F, t\}$ such that*

1. *$t = g(y, T, F)$ is optimal given F* (2)
2. *F is consistent with t* (5)

Figure 5 displays the best response of an individual household to the aggregate search behavior of all other households. Aggregate search behavior affects the offer price distribution and thereby the individual household, which summarizes the fixed-point problem that characterizes the equilibrium.

Many distributions F are consistent with zero search, for example a degenerate distribution with support only on r. This distribution would induce zero search and demonstrates that an equilibrium without search always exists.

When aggregate search t' is very low, most of the mass of the offer price distribution will still be near r, inducing a very little search: both the mean and the variance of F are such that search is not optimal. Following Proposition 3, a higher λ leads to distributions that are preferred by the agent. This induces additional search to draw from that distribution: $dg/dt > 0$. As t increases, the dispersion in F eventually starts decreasing, such that additional searches do not improve much on a first price draw. Since the variance motivation of search decreases, households reduce their search: $dg/dt < 0$.

Figure 5: Equilibrium as a fixed-point problem

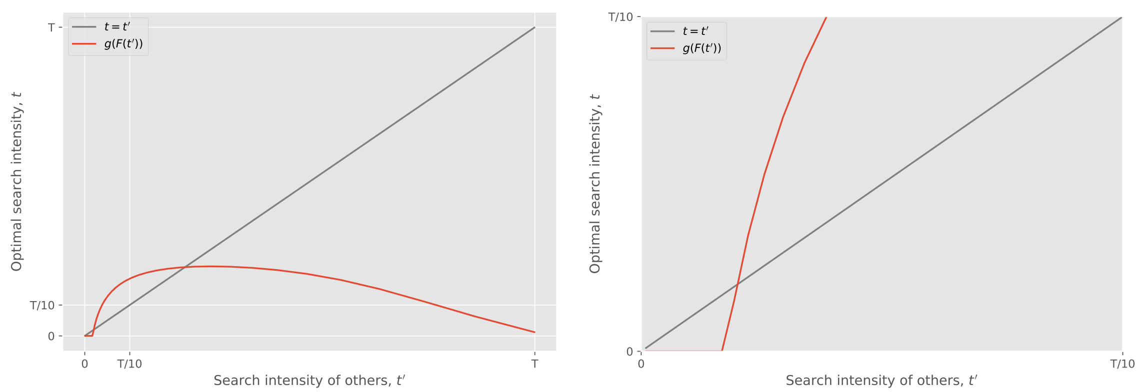

On the x-axis, we vary the aggregate search intensity which affects the offer price distribution. Against that, we plot the search intensity which is optimal given the offer price distribution. The right-hand panel zooms into the lower-left quadrant of the left-hand panel.

In conclusion, $g(F(t'))$ always first increases and then decreases. What does this mean for the number of potential search equilibria?

This particular example in Figure 5 features two search equilibria. Figure 6 changes the search efficiency parameter a which linearly scales $\lambda'(t)$. When search is very inefficient, no search equilibrium exists. An increase in the search efficiency shifts $g(F)$ upward, and eventually leads to the two familiar equilibria. An intermediate value of a exists such that $g(F)$ would be tangential to $t' = t$, implying that there was only a single equilibrium.

Which of these equilibria are more likely to be observed in the real world? To answer this question and select an equilibrium to which to calibrate, I focus on a particular type of trembling-hand mistake where *all* agents tremble at the same time. I call an equilibrium stable if a sequence of best responses to any tremble in the neighborhood around that equilibrium converges to the equilibrium.

Definition 3 (Stable equilibrium)**.** *A stable equilibrium $\{t, F\}$ is one where the sequence of best responses to any $\tilde{t}$ in a neighborhood around t converges to $\{t, F\}$. Denote $t^i = g(F(\lambda(t^{i-1}))$. Then, a stable equilibrium $\{t, F\}$ satisfies*

$$\lim_{x \to \infty} \tilde{t}^x \to t \quad \forall \tilde{t} \in (t - \epsilon, t + \epsilon) \quad , \epsilon > 0$$

From inspecting Figure 6, it is clear that in the two equilibria scenarios, only the latter is stable. When there is a single equilibrium, it is stable.

1.4 Normalizations

To inform the calibration, it is useful to analyze the impact of c and r on the offer price distribution. Increasing either of these will tilt the distribution to the right. However, a proportional scaling of c and r shifts and scales the offer price distribution proportionally, as claimed by Lemma 5.

Lemma 5. *A proportional increase in both r and c by a scaling factor $\psi > 0$ proportionally scales $F(p)$.*

$$\underline{p}(\lambda, \psi r, \psi r) = \psi \underline{p}(\lambda, r, c)$$
$$F(\psi p; \lambda, \psi r, \psi c) = F(p; \lambda, r, c)$$

Proof. In the appendix. □

Lemma 6. *Household income y does not affect the outcomes.*

Figure 6: Visualization of equilibria

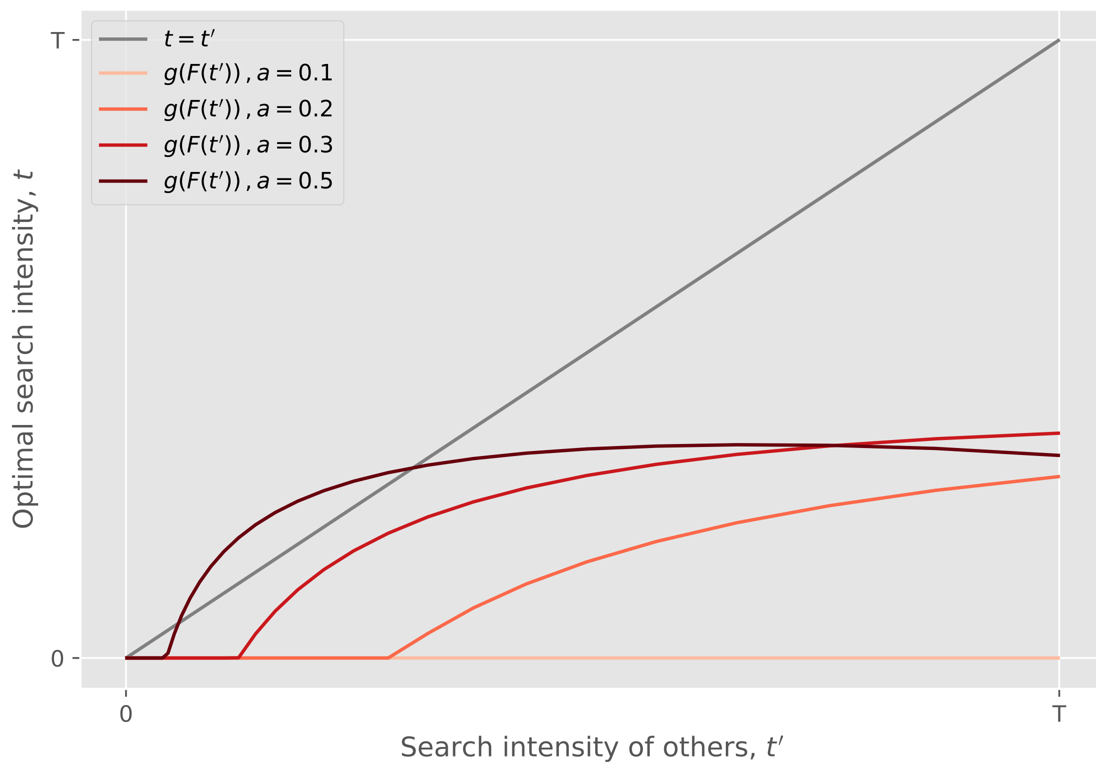

On the x-axis, we vary the aggregate search intensity which affects the offer price distribution. Against that, we plot the search intensity which is optimal given the offer price distribution.

