



CEP Discussion Paper No 1296

September 2014

**Consumer Search:
Evidence from Path-tracking Data**

Fabio Pinna and Stephan Seiler

Abstract

We estimate the effect of consumer search on the price of the purchased product in a physical store environment. We implement the analysis using a unique data set obtained from radio frequency identification tags, which are attached to supermarket shopping carts. This technology allows us to record consumers' purchases as well as the time they spent in front of the shelf when contemplating which product to buy, giving us a direct measure of search effort. Controlling for a host of confounding factors, we estimate that an additional minute spent searching lowers price paid by \$2.10 which represents 8 percent of average trip-level expenditure.

Key words: Consumer search, in-store marketing, path data
JEL: D12; D83; L11; L15

This paper was produced as part of the Centre's Productivity Programme. The Centre for Economic Performance is financed by the Economic and Social Research Council.

We thank Swati Yanamadala for excellent research assistance. We are grateful to Herb Sorenson for providing us access to the data and to Herb and Jamin Roth for helping us understand the data better. We would like to thank seminar participants at Toronto, Michigan, Boston College, Chicago, UC Davis, and conference participants at EARIE (Évora), the Choice Symposium (Noordwijk), IIOC (Chicago), Marketing Science (Atlanta), and SICS (Berkeley) for great feedback. We also benefitted greatly from discussions with Emek Basker, Eric Bradlow, JP Dube, Daria Dzyabura, Pedro Gardete, Matt Gentry, Jonathan Haskel, Guenter Hitsch, Ella Honka, Alessandro Iaria, Guy Michaels, Chris Nosko, David Rapson, Navdeep Sahni, John Van Reenen, Matthijs Wildenbeest, Song Yao, and Hema Yoganarasimhan. All remaining errors are our own.

Fabio Pinna is a PhD student at London School of Economics. Stephan Seiler is an Associate of the Centre for Economic Performance, London School of Economics and Political Science. He is also Assistant Professor of Marketing, Stanford Graduate School of Business.

Published by
Centre for Economic Performance
London School of Economics and Political Science
Houghton Street
London WC2A 2AE

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

1 Introduction

When consumers make a purchase decision, they might often not be aware of prices for all products due to informational and cognitive constraints. In many categories, a large number of products are available, and obtaining relevant information can be costly. In a grocery-shopping context, consumers can search across stores, time their purchase in order to benefit from temporary price reductions, and search across various products within a particular store when standing in front of the shelf. In this paper, we focus on the final part of this decision process: the consumer’s search effort when processing information and comparing products and prices immediately before putting the chosen product into her shopping cart. Specifically, our goal is to estimate the effect of the extent of consumers’ search activity within a particular product category on the price they pay.

A key challenge in analyzing consumer search behavior in a physical store environment lies in the fact that observing and recording which products the consumer was considering before picking one particular product from the shelf is hard. In studies using online data instead, one typically observes the sequence of searches, as for example in De Los Santos, Hortacsu, and Wildenbeest (2013), Koulayev (2013), or Chen and Yao (2014). An alternative in a brick-and-mortar environment would be to provide consumers with eye-tracking equipment as in Stüttgen, Boatwright, and Monroe (2012). This approach provides a great level of detail but has the disadvantage of disrupting the consumer’s “natural” shopping experience. In this paper, we propose an approach to understanding search behavior without such an intervention. To this end, we use “path-tracking” data obtained from shopping carts that are equipped with radio-frequency identification (RFID) tags combined with store-level data on purchases and product prices.¹ The data allow us to measure the time a consumer spends in front of a particular category before deciding to purchase a specific product, thus giving us a direct measure of the extent of the consumer’s search activity.²

The central contribution of the paper is to demonstrate how the monetary benefits from search (per unit of time) can be estimated using data on the total duration of search as well as the price of the chosen product. To the best of our knowledge this paper, in parallel with Jain, Misra, and Rudi (2014), is the first to gather data on search effort and to estimate search benefits in a physical store environment. Using a reduced-form approach, where we regress price paid on search time in a linear regression framework, we find that an additional minute spent searching lowers expenditure by \$2.10. The magnitude is economically significant: extending search time by one standard deviation in each product category lowers total trip-level expenditure on the average shopping trip by 7%. Given that we are analyzing frequently purchased products, the potential unrealized savings are large and suggest that consumers engage in a limited amount of search activity. We also find that search activity varies greatly across different areas of the store, which suggests one possible channel through which product

¹A further source of data on consumer search behavior / considerations is survey information directly levied from consumers. Draganska and Klapper (2011) and Honka (2014) use this kind of data.

²Apart from RFID, other technology such as video capture (see Jain, Misra, and Rudi (2014) or Hui, Huang, Suher, and Inman (2013)) or smart-phone wi-fi signals might also be used to measure search time in a similar fashion.

placement and store design can influence consumer behavior. Moving a product category from the area with the lowest to the area with the highest level of search activity leads to an increase in search time of almost 16 seconds (2 standard deviations), which decreases price paid by \$0.54 and increases the probability of purchasing a promoted product by 13 percentage points. The magnitude of this effect is large and relevant for manufacturers that pay slotting allowances to place their products in certain locations inside the store. Based on search time differences, some locations do engender closer competition with other brands due to consumers engaging in more search. Similarly, pricing decisions should arguably be a function of product location. In high search locations, running a promotion will be more effective than in areas of the store where consumers' search effort is lower.

To guide our empirical analysis, we rely on the canonical sequential model of consumer search (McCall (1970)), which we use to derive predictions for the relationship between search time and price paid in response to variation in consumers' search costs. Our objective is to empirically uncover how an increase in search time *due to lower consumer search costs* translates into a lower price paid. To this end, we need to isolate variation in search time that is due to search-cost differences rather than other factors. Although a more in-depth discussion is relegated to a later point, we present the key concerns and identification assumptions here. Using the sequential search model, we identify two other determinants of search time variation. First, consumers face different price distributions over time within each category due to promotional activity, and, most importantly, drawing prices from a more "favorable" price distribution leads to shorter search spells as well as a lower price paid. Second, consumers with the same search cost will in general draw different prices from a given distribution during their respective search process. In other words, pure chance when drawing prices will lead to variation in search time. We show that this variation in search time does not affect the price paid and will therefore lead to attenuation bias in our estimate.³ A final concern is due to the nature of our data: we are able to record the time a consumer spends in the vicinity of the product category, which is a noisy measure of actual category-level search activity. The presence of this measurement error will lead to attenuation bias in an OLS regression.