Proof. The two endogenous outcomes in the economy are t and F. y does not affect income, as shown in Proposition 1. Moreover, y does not appear in the expression for F. $\square$

Lemma 7. *For any* $\{H, S, a\} \exists a'$ *such that the equilibrium outcomes under* $\{H, S, a\}$ *and* $\{1, 1, a'\}$ *are identical.*

Proof. H and S do not directly affect either F or t. The affected variable is $\lambda = atS$, and rescaling $a' = aS$ will allow us to normalize $S = 1$ and keep λ at its previous level. $\square$

2 Good search by employment status

Next, we validate the model by matching essential properties of the consumer good search. In the model, households with a ceteris paribus higher time endowment decide to spend more time shopping. Therefore, I attempt to match the search behavior of the employed and the unemployed. The differential search behavior by employment status is an interesting moment to match for two reasons. First, as I will show empirically, there is a large difference between the search behavior of the employed and the unemployed. Second, these differences are economically meaningful: Kaplan and Menzio (2016) show that this differential search behavior by employment status can significantly amplify business cycle fluctuations.

Therefore, I will attempt to match these differences in the model. We will find that the calibrated model cannot match the data: it vastly overestimates how many additional hours the unemployed want to search, compared to the employed.

Before detailing the measurement of the search data and the calibration procedure, I need to extend the model to allow for both employed and unemployed households.

2.1 Model with employment status

Most of the model is very similar to the earlier homogeneous household framework. I will keep it in the static partial equilibrium and fix the unemployment rate at u. As argued before, I can normalize $H = 1$ and $S = 1$. The employed households have wage income w, while the unemployed worker's income is denoted as b. They have different time endowments available that I denote T^e and T^u. Following Kaplan and Menzio (2016), firms are not able to discriminate between their employed and unemployed customers: both the employed and the unemployed draw from the same offer price distribution F. For any given offer price distribution F, I denote the optimal search choice of the employed and the unemployed as $g(w, T^e, F)$ and $g(b, T^u, F)$.

The main difference as compared to the previous framework is that firms now have to consider the two different types of customers when setting their prices. Conditional on contact, the probability of meeting an agent of type i is denoted as ξ^i and is a function of both arrival rates and the relative shares. These probabilities naturally satisfy $\xi^e + \xi^u = 1$. The probability of capturing η is now type-dependent and denoted as η^i.

$$\xi^e = \frac{(1-u)\lambda^u}{(1-u)\lambda^e + u\lambda^e}$$
$$\xi^u = \frac{u\lambda^u}{(1-u)\lambda^e + u\lambda^e}$$
$$\eta^i = \frac{e^{-\lambda^i}}{1-e^{-\lambda^i}} \frac{1}{1-F(p)} \left[e^{\lambda^i(1-F(p))} - 1 \right]$$

$$\pi(p;\lambda^e,\lambda^u,u,F) = \lambda \cdot \left[\xi^e(\lambda^e,\lambda^u,u)\eta^e(p,F,\lambda^e)\frac{w}{p} + \xi^u(\lambda^e,\lambda^u,u)\eta^u(p,F,\lambda^u)\frac{b}{p} \right](p-c)$$

An equilibrium is now characterized by $\{t^e, t^u, F\}$, where $t^e = g(y, T^e, F)$, $t^u = g(b, T^u, F)$ and F has positive support for any price that maximizes $\pi(p, \lambda^e, \lambda^u)$. F can no longer be expressed in closed form. However, a closed-form solution does exist for $\underline{p}$, the lower bound of the support for the offer price distribution.

$$\underline{p} = \frac{\left[\xi^e w + (1-\xi^e)b \right] cr}{\left[\xi^e w + (1-\xi^e)b \right] r - \left[\xi^e \eta^e(r,F,\lambda^e) w + (1-\xi^e)\eta^u(r,F,\lambda^u) b \right](r-c)} \tag{6}$$

As in the simple model, the absolute levels of income do not affect F. Here, this implies a proportional scaling of $\{w, b\}$. Proportional increases in $\{c, r\}$ linearly scale $\underline{p}$ and F.

2.2 Measurement

I use the American Time Use Survey (ATUS) to measure the search intensity of the employed and the unemployed. Each individual that is interviewed for the ATUS provides a detailed record of all activities for a particular random day. I weight each individual by her ATUS record weight, and use all years between 2003 and 2017. I aggregate the reported activities into major groups and exclude some ambiguous activities that amount to a total of

Table 1: Time use by employment status

	Employed	Unemployed
Leisure	5.147	7.744
Shopping	0.544	0.642
Personal care	10.439	11.327
Home production	1.405	2.267
Work	5.495	0.862
Education	0.154	0.381

Source: American Time Use Survey. Measured in hours per day. Unassigned time: ca 40 minutes.

40 minutes per day on average. The resulting aggregated time use categories are summarized in Table 1.

Personal care includes sleep, and it is the largest category for both types. The model allows individuals to distribute time-at-hand into either leisure or search, and cannot speak to other margins of time use: I calibrate total time at hand T to the sum of leisure and shopping.

2.3 Calibration

I want to test whether the empirically observed search choices of employed and unemployed households are one equilibrium outcome of the model. I test whether these time allocations can be a fixed point by employing the following calibration strategy. I fix a number of preference and technology parameters. Importantly, I also fix t^e and t^u, the optimal search intensity of employed and unemployed households, to their empirical counterparts. Given these values, I calibrate the offer price distribution to match the empirical counterparts. Then, I test whether – given the calibrated distribution – I can recover t^e and t^u as solutions to the households' problem.

I assume the period length to be one week. Table 2 lists the chosen parameters. Time at hand T^i is calibrated to the sum of the household's leisure and shopping time, as sourced in Table 1. For unemployed and employed households, this amounts to 8.38 and 5.68 hours, respectively. The model is isomorphic in the absolute value of time endowments. Therefore,

Table 2: Selected parameters

Moment	Value	Description
T^e	0.679	Time endowment (employed)
T^u	1.000	Time endowment (unemployed)
σ	0.500	Curvature of utility
b	0.850	Expenditure of unemployed
t^e	0.022	Search (employed)
t^u	0.028	Search (unemployed)
u	0.050	Unemployment rate

Sources detailed in text.

I normalize $T^u = 1$ and set $T^e = 5.68/8.38 = 0.68$. I set the risk aversion parameter σ to 0.5 but will conduct a robustness check later. As argued before, offer price distributions are invariant to a proportional scaling of b and w. Moreover, search decisions are independent of incomes. Therefore, I normalize $w = 1$, and follow Kaplan and Menzio (2016) by setting $b = 0.85$ to match the relative expenditures of the unemployed and the employed. I fix the share of unemployed households at 0.05.

Two parameters that are related to the offer price distribution are calibrated to match moments in the data. First, a governs the translation of time spent searching into average price draws. For any fixed search intensities t^u, t^e, we can choose the difference in the average number of draws by selecting a appropriately. A direct implication of the difference in the average number of draws is the expected difference in average prices: a higher a will lead to a larger difference in the average expected prices between the employed and the unemployed. Following Kaplan and Menzio (2016), I calibrate a to match the fact that the unemployed spend on average 2% less on a comparable consumption basket. Second, the households' outside-option price r is calibrated to match the max-to-min ratio of the empirical offer price distribution. From (6), it is clear that $\underline{p}$ responds less than one-for-one to a change in r. Therefore, one can target $r/\underline{p}$ through the calibration of r. I follow Kaplan and Menzio (2016) by targeting a max-to-min ratio of 1.7. Table 3 displays the implied values for r and a.

Next, I want to test whether the model can produce the so-far fixed search choices t^e and t^u as optimal choices under the calibrated offer price distribution. The last free parameter

Table 3: Calibrated parameters

Moment	Target	Value	Parameter
$r/\underline{p}$	1.70	1.700	r
$E[p^u]/E[p^e]$	0.98	0.980	a

Table 4: Endogenous search intensities

Variable	Target	Value	Description
t^e	0.022	0.022	Search intensity, employed
t^u	0.028	0.052	Search intensity, unemployed

is γ. It governs the relative importance of leisure in preferences. I calibrate γ to match t^e and test how close the implied t^u is from its empirical counterpart. Table 4 documents the result: γ manages to pinpoint t^e exactly at its target, but t^u is twice as large as its empirical counterpart.

2.4 Mechanism

Why do the unemployed spend more time shopping in the model than in the data? Households that have a higher time endowment want to spend it on all available margins – leisure and search. Qualitatively the result makes sense: households with more time available want to spend more time on search. Quantitatively, the large extent to which an unemployed individual's time is devoted to search is not in line with the data. The reason is that – in the model – the gains from additional search are relatively high. Figure 7 displays the distribution of the minimum price of F for different search intensities t^i. It is clear that the effective price distributions for the targeted employed and unemployed households look very similar. In particular, both have a high rate of making zero draws leading to the high minimum price r. However, the t^u that is implied by optimal choice leads to a distribution that has a much larger mass at the lower end of the distribution, and a much lower weight on the maximum price. The gains associated with additional search appear high.

Why does the model predict large differences in optimal search by employment status?

Figure 7: Minimum-price distributions by search intensity

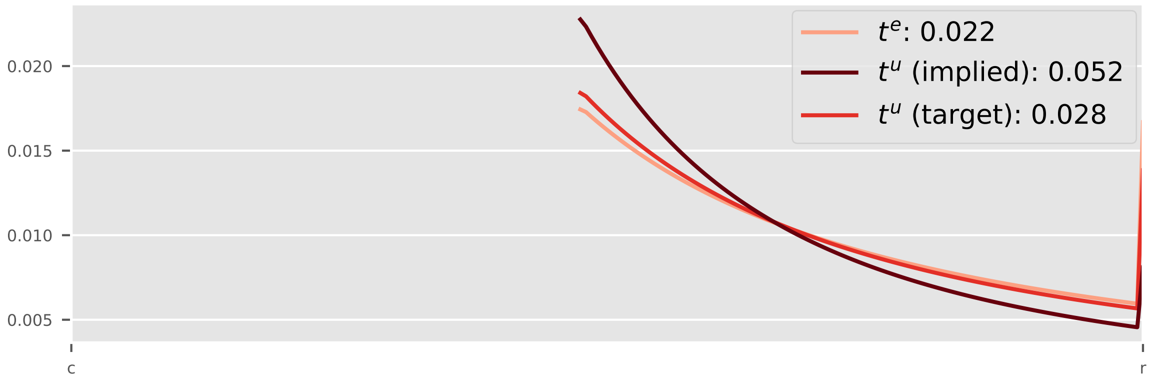

Density of the minimum-price distribution under the (fixed) equilibrium offer price distribution. Returns to search are partly driven by a reduction in the probability of paying r.

Figure 8: Variation of marginal cost and gains with T

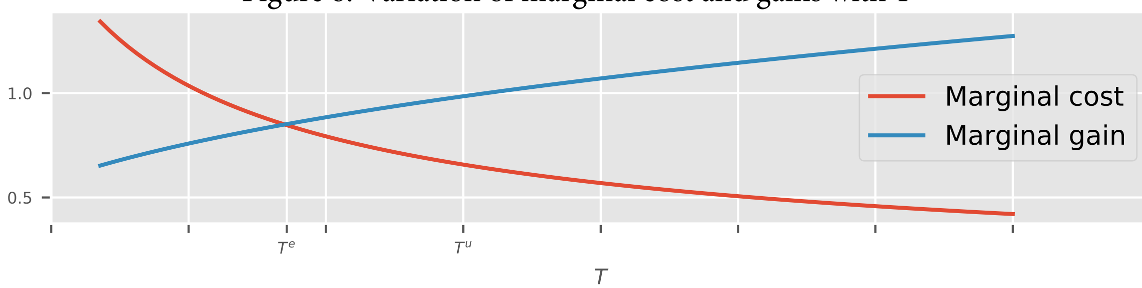

The employed and the unemployed both differ in time and income endowment. We know that in the model, the household's choice does not vary with income: the variation is purely caused by the time endowment. To analyze the relevance of the time endowment, we can decompose the objective function into the product of a leisure component and a consumption component – I refer to the latter as $A(t, y)$.

$$K(t, T, y) = \frac{(T-t)^{(1-\gamma)(1-\sigma)}}{1-\sigma} A(t, y)$$

$$A(t, y) \equiv e^{-\lambda(t)} \left(\frac{y}{r}\right)^{\gamma(1-\sigma)} + \lambda(t) \int \left(\frac{y}{p}\right)^{\gamma(1-\sigma)} f(p) e^{-\lambda(t)F(p)} dp$$

An interior solution requires that $K_t = 0$: the marginal cost in terms of leisure is equal to the marginal gains in terms of consumption. Equation (7) computes this derivative. The first term denotes the marginal cost associated with searching more, and the second term the corresponding consumption gain. Figure 8 displays these marginal costs and the marginal gains as we vary the time endowment T and keep the solution $t = t^e$ fixed. An increase in the time endowment naturally decreases the marginal cost of search. That the marginal gains vary with T is more surprising since T does not directly appear in $A_t(t, y)$. For a fixed t, a larger T does increase the utility derived from leisure, and it complements the gains from consumption. Both the decreased costs and the increased gains lead to the choice of a high t^u as implied by the calibration.

$$K_t(t, T, y) = -\frac{(1-\gamma)(1-\sigma)}{T-t} K(t, T, y) + \frac{(T-t)^{(1-\gamma)(1-\sigma)}}{1-\sigma} A_t(t, y) \tag{7}$$

$$A_t(t, y) = \left[-e^{\lambda(t)} \left(\frac{y}{r}\right)^{\gamma(1-\sigma)} + \int \left(\frac{y}{p}\right)^{\gamma(1-\sigma)} f(p) e^{-\lambda(t)F(p)} [1 - \lambda(t)F(p)] dp \right] \lambda'(t)$$

The only arbitrarily set parameter in the calibration strategy was the degree of risk aversion σ. We want to ensure that the riskiness of receiving zero draws together with the chosen degree of risk aversion is not the main driver behind the results. Therefore, I redo the calibration for a range of values for risk aversion. Figure 9 displays the results of this exercise. The top panel shows that the calibrated γ slightly increases in σ almost everywhere. The discontinuity of preferences at $\sigma = 1$ is also visible in the calibrated γ. To provide another testable prediction of the model, I compute the Frisch elasticity for each calibrated γ-σ combination. For very low degrees of risk aversion, the implied Frisch elasticity is high. For the more reasonable values of σ, the Frisch elasticity is around 1, in between its typical micro estimations and macro calibrations. The last panel displays the implied ratio of t^u/t^e. The model generates ratios around 2.2 for all σ values – far off the empirical ratio of 127%.