To address all of the concerns above, we need to use search-cost shifters as instruments for search duration. We leverage the fact that we have information on consumer purchases as well as in-store behavior for the whole trip of which search activity within each category makes up only a small part. Specifically, we use the consumers' walking speed over the course of the trip, the total number of items purchased, and a dummy for whether the consumer used a basket (rather than a shopping cart) as instruments for search time. The identifying assumption is that exogenous variation in consumption needs and search-costs drive overall trip behavior such as walking speed and basket size. For instance, consumers might go shopping on the weekend when they are not in a rush. A larger basket size and slower walking

³Note that promotional variation over time is variation that could be used under the assumption that the timing of promotional activity is exogenous. However, one would need to take a more structural approach in order to use this variation. For instance, the fact that in a week with more promotions, search time decreases by a certain amount, does tell us something about the search process. To make a statement about search costs, one would need to take a stance on how expectations are formed, which is something our approach avoids. Similarly, a structural model could explicitly account for the role of chance in the search process.

speed would characterize this trip, relative to a quick fill-in trip at lunchtime on a weekday. Most importantly, we assume the price distribution of any particular category and luck in the search process do not influence walking speed, basket size, and the choice between using a basket or cart.

Our analysis has two main caveats. First, our data cover only a short period of time, and although we do observe some consumers repeatedly in the path-tracking data, the panel dimension is too small to exploit. Our estimation can therefore be thought of as essentially cross-sectional. This feature of the data is not unique to our setting, and indeed all empirical papers on consumer search that we are aware of are cross-sectional in nature (Kim, Albuquerque, and Bronnenberg (2010), De Los Santos, Hortacsu, and Wildenbeest (2013), Koulayev (2013), Honka (2014)). The main drawback the lack of a panel dimension creates in our setting is the inability to control for consumer preferences through repeated observations for the same consumer. Not being able to control for preference heterogeneity is a concern to the extent that search costs and therefore search-spell duration are correlated with consumers' price sensitivity. If consumers who are more price sensitive also have lower search costs, we are likely to overestimate the effect of search on price paid. Although search time is conceivably correlated with price sensitivity, such a correlation is much less clear for our set of instruments.⁴ Nevertheless, we run a set of additional robustness checks to address this concern. Specifically, our setting allows us to control for heterogeneity in preferences by using variation in search activity within consumers across different categories (mostly on the same trip) as well as panel data on purchases (but not search).⁵ Second, we have to pool the search-data across categories because we do not have enough observations at the individual category level. Having to deal with data across 150 categories and 30,000 UPCs prevents us from modeling utility across products more broadly, and we instead focus on the effect of search on price. Undoubtedly, price is not the only relevant product characteristic in CPG categories, and our approach therefore only captures one aspect of the search process. However, the effect of search on price is relevant for informing product location and pricing decisions as we demonstrate in more detail later.

Our paper contributes to various streams of literature. It is closely related to a series of seminal papers by Hui, Bradlow, and Fader (Hui, Fader, and Bradlow (2009), Hui, Bradlow, and Fader (2009a), and Hui, Bradlow, and Fader (2009b)) that introduced path-tracking data to the marketing literature. Relative to their work, which jointly describes the path as well as purchase decisions of consumers, we make little use of the actual path the consumer takes. Instead, we focus more narrowly on the consumer's search process when standing in front of the shelf containing a particular product category.⁶ In addition to the path data, we also

⁴For instance, we do not have a strong prior as to whether consumers that purchase more items are more or less price sensitive. We are indeed hoping that most variation in basket size occurs within consumers across different trips.

⁵We also note that even if panel data was available, one would still need to take care of issues such as measurement error. Therefore, although helpful, panel data cannot be considered a substitute to our IV-approach.

⁶Another application of path data is the analysis of unplanned purchases in Hui, Inman, Huang, and Suher (2013).

make use of detailed product-level price and purchase data that we are able to link to the path-tracking data set. The combination of the two data sources allows us to analyze how consumers' search duration (recorded by the path data) impacts the purchases they make (measured in the sales data). In this way, we are able to link the novel information we can get out of the path-tracking data to the literature on consumer search and consideration-set formation. To the best of our knowledge, when analyzing consideration sets in a physical store context (see, e.g., Roberts and Lattin (1991), Andrews and Srinivasan (1995), Bronnenberg and Vanhonacker (1996), Mehta, Rajiv, and Srinivasan (2003), and Seiler (2013)), the search process was usually unobserved. In this paper, we instead have a direct measure of the extent of search activity.⁷ As mentioned before, one notable exception is Jain, Misra, and Rudi (2014), who also observe search behavior in a physical store environment via video capture.

The paper also contributes to a strand of literature on consumer search that uses data on the search process, such as Kim, Albuquerque, and Bronnenberg (2010), De Los Santos, Hortacsu, and Wildenbeest (2013), Koulayev (2013), Honka (2014), or Chen and Yao (2014), mostly in the online realm. Relative to those papers, we take a more reduced-form approach to modeling search benefits rather than estimating a search model structurally. Our approach has some advantages and drawbacks. Our setting allows us to deal with measurement error in search effort, which is presumably present in many settings.⁸ However, usually the search effort or search sequence enters a highly non-linear model. Instead, in our linear setting, instrumental variables provide a simple solution. Second, we do not need to make assumptions about consumers' information sets and expectation formation, which are crucial identifying assumptions in most structural search models. On the other hand, without a more structural approach, we are not able to model consumer utility and choice as a function of product characteristics more broadly and instead focus solely on price. Finally, our approach does not lend itself easily to interesting counterfactuals such as search-cost reductions that could be achieved through various marketing tools. To a large extent, the nature of our data motivates the approach. Nevertheless, we believe that our approach has certain advantages over more structural ones, and we see it as a novel way of using search data that is complimentary to previous approaches.

The remainder of the paper is organized as follows. Section 2 provides a detailed explanation of the data used in our analysis and descriptive statistics. In section 3, we provide a theoretical framework to guide our empirical strategy and discuss identification. In section 4, we present the main results, followed by robustness checks. In section 5, we provide some interpretation for the magnitude of the estimated effect. In section 6, we explore the effect of product location on search and purchase behavior. Finally, we make some concluding remarks.