Figure 9: Optimal allocation of leisure with varying σ

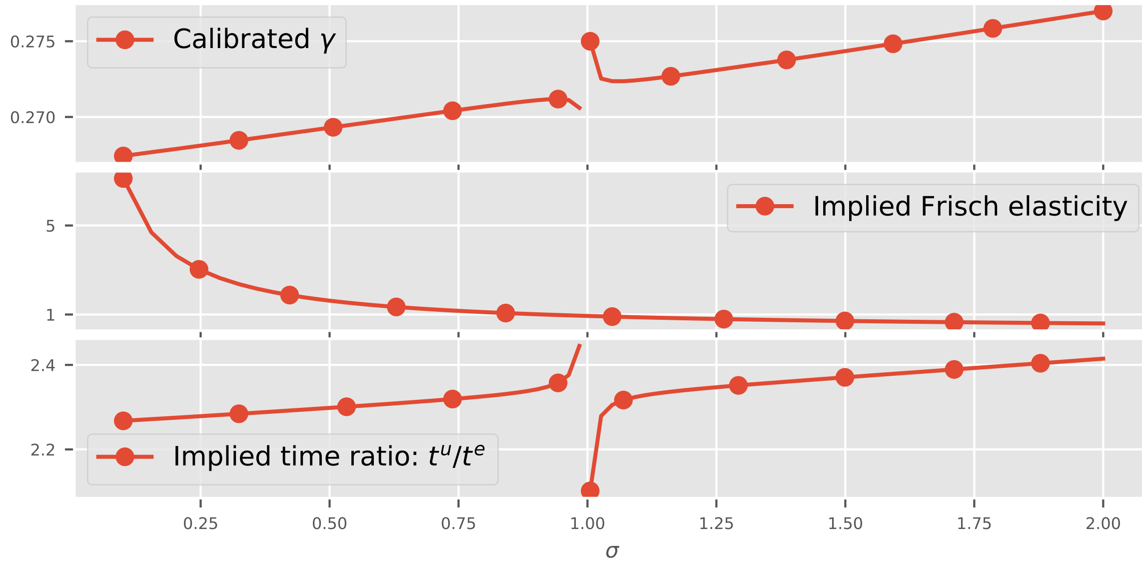

3 Discussion

The attempt to validate the model by targeting different search intensities by employment status fails. In the data, the unemployed spend 27% more time shopping than the employed, and on average spend 2% less on a similar consumption basket. When targeting the implied price distribution, the model generates ratios of search intensities that are around a factor of 2. Under reasonable calibrations and independently of the chosen degree of risk aversion, the model generates search intensities of the unemployed that are far beyond those measured empirically. Which key assumption(s) of the model are causing the disconnect?

First, households have perfect information about the distribution of prices, but no information about the actual prices at any given store. To the extent that prices are not completely unpredictable on a weekly basis, households could use information from previous periods to reduce the required search intensity – a feature missing in this static framework. However, it is unclear why the introduction of additional information would reduce the search gap between the unemployed and the employed.

3.1 Heterogeneous goods

In this model, households search to find low prices for a single representative good. In the data, unemployed households spend 15% less on nondurable consumption than the employed. If the model featured multiple varieties of consumption goods, and search was required for each of these, the observed small search gap could potentially be rationalized. For example, suppose households have non-homothetic preferences over two goods, "food" and "other". Unemployed households purchase food only, while employed households use their higher disposable income to purchase both types of goods. To phrase the solution in terms of Figure 8, if both types choose the same search intensity, the marginal cost of time will still be lower for unemployed households. However, their marginal gains from additional search will also be lower, since they only spend that additional search on a single good. Employed households spend their time on two goods: when they search as much as the unemployed, they effectively draw fewer prices for each good. The difference in total varieties purchased across the employed and the unemployed could be used to discipline the consumption good aggregator in the preferences.

3.2 Time does not equal search

A second approach involves the fact that price draws are probably not linear in the time spent searching. In the microfoundation for the Poisson draws, we assumed that households spend t traveling at constant speed on a unit circle, and contact stores randomly. This implied that the average number of prices drawn increases linearly in the time spent searching $\lambda(t) = atS$. Alternatively, the search process might involve a fixed sunk cost $\underline{t}$:

$$\lambda(t) = a(t - \underline{t})S \tag{8}$$

This additional technology parameter would indeed allow the model to fit any search gap $t^u - t^e$. One microfoundation involves the fact that stores are not randomly spread on a unit circle. Instead, several stores are located in the vicinity of a parking space. Households have to first spend $\underline{t}$ to reach that parking space, but can then access many stores at once.

3.3 Discrimination by employment status

The model assumes that the stores cannot discriminate prices by employment status: both the employed and the unemployed are drawing from the same distribution F. There is some suggestive evidence that firms are indeed able to discriminate. For example, employed households are typically constrained by the times of the day at which they can search for prices. Consistent with that type of discrimination, Kaplan and Menzio (2015) show that some prices vary within the same good and store over time. If such a mechanism was true, the employed and the unemployed would be searching from different distributions. This could explain why the unemployed only spend 27% more time shopping: the distribution of prices is very compressed both for the employed and the unemployed, which reduces the incentives for additional price draws.

In this model, even if F^e and F^u were both calibrated to the same technology parameters r and a, they would look different. To see this, assume by contradiction that both distributions were identical. In that case, the unemployed would search more, since they have the same marginal gains, but a smaller marginal cost of searching. By Proposition 3, F^e would then stochastically dominate F^u, which is a contradiction. Note that F is independent of income y, and different expenditures by the employed and the unemployed would not play any role here.

The question left to answer is whether the differences between F^u and F^e are such that they reduce the gap in search intensity between the two types. Recall that a household's search intensity decreases in the expected price of the distribution, and increases in its dispersion. F^e first-order stochastically dominates F^u, and so the difference in average prices would even increase the search gap. The dispersion could potentially offset this: if F^u has a smaller dispersion than F^e, the returns to additional price draws are smaller for the unemployed, which might overall shrink the search gap between the employed and the unemployed.

4 Conclusion

In this paper, I provide a microfoundation for the household's search decision that can be readily integrated into macroeconomic analysis. However, I caution against doing that because the model – taken at face value – does not match well the differential search behavior by employment status. More precisely, the model predicts a much larger ratio of search time between the unemployed and the employed than what is observed in the data.

I discuss several potential mechanisms that can reduce this search gap. I argue that a model could make sense of the data if searching households have to sink a fixed cost of time prior to receiving price draws. A second approach would be to incorporate non-homothetic preferences. Finally, I argue that a model that allows firms to discriminate by employment status could potentially rationalize the empirical findings.

Given the limited empirical data available, it is difficult to ascertain which of these mechanisms are at play. However, different implementations of the search environment in a macroeconomic model will likely lead to different aggregate behavior of the model.

For example, a model that follows the fixed-cost approach would understand the rise in internet shopping as a decrease in the fixed-cost component and predict that the search gap has increased in recent years – not entirely in line with empirical observations. Also, Kaplan and Menzio (2016) emphasize a business cycle mechanism where the search behavior of the unemployed affects the revenue that the firms receive from employed shoppers. If firms are indeed able to discriminate by employment status, that mechanism is moot and the model's dynamics become very similar to those of a simpler version that does not include consumer good search (Pissarides, 1979). Therefore, it is important to understand which of these mechanisms is actually the most critical for bringing the model's microfoundation closer to the observed search behavior. So far, our understanding of these mechanisms is limited: little is known about stores' ability to discriminate across their customers, or how additional time spent searching transforms to lower effective prices. This paper emphasizes the value of additional empirical work on that front to ensure that the resulting models are less susceptible to the Lucas (1976) critique.

References

Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis (2013). "Time use during the great recession". In: *The American Economic Review* 103.5, pp. 1664–1696. ISSN: 0002-8282.

Alessandria, George (Aug. 2009). "Consumer search, price dispersion, and international relative price fluctuations". In: *International Economic Review* 50.3, pp. 803–829. ISSN: 00206598.

Bai, Yan, Jose-Victor Rios-Rull, and Kjetil Storesletten (2011). "Demand shocks as productivity shocks". In: *Federal Reserve Board of Minneapolis*.

Burdett, Kenneth and Kenneth L. Judd (July 1983). "Equilibrium price dispersion". In: *Econometrica* 51.4, p. 955. ISSN: 00129682.

Head, Allen, Lucy Qian Liu, Guido Menzio, and Randall Wright (Oct. 2012). "Sticky prices: a new monetarist approach". In: *Journal of the European Economic Association* 10.5, pp. 939–973. ISSN: 15424766.