⁷A small number of studies on consumer search in a physical store environment, such as Cobb and Hoyer (1985) or Dickson and Sawyer (1990) and Hoyer (1984), employed teams of trained investigators who observed consumers in the store and recorded their search time manually. This approach allowed them to record search duration, albeit only for a relatively small sample of consumers.

⁸For instance, Honka (2014) uses self-reported data on which products consumers considered. De Los Santos, Hortacsu, and Wildenbeest (2013) assume that every visit to an online bookstore in the week prior to the purchase of a specific book is part of the search history for that specific title.

2 Data

We use data from a large store in Northern California that belongs to a major supermarket chain.⁹ The complete data set comprises three pieces: (1) sales data from the supermarket, (2) a store map with information on product-locations, and (3) data on the path a consumer took through the store for a subset of trips over a period of 26 non-consecutive days.¹⁰ Importantly, we are able to link the path data to the corresponding purchase baskets from the sales data with the help of the store map. In section (A.1) of the appendix, we provide details on how the two pieces of data are combined.

We have complete purchase data for all consumers that visited the store during a six-week window that comprises the 26 days for which we also observe the path data. This part of the data is a standard supermarket scanner data set similar to the IRI data set (see Bronnenberg, Kruger, and Mela (2008)). At the consumer-level, we observe the full basket of products as well as the price paid for each item. Unfortunately, prices for items that do not come in specific pack-sizes (e.g., fresh fruit, vegetables, meat) are not reported in meaningful units (e.g., per kilogram). We are therefore unable to use those products in our analysis. Apart from these problematic products, we are going to use data from about 150 different product categories that are stocked in the store. Over our sample period, we observe a total of about 220,000 shopping baskets. However, the path data are only available for a subset of those.

2.1 Path data

In addition to the sales data, we also have data on the path consumers took when walking through the store. We record the paths using RFID tags that are attached to consumers' shopping carts and baskets (see Sorensen (2003)). Each RFID tag emits a signal about every four seconds that is received by a set of antennas throughout the store. Based on the signal, triangulation from multiple antennas is used to pinpoint the consumer's precise location. The consumer's location is then assigned to a particular point on a grid of so-called "traffic points," which is overlaid onto the store map. The points used to assign consumers' locations are four feet apart from each other, allowing for a fairly granular tracking of the consumer. For every path, we observe a sequence of consecutive traffic points with a time stamp associated with each point.¹¹

However, not all shopping carts and baskets in the store are equipped with RFID tags. We only observe path data for a subset of about 7% of all store visits. We therefore rarely observe multiple trips for the same consumer despite the fact that we have more of a panel dimension in the purchase data. We will discuss how this limitation affects our analysis later when we present the empirical strategy. Second, even if a shopping basket is matched to the path data, not all purchased items in the basket necessarily have a match in the path data.

⁹We are not able to disclose the identity of the supermarket. The store has a fairly typical format with a trading area of about 45,000 square-feet and a product range of 30,000 UPCs.

¹⁰The days in the path data are 8/24/2006 - 8/29/2006 and 9/7/2006 - 9/26/2006.

¹¹If a consumer moves further than to an adjacent traffic point between signals, the movement over traffic points in between the signals is interpolated. As the signal is emitted at a high frequency little interpolation is necessary for most trips.

This lack of a match can happen if the consumer leaves her cart or basket behind and thus the data do not capture the item pickup.

The primary variable of interest derived from the path data is the time a consumer spends stationary at a certain point in the store when picking up a product. An individual item purchase, or, more precisely, the “pickup” of the item from the shelf, constitutes the unit of observation in our regressions, and we observe a total of around 34,000 pickups in the data. Using the store map, we match the grid of traffic points to product locations that are within reach of the consumer from a given traffic point.¹² For a given path and set of products in the basket at the checkout, we can then use the store map to determine when the consumer picked up the product, as well as how long she spent in front of the shelf. In other words, the item pickup is defined as the moment the consumer walked past a specific product that we later see in her purchase basket. To compute search time, we measure the time elapsed between (1) the moment the consumer is first located at a traffic point assigned to the product and (2) the point in time when she moves on to a traffic point outside of the assigned area. Figure (1) illustrates how search time is assigned to a product pickup. This metric gives us a measure of time spent in the vicinity of the product that was ultimately purchased. For convenience of exposition, we will refer to this metric as search time. However, we recognize that it is a noisy measure of actual search activity, and the consumer might have been doing other things at the same time. The presence of such measurement error will inform our empirical strategy later.¹³ Figure (2) shows the histogram for our search metric across item pick-ups. The variable is roughly log-normally distributed with a mean of 10.3 seconds and a standard deviation of 8.5 seconds.

Furthermore, we also compute the speed at which the consumer moves over the course of the trip, using time stamps and distances between consecutive traffic points. Speed, although not the primary focus of this paper, will play a role in our empirical strategy. Basic descriptive statistics for the key variables used in the empirical analysis are reported in the top two panels of Table (1). These variables include trip characteristics such as average speed throughout the trip and trip duration as well as the pickup-specific measure of search time and price paid.

2.2 Price Dispersion and Possible Savings from Search

We conduct the empirical analysis using data that are pooled across product categories. To control for category-specific differences in price levels and average search-spell duration, we include a set of category fixed effects in all our regressions. In other words, we model how a consumer’s search activity *within* a category affects the particular product she buys from that category. In total, we have around 150 categories that are defined as groups of products that are naturally substitutes for each other but not with other products outside of the category. Examples for categories defined in this way are Bacon, Beer and Bird Food.

¹²The data provide the linkage between traffic and product points. Mostly any product location is associated with two or three traffic points.

¹³Furthermore, we only observe the movement and stationarity of the cart, and not the consumer herself. The fact that the consumer might leave her cart or basket behind contributes to measurement error in the search time measure.

To quantify the possible benefits of search, we report the category-specific differences between the highest and the lowest price in the category. Because prices for the same product vary substantially over time, we compute the difference between the minimum and maximum price for each day/category combination. We then compute the average of this variable across days for each category. The first row in the bottom panel of Table (1) reports the distribution of the min-max price difference across categories. On average, we find a price difference of \$4.21, but this difference varies across categories. At the 25th percentile, the price difference is equal to \$1.78 and it rises to \$5.26 at the 75th percentile. We also report the *percentage* difference of the lowest daily price relative to the highest daily price in the category in the second row of the same panel.

Because prices vary substantially due to promotional activity, we also report some descriptive statistics on the time-series variation in prices. For the purpose of this exercise, we define a promotion as a daily price that lies at least 15% below the maximum price of that product over our sample period. Similar to the calculation for the price difference, we compute the share of promoted products for each day/category pair and then take the average across days for each category. The distribution across categories is reported in the third row. On average, about 30% of UPCs within a category are on promotion. Furthermore, even within our short time window, many different products go on promotion. To document this feature of the data, we compute for each category the percentage of UPCs that went on promotion *at some point* during the six-week sample period of the sales data. The average across categories is almost 60% which is substantially higher than the daily share of promoted products, indicating that the identity of the set of promoted products changed frequently.

Taken together, the large within-category price dispersion as well as the substantial degree of promotional activity suggest there are gains from search. The average category-level saving of \$4.21 might seem relatively small compared to other (non-CPG) product categories, but relative to the amount of total shopping expenditure, it is not trivial. Consumers buy on average in seven product categories on a shopping trip, which would allow for maximum savings of roughly \$29. Furthermore, consumers can realize these gains on a regular basis on each shopping trip, and they are therefore of a large overall magnitude.

3 Model and Identification Strategy

In this section, we outline the predictions of the canonical sequential search model described in McCall (1970) and describe how the model maps onto our specific context and data. In the sequential model, consumers receive draws from a distribution of utilities and optimally decide when to stop searching. In our context, consumers search across products within a category, and we assume they care only about price and not about other product characteristics. The fact that we are not estimating the model structurally motivates this simplification. We therefore choose a simple setup that is rich enough to illustrate the key aspects that play a role in the identification strategy. Later, we extend the model to search over product characteristics other than price, and analyze the implications of such a modified model. For

now, we assume a consumer gets gross utility v if she consumes any product within the category. Further, the consumer incurs a search cost $c_{product}$ when evaluating an additional product and receives a draw from the price distribution $F(p)$ with support $[p, \bar{p}]$ for each search attempt. The optimal stopping rule is a time-invariant threshold rule λ (i.e., the consumer will accept any price below λ) that maximizes the consumer's value function¹⁴

$$EV = -c_{product} + \int_p^\lambda (v - p)dF(p) + (1 - F(\lambda))EV. \quad (1)$$

Alternatively, one can interpret the optimal stopping rule as the value of λ , which equates the marginal benefit with the marginal cost of searching:

$$\int_p^\lambda (\lambda - p)dF(p) = c_{product}.$$

One can easily see that the optimal threshold λ is increasing in search costs $c_{product}$. Intuitively, a higher search cost will make the consumer less picky and therefore willing to accept a higher price.

In the standard search model, we can think of $c_{product}$ as representing the cost of resolving uncertainty about one more option. However, in our data, we are not able to measure the number of options evaluated; instead, we only know the extent of search activity measured in real time. To adapt the model to our setting, we model the search cost of evaluating one more alternative as $c_{product} = TimePerSearch * c_{time}$, the product of time needed to search one option ($TimePerSearch$), and the opportunity cost of time (c_{time}). $TimePerSearch$ represents the efficiency of the search process. It might (as any of the other model primitives) vary across consumers. For simplicity of exposition, we ignore any consumer i subscripts. The conversion into real time leads to a slightly modified optimality condition:

$$\int_p^\lambda (\lambda - p)dF(p) = TimePerSearch * c_{time}. \quad (2)$$

Using this condition, it is easy to show that the expected price paid is equal to

$$E(p) = \frac{1}{F(\lambda)} \int_p^\lambda pdF(p), \quad (3)$$

and the expected number of searches is determined by

$$E(N) = \frac{1}{F(\lambda)};$$

therefore, the expected time spent searching is given by

¹⁴We ignore discounting due to the short amount of time consumers spent searching in a given category in our data. We also assume v is high enough such that the consumer searches at least one option. We interpret this assumption as the consumer being committed to buying in the category but deciding which specific product to pick from within the category.

$$E(\text{SearchTime}) = \text{TimePerSearch} * E(N) = \frac{\text{TimePerSearch}}{F(\lambda)}. \quad (4)$$

Note that λ is the optimal stopping rule defined by equation (2) and is therefore a function of *TimePerSearch*. A larger amount of time needed to make an additional search will increase search costs and therefore increase the stopping threshold λ , making the consumer willing to accept higher prices.

Based on the relationships derived above, we can trace out how search time and price paid change when varying the consumer’s search costs. Lowering search costs leads to a lower stopping threshold (λ) as determined by equation (2), which in turn increases expected search time and decreases expected price (see equations (3) and (4)). The relationship is non-linear with extensions in search time from a lower level being associated with larger gains in terms of finding lower prices. Moreover, the lower bound of the price distribution bounds the potential gains within a category. Figure (3) illustrates this relationship graphically. This relationship between search time and price paid is what we aim to uncover: we want to quantify the impact of an additional second spent searching *due to a lower search cost value* on price paid.¹⁵ This relationship can also be thought of as an estimate of the “search technology” that allows consumers to convert the invested search time into price savings.

We reiterate that because we do not structurally estimate the search model, its primary role is to guide our thinking regarding identification. Most search models that we can think of generate a relationship between search time and price paid when varying search costs as the one described in Figure (3). For instance, a search model for differentiated goods where consumers are aware of brand-specific utilities but have to resolve uncertainty over price as in Honka (2014) (based on Weitzman (1979)) would also predict longer search spells and a lower price paid when search costs decrease. Similarly, a satisficing model as in Stüttgen, Boatwright, and Monroe (2012) would lead to consumers evaluating more options and finding a lower price when their search costs are lower. Because the search literature features a multitude of different modeling approaches, we are unable to outline the implications of each approach in detail here. Most of the discussion on identification below, however, does not hinge on the model specifics. We therefore think of our model framework as the simplest one to illustrate the relevant aspects of the empirical strategy. The main shortcoming is that the model, for the sake of simplicity, focuses on price search. In a later section, we extend the model to allow for search over other product characteristics. This aspect has important consequences for our estimation and we discuss it in greater detail later.