Kaplan, Greg and Guido Menzio (Nov. 2015). "The morphology of price dispersion". In: *International Economic Review* 56.4, pp. 1165–1206. ISSN: 00206598.

— (June 2016). "Shopping externalities and self-fulfilling unemployment fluctuations". In: *Journal of Political Economy* 124.3, pp. 771–825. ISSN: 0022-3808.

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo (1988). "Production, growth and business cycles: i. the basic neoclassical model". In: *Journal of Monetary Economics* 21.2–3, pp. 195–232. ISSN: 0304-3932.

Klenow, Peter J. and Benjamin A. Malin (2010). "Microeconomic evidence on price-setting". In: *Handbook of Monetary Economics*. Vol. 3. C, pp. 231–284.

Lucas, Robert E. (Jan. 1976). "Econometric policy evaluation: a critique". In: *Carnegie-Rochester Conference Series on Public Policy* 1.5, pp. 19–46. ISSN: 01672231.

Menzio, Guido and Nicholas Trachter (2018). "Equilibrium price dispersion across and within stores ". In: *Review of Economic Dynamics* 28, pp. 205–220. ISSN: 1094-2025.

Petrosky-Nadeau, Nicolas, Etienne Wasmer, and Shutian Zeng (2014). "Shopping time". In:

Pissarides, C. A. (Dec. 1979). "Job matchings with state employment agencies and random search". In: *The Economic Journal* 89.356, p. 818. ISSN: 00130133.

Prescott, Edward C (1986). “Theory ahead of business cycle measurement”. In: *Quarterly Review* 10.4, pp. 9–23.

Qiu, Zhesheng and José-Víctor Ríos-Rull (2019). “Directed search, nominal rigidities and markup cyclicality”.

A Shopping time over the business cycle

Individuals chosen for participation in the American Time Use Survey (ATUS) are randomly drawn from the Current Population Survey (CPS). They are then each asked to provide detailed diaries about their activities during the previous day. As individuals have a tendency to lump together shopping on few days of the week, this means that many individuals will report zero minutes of shopping. Additionally, the unemployment rate is low throughout parts of the sample period. Finally, shopping time has a lot of cyclical variation within weeks, months, and seasons. Therefore, it is a challenge to document the shopping time of the unemployed precisely and simultaneously capture changes of that shopping time over the business cycle.

I start by computing an aggregate shopping index by adding up the following activity codes:

t070201	Comparison shopping
t070299	Researching purchases, n.e.c.*
t070301	Security procedures rel. to consumer purchases
t070399	Security procedures rel. to consumer purchases, n.e.c.*
t180701	Travel related to grocery shopping
t180782	Travel related to shopping (except grocery shopping)
t070101	Grocery shopping
t070104	Shopping, except groceries, food and gas
t070105	Waiting associated with shopping
t070199	Shopping, n.e.c.*

This list entails most shopping activities in the ATUS. I exclude purchasing gas and food (not groceries) since the demand among the employed and the unemployed for these may make the comparison very difficult.

I then compute a running mean of the aggregated shopping time for both the employed and the unemployed with a centered uniform window length of two years. Each observation is weighted by TUFNWGTP, the weighted variable provided by the ATUS. I exclude three outliers in the year 2018 where the weighted shopping time is 30% higher than in the next

Figure A.1: residualized shopping time not constant over the recent business cycle

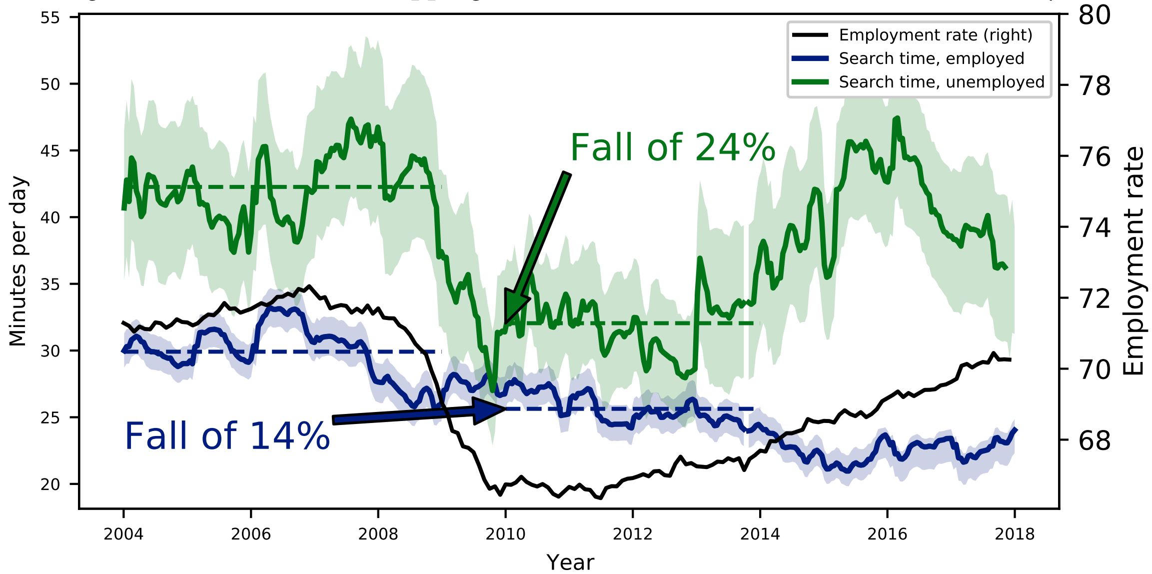

Shaded areas are 95% confidence bands.

one. When included, these weighted outliers are strong enough to significantly change the shopping time mean in the years after 2015, since there are few unemployed sampled in that period and the outliers also affect the nearby periods through the running mean. Finally, I compute 95% confidence bands by sampling the dataset separately by employment status. This results in Figure 1.

It remains to be answered whether the change in average shopping time by employment status actually reflects a change in the shopping decision or is simply driven by a change in the composition of the pool of the unemployed. To answer this, I residualize the shopping time variable by adding fixed effects for income, sex, age, and income. I add the mean of shopping time by employment status to the resulting residualized variable and proceed with the same running mean as before. This results in Figure A.1.

B Proofs

B.1 Proof of Lemma 1

After inserting the Poisson probabilities, we get

$$
\begin{aligned}
K(t; y, T, F) &= e^{-\lambda(t)} U(r, t; y, T) + \sum_{s=1}^{\infty} e^{-\lambda(t)} \frac{\lambda^s}{s!} \int U(p, t; y, T) s(1 - F(p))^{s-1} f(p) dp \\
&= e^{-\lambda(t)} U(r, t; y, T) + e^{-\lambda(t)} \int U(p, t; y, T) \lambda(t) \sum_{s=1}^{\infty} \frac{\lambda(t)^{s-1}}{(s-1)!} (1 - F(p))^{s-1} f dp \\
&= e^{-\lambda(t)} U(r, t; y, T) + e^{-\lambda(t)} \int U(r, t; y, T) \lambda(t) e^{\lambda(t)(1-F(p))} f(p) dp \\
&= e^{-\lambda(t)} U(r, t; y, T) + \int U(r, t; y, T) \lambda(t) e^{-\lambda(t)F(p)} f(p) dp
\end{aligned}
$$

where the third line uses one definition of the exponential. Replacing again $U(r, t; y, T)$ yields the expression in the Lemma.