3.1 Identification

In line with the reasoning outlined above, one way to think about our empirical strategy is the following: if all the observed differences in search time were caused solely by differences

¹⁵For specific distributional assumptions on $F(p)$, it is easy to rewrite expected price explicitly as a function of search time. For a uniform $[0, 1]$ distribution, for instance, we obtain $E(p) = \frac{\lambda}{2} = \frac{\text{TimePerSearch}}{2 * E(\text{SearchTime})}$ and $E(\text{SearchTime}) = \frac{\text{TimePerSearch}}{\lambda} = \sqrt{\frac{\text{TimePerSearch}}{2 * c_{time}}}$

in consumers’ search costs c_{time} , our data would trace out the relationship depicted in Figure (3). A simple OLS regression, possibly allowing for a non-linear effect, could be used in this case. As we outline in more detail below, factors other than search cost variation are likely to affect search time (and price paid). The OLS estimate is therefore unlikely to allow us to estimate the desired relationship.

In the absence of direct information on search costs, we need to find variables that are correlated with search costs and can thus be used as instruments. In other words, we need search-cost shifters. In our baseline specifications, we use the consumer’s walking speed over the course of the trip, the total number of items purchased on the shopping trip, and a dummy for whether the consumer used a basket (rather than a shopping cart) as instruments for search time. The identifying assumption is that exogenous variation in consumption needs and search-costs drivers overall trip behavior such as walking speed and basket size. In other words, search time and our instruments are correlated because they are both affected by a latent third variable: search costs. As we argue in more detail below, our instruments only affect search time due to their correlation with search costs and not any other factors affecting the search process.

In the following sections, we lay out the evidence in support of our exclusion restriction. In particular, several factors will lead to bias in an OLS regression: (1) variation in category-level promotional activity over time, (2) the role of chance in the search process, and (3) measurement error in search time. We argue that none of these three factors is likely to be correlated with our instruments. The main reasoning with respect to all three confounds is that they constitute relatively localized factors that influence the category-level search process and the measurement of it. However, the typical search spell last 10 seconds, whereas the average trip duration is 23 minutes. We thus argue that factors that are specific to search within any particular category are unlikely to affect trip-level characteristics such as the number of purchased items, walking speed, and the choice to use a basket or cart.

The main confound that our instruments might not fully address is a potential correlation between search time and price sensitivity. If consumers who are more price sensitive also have lower search costs, we are likely to overestimate the effect of search on price paid. Although search time is conceivably correlated with price sensitivity, such a correlation is much less clear for our set of instruments. Nevertheless, we later run several additional checks to probe whether heterogeneity in preferences over product characteristics other than price might cause bias in our estimate.

3.2 Category-level Price Variation over Time

Consumers form expectations knowing that prices vary both across products and over time. The latter dimension is particularly important in the grocery-shopping context due to the presence of high-frequency price movements. As shown in section (2.2), price reductions due to promotions are common in our data. Both dimensions are embodied in the price distribution governing the expectation process $F(p)$. On any given day t , a price distribution

$F_t(p)$ across products exists that is (in most cases) not known to the consumer¹⁶, but that will influence the length of the search process as well as the expected price. Formally, this situation corresponds to the threshold value of the stopping rule being determined by $F(p)$, whereas the expected number of searches and the expected price are a function of $F_t(p)$ ¹⁷:

$$E(p) = \frac{1}{F_t(\lambda)} \int_{\underline{p}}^{\lambda} p dF_t(p)$$

$$E(SearchTime) = \frac{TimePerSearch}{F_t(\lambda)}$$

A price CDF with more weight in the left part of the distribution characterizes days with more promotional activity. This leads to a lower expected search duration, as can be seen from the equation above. The impact on $E(p)$ is in principle ambiguous and depends on how the mass of the probability density function moves with respect to the threshold. When more products are promoted, more prices will lie below λ . However, depending on where those prices lie within the *truncated* distribution, the expected price paid will increase or decrease. With our data, we are able to directly test whether changes in $F_t(p)$ have any impact on search time and price paid. We implement such a test by regressing time spent searching (and price paid) on the fraction of products promoted within the category and a set of category fixed effects. Doing so, we find a negative and significant effect in both regressions.¹⁸

This result shows that changes in the price distribution lead to a movement of price paid and search time in the same direction.¹⁹ This co-movement could potentially mask the negative effect of search time on price that we are seeking to uncover. Put differently, our object of interest is the effect of extending search time on the price a consumer obtains from a *given* price distribution. That is, we want to estimate the relationship depicted in Figure (3) by isolating variation *along the curve*. Changes in the price distribution instead constitute a movement of the curve itself. We therefore need the instruments to be uncorrelated with promotional activity in the specific category and only shift search costs. This condition is

¹⁶It could be known to the consumer in some circumstances, such as information about promotions being available through feature advertising. We will address this issue later.

¹⁷Strictly speaking, both $E(SearchTime)$ and $E(p)$ are also still a function of $F(p)$, which determines the optimal stopping threshold λ .

¹⁸More specifically, we regress time spent searching (and price paid) by consumer i in category c on day t on the fraction of products promoted within the category and a set of category fixed effects as well as day fixed effects:

$$SearchTime_{ict} = \alpha * FractionPromotedProducts_{sct} + \xi_c + \delta_t + \varepsilon_{ict} \quad (5)$$

$$Price_{ict} = \tilde{\alpha} * FractionPromotedProducts_{sct} + \tilde{\xi}_c + \tilde{\delta}_t + \tilde{\varepsilon}_{ict}, \quad (6)$$

where ξ_c ($\tilde{\xi}_c$) denotes the category fixed effect and δ_t ($\tilde{\delta}_t$) denotes the day fixed effect. The predictions outlined above correspond to a negative coefficient α in the first regression. The prediction for $\tilde{\alpha}$ instead is ambiguous. Note that controlling for category fixed effects is important here because promotional activity and search might vary across categories for a host of reasons. Table (B1) in the appendix reports the results from both regressions. Note also that we are less concerned with measurement error in search time in this regression because search is used as the dependent variable.

¹⁹Note that although the impact on price is theoretically ambiguous, we do find a significant negative effect of promotional activity in the regression.

likely to be fulfilled as long as consumers do not have any price information before arriving at the shelf. Even if consumers obtain information about pricing from promotional flyers and/or in-store displays, the IV is only invalid if consumers adjust their walking speed and/or basket size to the price information for any specific category, which we don't consider to be a likely scenario.²⁰

3.3 Measurement Error

In our data, we are able to measure time spent in the vicinity of the product category, which presumably is a noisy measure of actual category-level search activity. In particular, measurement error in search time might arise for a variety of reasons: the consumer might be looking at other categories nearby, leave her cart behind, or simply spend part of the time engaging in search-unrelated activity. In other words, we are not dealing with measurement error that arises simply from imperfections in the data-recording process. Instead, search time as recorded in the data can be seen as a proxy for actual search effort. As usual, the presence of this measurement error will lead to attenuation bias in an OLS regression setup. Given the nature of our data, this issue could potentially be quite severe. However, as just outlined, we think of measurement error as arising from localized and relatively isolated occurrences such as the consumer leaving the cart behind or contemplating a purchase in another nearby category. The measurement error is therefore unlikely to be correlated with the trip-level variables used as instruments.

3.4 Chance and Search-Spell Duration

Finally, another issue, specific to our context, might cause attenuation bias in a similar way as measurement error. A sequential searcher can be more or less lucky in how quickly she comes across a price draw that lies below her stopping threshold. However, the expected price conditional on having already searched a certain number of times remains unchanged. In other words, whether the consumer searched only once or 10 times, conditional on not having stopped yet, the expected price is always equal to the unconditional price expectation at the beginning of the search process:

$$E(p|p_1 > \lambda, \dots, p_k > \lambda) = E(p) = \frac{1}{F(\lambda)} \int_{\underline{p}}^{\lambda} p dF(p),$$

where p denotes the price of the actually purchased product, and p_1 to p_k denotes the price draws for the k options searched so far (without having stopped). The intuition for this result can be easily obtained from the basic dynamic optimization problem in equation (1).

²⁰Note that variation in pricing over time could in principle be exploited in order to study the search process. More specifically, search time differences for consumers facing different price distributions can be informative about consumers' stopping thresholds and therefore ultimately their search costs. However, we do not see a simple way to use this variation without employing a more structural framework that allows us to invert search time differences into the implied stopping thresholds and therefore search costs. Although such an analysis is in principle feasible, it is not the approach we are taking in this paper.

As long as prices above the threshold are drawn, the consumer always finds herself back in the same situation with an unchanged value function when making the decision to continue searching.

To fix ideas, assume a set of consumers exists with identical search costs (in terms of both c_{time} and $TimePerSearch$), and therefore identical threshold value λ . The actual duration of their respective search spells will in general be different although the expected duration is the same, and this difference depends entirely on the sequence of price draws they receive. Furthermore, consumers with longer spells will not pay different prices on average, because the expected price conditional on the number of unsuccessful searches is the same as the unconditional expected price. This mechanism generates variation in the duration of search spells, which is uncorrelated with price.

Remember we want to find the effect of search time on price caused by a change in search costs. In other words, we want to know how much less a consumer pays who searches more *on average* because she is pickier. Therefore, we want to get rid of the variation in search duration that is caused by similarly picky consumers being more or less lucky with their price draws. In a similar vein as measurement error, the chance-induced variation in search-spell duration would bias the effect of search time on price paid toward zero. It seems safe to assume our instruments are not correlated with chance during the search process and the IV should therefore deal with this issue.

4 Main Results and Robustness Checks

To analyze the impact of search time on the price paid within a category, we run the following regression:

$$p_{ijt} = \beta * SearchTime_{ijt} + \zeta_c + \varepsilon_{ijt}, \quad (7)$$

where p_{ijt} denotes the price consumer i pays for product j that she purchased on day t . ζ_c denotes a category fixed effect, the subscript c denotes the category to which product j belongs, and ε_{ijt} denotes the error term. We include a full set of category fixed effects in all our specifications because we want know whether *within* a given category, longer search leads to a consumer picking a lower-priced product. We cluster standard errors at the consumer level to allow for an arbitrary within-consumer correlation of the error terms.

A few comments are in order regarding the interpretation of the coefficient on search time in a regression that pools observations across different categories. We framed the discussion in the theoretical model around search within one category; it is therefore reasonable to ask how to interpret the effect of search on price paid when this relationship might differ across categories. In terms of our model, differences in the estimated effect across categories could be due to either differences on the benefit side of search or differences in search costs. Differences in the price distribution drive the former, with benefits being larger in categories with high price dispersion. To illustrate the aggregation across categories, consider the relationship

between search time and price within each category as illustrated in Figure (3). Given the shape of this relationship, our linear estimate will recover the slope of the curve for the average consumer in the sample. Figure (4) illustrates the local, average nature of our estimate in more detail. In particular, the magnitude of our estimate depends on whether consumers in our data search relatively little (represented by the red scatter plot) or a lot (the blue scatter plot). In the latter case, the average consumer realizes more of the potential gains from search, and the incremental benefit at the margin is smaller. We can therefore interpret our estimates as the average consumer’s marginal benefit from search or the unrealized potential gain from extending search by another second.²¹ We would expect a rational consumer to equalize the marginal benefits of search across categories. This behavior implies that even if the shape of the relationship between price and search differs across categories, the slope at the optimal stopping point will not. Therefore, differences on the benefit side will in general not lead to different effect magnitudes across categories.²² Cross-category heterogeneity in the effect could, however, arise due to search-cost difference caused by differences in product locations and placement. For instance, categories with more facings per UPC might have higher search costs because different UPCs are farther away from each other. To the extent that such differences exist, our estimate can be interpreted as an average treatment effect across categories. Note that we tried to estimate the effect of search on price for individual categories, but were unsuccessful due to an insufficient number of observations per category.²³ On the positive side, our estimates allow us to aggregate savings to the trip level, which would be difficult with estimates from individual categories.

Results using the regression presented in equation (7) are reported in Table (2). We start by running the regression by OLS, which yields a negative and significant coefficient of search time on price. The coefficient is equal to -0.0071; in other words, an additional minute spent searching would lower the price paid by about 40 cents. To deal with the various issues described in the previous section, we implement our IV-strategy. In the baseline regression, we use the consumers’ walking speed over the whole trip as our instrument. This instrument constitutes our preferred specification because speed presumably most directly reflects the extent to which the consumer is in a hurry and therefore her search costs on the particular trip. With walking speed as an instrument, we find that the first-stage coefficient, reported in column (2), is highly significant with an F-stat of 619.62. Column (3) reports the coefficient of the effect of our (instrumented) measure of search time on price. We find a negative and significant effect of -0.0344, which is substantially larger than the OLS estimate of -0.0071, showing that the issues described previously had a substantial impact on the magnitude of the OLS coefficient. Quantitatively, the point estimate of the IV corresponds to about a \$2.10

²¹Note that based on our model, we think of the consumer as having rationally decided not to continue searching, because the expected gains were lower than her search costs. Our estimate therefore represents the potential gains the consumer optimally decided not to realize in the search process.