B.2 Proof of Lemma 2

Denote by w a fixed wage rate and by h the amount of hours worked. Then, the adjusted objective function can be written as previously

$$
\tilde{K}(t, \ell, h) = e^{-\lambda(t)} \frac{1}{1-\sigma} \left[\left(\frac{wh}{r} \right)^{\gamma} \ell^{1-\gamma} \right]^{1-\sigma} + \lambda(t) \int \frac{1}{1-\sigma} \left[\left(\frac{wh}{p} \right)^{\gamma} \ell^{1-\gamma} \right]^{1-\sigma} e^{-\lambda(t)F(p)} f(p) dp
$$

It will be instructive to refactor this as

$$
\begin{aligned}
\tilde{K}(t, \ell, h) &= \frac{1}{1-\sigma} \left[(wh)^{\gamma} \ell^{1-\gamma} \right]^{1-\sigma} Z(t), \\
Z(t) &\equiv e^{-\lambda(t)} \left(\frac{1}{r} \right)^{\gamma(1-\sigma)}) + \lambda(t) \int \left(\frac{1}{p} \right)^{\gamma(1-\sigma)} e^{-\lambda(t)F(p)} f(p) dp.
\end{aligned}
$$

The household's problem is then to

$$\max_{t,\ell,h} \tilde{K}(t,\ell,h) \text{ s.t. } T \geq t + \ell + h$$

Since we can factor out the search component $Z(t)$, the leisure-labor decision can be solved independently of the search decision. That is, the first-order conditions for ℓ and h together imply the following optimality condition:

$$(1-\gamma)h = \gamma\ell$$

We can then simplify the original objective function by substituting in the optimal hours decision:

$$\tilde{K}(t,\ell,h) = e^{-\lambda(t)}\frac{1}{1-\sigma}\left[\left(\frac{w\frac{\gamma}{1-\gamma}}{r}\right)^{\gamma}\ell\right]^{1-\sigma} + \lambda(t)\int\frac{1}{1-\sigma}\left[\left(\frac{w\frac{\gamma}{1-\gamma}}{p}\right)^{\gamma}\ell\right]^{1-\sigma} e^{-\lambda(t)F(p)}f(p)dp,$$

which is homogeneous of degree one in the wage rate. Therefore, the shopping decision is also independent of the wage rate if a labor choice is considered.

B.3 Proof of Lemma 3

To compute η, we expand $P(\text{p is lowest} \mid \text{contact})$ using the law of total probability:

$$\begin{aligned}
\eta(p;\lambda,F) &= \sum_{s=0}^{\infty} P(\text{p is lowest}|\text{draws}=s)\cdot P(\text{draws}=s|\text{contact}) \\
&= \sum_{s=1}^{\infty} P(\text{p is lowest}|\text{draws}=s)\cdot P(\text{draws}=s) \\
&= \sum_{s=1}^{\infty}(1-F(p))^{s-1}\cdot e^{-\lambda}\frac{\lambda^s}{s!} \\
&= e^{-\lambda}\sum_{s=1}^{\infty}\frac{[\lambda(1-F(p))]^s}{s!}\frac{1}{1-F(p)} \\
&= \frac{e^{-\lambda}}{1-e^{-\lambda}}\frac{1}{1-F(p)}\left[e^{\lambda(1-F(p))}-1\right]
\end{aligned}$$

The second line uses the fact that contacted customers cannot have zero draws. The final line uses the definition of the exponent.

Notice that $F \to 0$ as $p \to \underline{p}$. In that case, the expression simplifies to 1. As $p \to r$, $F \to 1$. L'Hôpital's rule is applied:

$$\frac{e^{-\lambda}}{1-e^{-\lambda}} \cdot \frac{\lim_{F\to 1} -\lambda e^{\lambda(1-F)}}{\lim_{F\to 1} -1} = \frac{\lambda e^{-\lambda}}{1-e^{-\lambda}} = \frac{Poisson(1)}{1-Poisson(0)}$$

The latter expression is the probability of the customer having exactly one draw, conditional on having at least one.

B.4 Proof of Lemma 4

The price distribution F is consistent with the firm's optimal pricing strategy only if

$$\pi(p; F, \lambda) = \pi^* \, \forall p : f(p; \lambda) > 0$$
$$\pi^* \equiv \max p\pi(p; F, \lambda)$$

Profits are strictly positive. Profits at the reservation price are given by

$$\pi(r, F, \lambda) = \frac{H\lambda}{S} \frac{\lambda e^{-\lambda}}{1-e^{-\lambda}} \frac{y}{p}(r-c)$$

With $\lambda > 0$, these are strictly positive since $r > c$. Therefore, $\pi^* > 0$.

The distribution is continuous. To show this, we need a few additional results. First, consider the general case of $\mu(p, F, \lambda)$ under the presence of potentially multiple point masses. Let $\psi(p)$ denote the point mass of prices at p.

Lemma 8. *If shoppers with n contacts with the same price randomly purchase from one of them, the capturing probability is given by*

$$\tilde{\eta}(p, F, \lambda, \psi) = \frac{e^{-\lambda}}{1 - e^{-\lambda}} \sum_{s=1}^{\infty} \frac{\lambda^s}{s!} \frac{(1 - F(p) + \psi(p))^s - (1 - F(p))^s}{sx}$$

Proof. Write the probability as

$$\begin{aligned}
&\eta(p, F, \lambda, \psi) \\
&= P(\text{Capture a household at CDF } F \text{ and mass } \psi(p) | \text{contact}) \\
&= \sum_{s=0}^{\infty} P(\text{Capture a household at CDF } F \text{ and mass } \psi(p) | \& s \text{ draws}) \cdot P(s \text{ draws} | \text{contact}) \\
&= \sum_{s=0}^{\infty} P(\text{Capture a household at CDF } F \text{ and mass } \psi(p) | \& s \text{ draws}) \cdot \frac{P(s \text{ draws}) \cdot P(\text{contact} | s \text{ draws})}{P(\text{contact})} \\
&= \sum_{s=1}^{\infty} P(\text{Capture a household at CDF } F \text{ and mass } \psi(p) | \& s \text{ draws}) \cdot \frac{P(s \text{ draws})}{P(s > 0)} \\
&= \sum_{s=1}^{\infty} \frac{P(s \text{ draws})}{P(s > 0)} \sum_{y=0}^{s-1} P(s - y - 1 \text{ draws from higher prices} \wedge y \text{ prices from mass } \psi(p)) \cdot P(\text{capturing under } y + \\
&= \sum_{s=1}^{\infty} \frac{e^{-\lambda}}{1 - e^{-\lambda}} \frac{\lambda^s}{s!} \sum_{y=0}^{s-1} \binom{s-1}{y} F^y (1 - F)^{s-1-y} \left(\frac{\psi(p)}{F} \right)^y \frac{1}{y+1} \\
&= \sum_{s=1}^{\infty} \frac{e^{-\lambda}}{1 - e^{-\lambda}} \frac{\lambda^s}{s!} \sum_{y=0}^{s-1} \binom{s-1}{y} \psi(p)^y (1 - F)^{s-1-y} \frac{1}{y+1} \\
&= \sum_{s=1}^{\infty} \frac{e^{-\lambda}}{1 - e^{-\lambda}} \frac{\lambda^s}{s!} \frac{(1 - F + \psi(p))^s - (1 - F)^s}{s\psi(p)}
\end{aligned}$$

□

For convenience, we rewrite this as

$$\eta(p, F, \lambda, \psi) = \sum_{s=1}^{\infty} \frac{e^{-\lambda}}{1 - e^{-\lambda}} \frac{\lambda^s}{s!} \zeta(p, F, \psi, s)$$

$$\zeta(p, F, \psi, s) = \frac{(1 - F + \psi(p))^s - (1 - F)^s}{s\psi(p)}$$

Lemma 9. $\lim_{\psi(p) \to 0} \eta(p, F, \lambda, \psi) = \eta(p, F, \lambda)$, *the special continuous case.*

Proof. To show this, it is sufficient to show that $\lim_{\psi\to 0}\tilde{\zeta}(p,F,\psi,s) = (1-F(p))^{s-1}$.