²²This logic is only true under the assumption that consumers have correct expectations about prices across categories. If consumers over- or underestimate their benefits from search (due to incorrect price expectations) differentially across categories, we would see heterogeneity in the estimated effect.

²³We also tried to investigate heterogeneity across different groups of categories such as categories with high and low price dispersion. Again, we found that we did not have enough data to draw any reliable conclusions.

drop in price paid for an additional minute of search. We will return to the interpretation of the effect magnitude in more detail later.

As outlined before, we think of walking speed as being reflective of the consumers' underlying search costs. Apart from speed, we would also expect other trip characteristics to be correlated with consumers' search costs, which allows us to test the robustness of our results to alternative instruments. Specifically, we expect search costs to be lower on trips with a larger overall basket size, because consumers are more likely to engage in such trips when they are under less time pressure. We operationalize this idea using two variables as instruments: the number of purchased items and a dummy for whether the consumer used a basket rather than a shopping cart. The first and second stage for this alternative specification are reported in columns (4) and (5) of Table (2). Both instruments are significant and have the expected sign. The joint F-stat is equal to 63.44 and thus weaker than our specification using speed as an instrument. The second-stage coefficient is equal to -0.0528 and statistically significant. Although larger in magnitude than the coefficient reported in our baseline specification in column (3), the two coefficients are not significantly different from each other.

In Table (B2) in the appendix, we further explore specifications with alternative instruments. We report first- and second-stage results using each of the above instruments on its own as well as all three instruments from the previous specifications together. We also employ two further instruments that capture "trip size" in a similar vein as the number of items purchased and the basket dummy. In particular, we use the duration of the trip (in minutes) and the in-store walking distance from the beginning to the end of the trip. Both instruments predict search time well and yield a similar second-stage coefficient. Once we include them together with the other three instruments, however, they become insignificant. In other words, they do not provide much additional explanatory power. We also note that whereas the second-stage coefficients are not significantly different from each other, the point estimates do vary somewhat in magnitude. Our baseline coefficient has the smallest magnitude among all specifications and we are therefore, if anything, likely to underestimate the impact of search.²⁴

As a further robustness check regarding our choice of instruments, we also run a specification in which we use the consumer's walking speed in the minute preceding a specific item pickup. Relative to the three instruments used previously, this instrument has the advantage that it varies over the course of the trip and therefore differs across item pickups on the same trip. This feature will be useful for a robustness check later. However, the validity of the instrument hinges on a clear delineation of the actual search process around a particular pickup. If the beginning of the search spell is defined incorrectly, we might capture some part of the search process in the speed measurement leading up to the pickup. Any measurement error in

²⁴As an instrument for search time, we also considered using congestion in the aisles, measured by the number of carts present in that area within a particular time window. However, this instrument presumably shifts *TimePerSearch* (by making it more cumbersome to uncover prices) rather than c_{time} . As we explain in more detail in section (4.3), this is not the type of variation we want to use. Furthermore, one could imagine that congestion leads to some form of social interaction when seeing other consumers buying particular brands. Due to these concerns, we do not consider congestion a suitable instrument for our situation. We also note that including congestion as an additional control variable in our baseline specification does not alter the results.

search time might therefore also affect speed immediately prior to the pickup. For this reason, we consider this instrument to be potentially more problematic than our baseline trip-level speed instrument. Note that trip-level speed is calculated over an average total trip length of 23 minutes and is therefore unlikely to be affected by individual search spells that last only about 10 seconds. When running the regression using speed before the pickup as an instrument, we obtain a highly significant first-stage coefficient of -3.577 with an F-stat of 1289.75. This level of significance is stronger than our baseline instrument, presumably because speed prior to a pickup varies across purchases within a trip and because it is more predictive of search time than speed over other segments of the trip. The second-stage coefficient is equal to -0.0278 (standard error of 0.0061) and not significantly different from our baseline result.

Finally, we re-run our main specification but change the dependent variable: instead of price paid, we use an indicator variable that is equal to 1 if the consumer picked a product that was on promotion. Note that the number of observations is smaller because we need to observe regular purchases of a particular product in order to define when it went on promotion.²⁵ We therefore drop pickups of products for which we cannot compute the promotion indicator. As before, the instrument is strongly correlated with search time with an F-stat of 385.61. The results differ slightly from our baseline first stage only due to the difference in the number of observations used.²⁶ In the second stage, the magnitude of the coefficient (standard error) on search time, reported in column (6) of Table (2), is 0.0082 (0.0046); that is, an additional minute spent searching increases the likelihood of finding a promotion by 50 percentage points ($0.0082 * 60 = 0.492$). This specification shows that our effect is not estimated from consumers with longer search spells buying possibly lower-quality products with lower base prices. Instead, longer search spells make consumers more likely to buy a promoted product.²⁷

Next, we use the sequential search model to systematically run through a battery of robustness checks. Despite the fact that we do not structurally estimate the search model, it nevertheless provides a natural starting point to guide the sensitivity analysis. In particular, we consider how variation in each of the model primitives influences search time and price paid as well as how it relates to the consumer’s walking speed and our other instruments. The search model is quite parsimonious; therefore, the set of model primitives we have to consider is small and comprises the price distribution F_p and search efficiency ($TimePerSearch$). We further investigate several extensions of the simple model: (1) a model in which consumers have preferences over non-price characteristics and therefore do not search only for a lower price, (2) deviations from rational expectations which influence the consumer’s perceived benefit

²⁵We define a promotion as a price reduction of at least 15% relative to the product’s base price.

²⁶We replicated the baseline regression using only the observations for which the promotion dummy is defined, and find results that are not significantly different from the ones using the full sample. This finding reassures us that issues of sample selection are unlikely to contaminate the analysis.