For any particular s, application of L'Hôpital's rule shows that:

$$\begin{aligned}\lim_{\psi\to 0}\frac{(1-F+\psi)^s-(1-F)^s}{s\psi} &= \frac{\lim_{\psi\to 0}\frac{\partial}{\partial\psi}(1-F+\psi)^s-(1-F)^s}{\lim_{\psi\to 0}\frac{\partial}{\partial\psi}s\psi}\\ &= \frac{\lim_{\psi\to 0}s(1-F+\psi)^{s-1}}{s} = (1-F)^{s-1}\end{aligned}$$

□

Lemma 10. *The probability of capturing a household increases in* ψ*:* $\frac{\partial\eta(p,F,\lambda,\psi)}{\partial\psi} > 0 \qquad ,\psi \in (0,F)$.

Proof. Sufficient to show that $\zeta_\psi(p,F,\lambda,\psi,s) > 0$ for $s > 1$:

$$\zeta_\psi(p,F,\lambda,\psi) = \frac{s(1-F+\psi)^{s-1}s\psi - s[(1-F+\psi)^s-(1-F)^s]}{(s\psi)^2}$$

Let $C(\psi) = 1-F+\psi$, and $H(\psi) = s^2\psi C(\psi)^{s-1} - sC(\psi)^s + (1-F)^s$.

We need to show that $H(x) \geq 0$ and have $H(0) = 0$. We have $C'(\psi) = 1$, so

$$\begin{aligned}H'(\psi) &= s^2C(\psi)^{s-1} + (s-1)s^2\psi C(\psi)^{s-2} - s^2C(\psi)^{s-1}\\ &= (s-1)s^2\psi C(\psi)^{s-2} > 0\end{aligned}$$

□

Lemma 11. *Conditional on the mass of prices below p, any point masses of prices at p reduces capturing likelihood:* $\mu(p,F-\psi(p),\lambda) = \tilde{\mu}(p,F-\psi(p),\lambda,0) > \mu(p,F,\psi(p)))$

Proof. Sufficient to show for all $s > 1$ and $\psi \in (0,F]$ that

$$\begin{aligned}\zeta(p,F-\psi,0,s) &> \zeta(p,F,\psi,s)\\ (1-F+\psi)^{s-1} &> \frac{(1-F+\psi(p))^s-(1-F)^s}{s\psi(p)}\end{aligned}$$

Define

$$C(\psi) = 1 - F + \psi$$
$$H(\psi) = s\psi C(\psi)^{s-1} - C(\psi)^s + (1-F)^s$$

The approach will be similar to the one in the proof of Lemma 10.[2] As before $C'(\psi) = 1$ and $H(0) = 0$. To show that $H(\psi) > 0$ for positive ψ, it is sufficient to show that the derivative is positive for $\psi > 0$:

$$\begin{aligned} H'(\psi) &= sC(\psi)^{s-1} + (s-1)s\psi C(\psi)^{s-2} - sC(\psi)^{s-1} \\ &= (s-1)s\psi C(\psi)^{s-2} > 0 \end{aligned}$$

, where the last inequality holds since $s > 1$ and $\psi > 0$.

□

Finally, we can turn to the main Lemma to be proven:

Lemma 12. *The distribution is continuous.*

Proof. Assume that the distribution of prices that maximize profits is not continuous : $\exists p_0 \in (\underline{p}, \overline{p}) : \quad \psi(p_0) > 0$. First, for arbitrary $\epsilon > 0$, the capturing probability is higher at $p_0 - \epsilon$:

$$\begin{aligned} \eta(p_0 - \epsilon, F(p_0 - \epsilon), \lambda, \epsilon(p_0 - \epsilon)) &> \eta(p_0 - \epsilon, F(p_0), \lambda, 0) \quad (9) \\ &> \eta(p_0, F(p_0) - \psi(p_0), \lambda, 0) \quad (10) \\ &> \mu(p_0, F(p_0), \lambda, \psi(p_0)) \quad (11) \end{aligned}$$

The first inequality (9) comes from Lemma 10. To see the second inequality, note that for all s, $\zeta(p_0 - \epsilon, F(p_0 - \epsilon), 0, s) > \zeta(p_0, F(p_0), 0, s)$, since $F(p_0) > F(p_0 - \epsilon)$. The third inequality, (11), comes from Lemma 11.

Using the previous inequalities, it is clear that firms make a strictly larger profit at price $p_0 - \epsilon$ and the implied probabilities, than at the price $p_0 - \epsilon$ and the implied probability of

[2]This elegant approach is suggested by Daniel Wainfleet.

$p_0 - \epsilon$:

$$\pi(p_0 - \epsilon, F, \lambda) = \frac{H\lambda}{S}\eta(p_0 - \epsilon, F(p_0 - \epsilon), \lambda, \epsilon(p_0 - \epsilon))\frac{y}{p_0 - \epsilon}(p_0 - \epsilon - c)$$
$$> \frac{H\lambda}{S}\eta(p_0, F(p_0), \lambda, \epsilon(p_0))\frac{y}{p_0 - \epsilon}(p_0 - \epsilon - c)$$

$p_0 > \underline{p}$ implies $p_0 > c$. This implies that there is a sufficiently small ϵ such that $p_0 - \epsilon > c$, and

$$\pi(p_0 - \epsilon, F, \lambda) > \pi(p_0, F, \lambda)$$

This implies that p_0 does not maximize profits, a contradiction. □

Intuitively, there is no price with a positive mass of sellers. If there existed a price p_0 with a positive mass of sellers, firms setting $p = p_0 - \epsilon$ for some small ϵ should make a second-order loss on the price per sold unit, but a first-order gain from the share of consumers captured.

Upper bound satisfies $\overline{p} = r$. Suppose that $\overline{p} < r$. At $\overline{p}$, we have $F = 1$ and the profits are given by

$$\pi(\overline{p}, F, \lambda) = \frac{H\lambda}{S}\frac{\lambda e^{-\lambda}}{1 - e^{-\lambda}}\frac{y}{r}(\overline{p} - c)$$

If a firm was to sell at price r, it would make profits of

$$\pi(r, F, \lambda) = \frac{H\lambda}{S}\frac{\lambda e^{-\lambda}}{1 - e^{-\lambda}}\frac{y}{r}(r - c)$$

Clearly, $\pi(\overline{p}, F, \lambda) < \pi(r, F, \lambda) \leq \pi^*$, contradicting the fact that F is consistent with profit maximization.

Now suppose that $\overline{p} > r$. In that case, firms that set $p = \overline{p}$ make zero profits, since households will not be buying. Again, $\pi(\overline{p}, F, \lambda) = 0 < \pi^*$, which is a contradiction.

The lower bound satisfies $\underline{p} \in (c, r)$. We know that $\underline{p} \leq \overline{p} = r$. If $\underline{p} = r$, the distribution would have a positive mass at r, a contradiction. If $\underline{p} \leq c$, we have $\pi(\underline{p}, F, lambda) < 0 < \pi^*$, inconsistent with profit maximization. The only remaining possibility is $\underline{p} \in (c, r)$.

The support is connected. Suppose that the support of F is not connected. Then, there exists $p_0 < p_1$ such that $F(p_0) = F(p_1)$. In that case, we have

$$\begin{aligned}
\pi(p_0, F, \lambda) &= \frac{H\lambda}{S}\eta(p_0; \lambda, F)\frac{y}{p_0}(p_0 - c) \\
&< \frac{H\lambda}{S}\eta(p_0; \lambda, F)\frac{y}{p_1}(p_1 - c) = \pi(p_1, F, \lambda) \leq \pi^*
\end{aligned}$$

This is a contradiction: no firm that sets p_0 can be maximizing profits under this F.

B.5 Proof of Proposition 2

Lemma 4 has established that $\pi(r; F, \lambda) = \pi^*$. Because the support is connected, we have that $\pi(p, F, \lambda) = \pi^* \, \forall p \in [\underline{p}, r]$. Therefore, $\pi(p; F, \lambda) = \pi(r; F, \lambda) \forall p \in [\underline{p}, r]$. Solving that identity delivers:

$$\begin{aligned}
\frac{H\lambda}{S}\eta(p; \lambda, F)\frac{y}{p}(p - c) &= \frac{H\lambda}{S}\eta(r; \lambda, F)\frac{y}{r}(r - c) \\
\frac{1}{1 - F(p)}\left[e^{\lambda(1-F(p))} - 1\right]\frac{y}{p}(p - c) &= \lambda\frac{y}{r}(r - c)
\end{aligned}$$

This provides an indirect representation of $F(p; \lambda)$. An explicit form of F is given by (5).

Lower bound $\underline{p}(\lambda)$ We solve for $\underline{p}$ using the fact that $\pi(\underline{p}; F, \lambda) = \pi(r; F, \lambda)$:

$$
\begin{aligned}
\frac{H\lambda}{S}\eta(\underline{p};\lambda,F)\frac{y}{\underline{p}}(\underline{p}-c) &= \frac{H\lambda}{S}\eta(r;\lambda,F)\frac{y}{r}(r-c) \\
\frac{y}{\underline{p}}(\underline{p}-c) &= \lambda\frac{e^{-\lambda}}{1-e^{-\lambda}}\frac{y}{r}(r-c) \\
r(\underline{p}-c) &= \lambda\frac{e^{-\lambda}}{1-e^{-\lambda}}(r-c)\underline{p} \\
\underline{p} &= \frac{rc}{r-\lambda\frac{e^{-\lambda}}{1-e^{-\lambda}}(r-c)}
\end{aligned}
$$

B.6 Proof of Lemma 3

To see the first part, note that $F_\lambda(p;\lambda) < 0\,\forall p$. For the second part, we need to show that $p'(\lambda)$. Before computing the derivative, note that

$$
\begin{aligned}
\frac{\partial\lambda e^{-\lambda}1-e^{-\lambda}}{\partial\lambda} &= \frac{(e^{-\lambda}-\lambda e^{-\lambda})(1-e^{-\lambda})-(e^{-\lambda}(\lambda e^{-\lambda}))}{(1-e^{-\lambda})^2} \\
&= \frac{e^{-\lambda}(1-e^{-\lambda})-\lambda e^{-\lambda}(1-e^{-\lambda})-e^{-\lambda}\lambda e^{-\lambda}}{(1-e^{-\lambda})^2} \\
&= \frac{e^{-\lambda}(1-e^{-\lambda}-\lambda)}{(1-e^{-\lambda})^2} \\
&= \frac{e^{\lambda}(1-\lambda)-1}{(e^{\lambda}-1)^2}
\end{aligned}
$$

Using this, we can compute the derivative as

$$
\begin{aligned}
\underline{p}'(\lambda) &= \frac{0 - (\frac{e^{\lambda}(1-\lambda)-1}{(e^{\lambda}-1)^2})(r-c)rc}{\left(r - \lambda\frac{1}{e^{\lambda}-1}(r-c)\right)^2} \\
&= \frac{[1 - e^{\lambda}(1-\lambda)](r-c)rc}{(e^{\lambda}-1)^2\left(r - \lambda\frac{1}{e^{\lambda}-1}(r-c)\right)^2} \\
&= \frac{e^{\lambda}(1-\lambda)-1}{(e^{\lambda}-1)^2}\frac{r-c}{rc}\underline{p}(\lambda)^2
\end{aligned}
$$

To sign this expression, note that the denominator is always positive. As for the numerator,

$$
\frac{\partial e^{\lambda}(1-\lambda)-1}{\partial\lambda} = e^{\lambda}\lambda
$$

The numerator has a single maximum/minimum at $\lambda = 0$. Notice that

$$
e^{\lambda}(1-\lambda) - 1|_{\lambda=1} = -1 < 0
$$

This implies that $\lambda = 0$ is a maximum, and the derivative $p'(\lambda)$ is negative everywhere else.

B.7 Proof of Lemma 5

For the first part, note that

$$
\begin{aligned}
\underline{p}(\lambda, \psi r, \psi c) &= \frac{\psi c}{1 - \lambda\frac{e^{-\lambda}}{1-e^{-\lambda}}\frac{r-c}{r}} \\
&= \psi\frac{c}{1 - \lambda\frac{e^{-\lambda}}{1-e^{-\lambda}}\frac{r-c}{r}}
\end{aligned}
$$

For the second part, recall that F is given by

$$
\begin{aligned}
F(p;\lambda,r,c) &= \frac{1}{\lambda}(z+\mathrm{LambertW}(-ze^{-z}))+1 \\
z(p;r,c) &= \frac{r}{p}\frac{p-c}{r-c}
\end{aligned} \tag{12}
$$

It is sufficient to see that $z(\psi p;\psi r,\psi c) = z(p,r,c) \Rightarrow F(\psi p;\lambda,\psi r,\psi c) = F(p;\lambda,r,c)$.

C Microfoundation of matching

Let S be the (integer) number of stores uniformly distributed on a unit circle. Before the start of the period, households commit to search for time duration t. During that time, they start at a random location on the unit circle and walk at speed a – distance/time – in a random direction.

We divide t into N subperiods of length $\Delta_N = t/N$. The number of stores met during that subperiod are binomially distributed: the probability of meeting x stores is given by

$$
\tilde{p}_N(x) = \binom{S}{x} p^x (1-p)^{S-x}
$$

What is the probability of any arbitrary store being contacted within Δ_N? Stores are uniformly distributed over the unit circle, and within Δ_N we travel at/N. Hence

$$
p = at/N
$$

Claim: as $\Delta \to 0$, the expected number of multiple draws of stores within same Δ vanishes.

Proof. We have N subperiods. For $x > 1$, the expected number of multiple draws is given by

$$
\lim_{N\to\infty} N \cdot \binom{S}{x}(at/N)^x(1-at/N)^{S-x} = \lim_{N\to\infty} N^{1-x}\binom{S}{x}(at)^x(1-at/N)^{S-x} \to 0
$$

where the N^{1-x} term vanishes, and the remainder remains constant. □

Intuitively, since S are uniformly distributed, there is a measure 0 of stores at exactly the same physical location. Therefore, for any ex-post distribution of S over the unit circle, we can pick large enough N s.t. the probability of 2 or more stores on length Δ_N is negligible.

We will hence focus on computing the probability of contacting 1 store in a subinterval Δ_N. Define

$$\begin{aligned} \lambda_N &\equiv N p_N(1) = NS(at/N)^1(1-(at/N)^{S-1}) \\ \lambda &\equiv \lim_{N\to\infty} \lambda_N = atS \end{aligned}$$

With the machinery in place, we can compute the probability of meeting s stores during the total search time t. Given that each subperiod has a binary outcome – meeting either one store with probability p_N or meeting none – the total number of stores met is also binomially distributed with the total number of draws N.

$$\begin{aligned} P(s) &= \frac{N!}{(N-s)!s!} p_N(1)^s(1-p_N)^{S-s} \\ &= \frac{\mu_N^s}{s!}\frac{N!}{(N-s)!N^s}(1-\frac{\mu_N}{N})^S(1-\frac{\mu_N}{N})^{-s} \end{aligned}$$

where the rearrangement allows me to study the limits for each component:

$$\begin{aligned} &\lim_{N\to 0}\frac{N!}{(N-s)!N^x} = 1 \\ &\lim_{N\to 0}(1-\frac{\mu_N}{N})^S = e^{-\lambda} \\ &\lim_{N\to 0}(1-\frac{\mu_N}{N})^{-s} = 1 \end{aligned}$$

Assembling the parts yields that the total number of stores contacted is Poisson distributed:

$$P(s) = e^{-\lambda(t)} \frac{\lambda(t)^s}{s!}$$

$$\lambda(t) = atS$$