²⁷Product with different quality levels might go on promotion more or less often. However, in our data, we find no relationship between base price (which is presumably reflective of product quality) and promotional frequency. At the product level we regress the fraction of days a product is promoted on the baseline price and a set of category dummies. We run the regression for the set of 5,848 UPCs for which we are able to define the promotion dummy. The coefficient on the baseline price is very small and insignificant with a coefficient (standard error) of 0.0016 (0.0018).

from searching, and (3) the scenario in which consumers have information about prices before arriving at the category location in the store.

4.1 Search over other product attributes

One threat to the validity of our estimation lies in the fact that consumers are likely to not only consider price, but also to search over a broader set of product characteristics. Because products in most categories are quite differentiated and consumers presumably have heterogeneous tastes over product attributes, it is natural to ask how search over multiple product attributes interferes with our analysis. In the search model, we can capture multi-attribute search by making the product valuation term v consumer/product-pair specific:

$$EV = -c_{product,i} + \int_{\lambda}^{\bar{u}} (u_{ij})dG_i(u_{ij}) + G_i(\lambda)EV,$$

where $u_{ij} = (v_{ij} - \alpha_i p)$ denotes utility, which is a function of both price and brand preferences. α_i denotes the individual-specific price coefficient and v_{ij} represents the consumer-specific valuation of product j . $G_i(u_{ij})$ is the cumulative density function that describes the distribution of utilities across products for consumer i .²⁸ In this framework, consumers will find higher-utility products as they search longer. A higher utility could be achieved either by a lower price or by finding a product that is preferable along other product dimensions, that is, that has a higher realization of v_{ij} .

First, note that the presence of preferences over other product characteristics does not necessarily invalidate our analysis. For instance, consider the situation in which price sensitivity and brand preferences are randomly distributed across consumers. The higher the weight on brand preferences (relative to price), the lower will be the effect of search time on price. However, this scenario does not introduce bias into our analysis. Instead, if preferences are correlated with search costs across consumers, this correlation could cause a problem for our estimation. For instance, one could imagine lower-income consumers have a stronger preference for lower prices relative to quality and also have lower search costs. These consumers would be searching longer as well as picking a lower-priced product from a given set of searched products due to their preferences. More formally, this pattern implies that $c_{product,i}$ and α_i are negatively correlated across consumers. Such a correlation would lead to an upward bias (in absolute terms) in the effect of search time on price. This issue is not specific to our setting, and as in our paper, the data used in other search-related papers have usually been cross-sectional, thus making it difficult to disentangle search costs and preferences.

To deal with the issue of separating price sensitivity from search-cost differences, one would ideally want to observe the same consumer searching multiple times. Under the assumption that preferences are time invariant but search costs are not, one could then identify the effect of the latter by comparing search-spell length and price paid across purchases and

²⁸Note that the threshold now denotes the *minimum* utility level at which the consumer will stop searching. In the price-search model, the threshold denoted the *maximum* price at which to stop.

searches of the same consumer. Unfortunately, such data are not available to us due to the short time window of our sample. Instead, we hope our instruments provide a cross-sectional substitute for identification through panel data. This approach is valid as long as our instruments vary primarily within but not across consumers. To illustrate the idea, assume that all consumers have small and large basket trips, but average basket size does not vary across consumers. In this case, our number-of-items instrument would randomly pick some consumers that happened to be on a large-basket trip and others on a small-basket one. The two groups would, however, not differ by their price sensitivity. We think our specification using the number of purchased items and the basket dummy fares well in this respect. Both variables are likely to mostly vary *within* consumers as a function of consumption needs and search costs. Even if average basket size differs across consumers, we see no reason why the across-consumer variation would be correlated with price sensitivity. Our baseline speed instrument is somewhat more problematic in this respect because speed might be correlated with age, which is likely to also affect price sensitivity. The fact that results are similar to the basket-size IV-specification is therefore reassuring.

To further probe the robustness of our results, we run two additional tests that leverage *within-consumer* variation in our data. Although second best to having an actual panel of consumer search behavior, our data do provide us with a panel aspect along two other dimensions that was not available in previous research on consumers search. First, we do have a panel of about six weeks for the purchase data.²⁹ Second, within a given trip (and occasionally across trips), we observe the same consumer searching and purchasing in multiple categories. The within-trip dimension thus provides us with repeated observations of search behavior for the same consumer, albeit in different categories.

To exploit the panel variation in the purchase data in a simple way, we compute for every UPC/day pair the percentile of each UPC’s price in the respective category’s (day-specific) price distribution. We then take the average of the price percentiles for all purchases we observe for the same consumer,³⁰ which gives us a simple measure of consumer-specific price sensitivity. We include this metric as an additional variable in our baseline IV-specification. We can only compute this metric for the set of consumers for which we have multiple observations and loyalty-card information that allows us to link multiple trips of the same consumer. The elimination of households without multiple observations leads to a reduction in sample size. In columns (2) and (3) of Table (3), we report the first and second stage for our baseline specification using the smaller sample. We then re-run the IV with the additional price-percentile control variable in columns (4) and (5). Doing so, we find that price sensitivity does not predict search time and is insignificant in the first stage. The coefficient on walking speed hardly changes due to the additional control. In the second stage, we find, unsurprisingly, that the consumer’s average price percentile is a strong predictor for price paid. However, the coefficient on search time remains almost unchanged. As the comparison between column (3) and our full-sample baseline regression in column (1) shows, the slight

²⁹Because only a small set of cart and baskets is equipped with the RFID, the panel dimension does not extend to the search data.

³⁰To avoid circularity, we omit purchases from trips for which we measure search time in the path data.

change in magnitude is primarily due to the change in sample size. We also computed the absolute and percentage difference of a UPC’s price to the maximum price in the respective distribution in order to ensure the functional form of the price-sensitivity variable does not drive our result. Using these alternative measures as control variables yields very similar results to the price-percentile control.

Next, we run a robustness check that controls for individual-specific differences in search and purchase behavior by including a set of consumer fixed effects. In this way, we are only identifying the effect of search from *within-consumer* variation in search time. However, note that we rarely observe the same consumer searching in the same category repeatedly. Most of the identification does in fact come from *within-trip* variation in search behavior across different categories. Although we have a small number of consumer for which we observe multiple trips, this dimension provides relatively little variation. To implement a regression with consumer fixed effects, we therefore need an instrument that varies at a more granular level than the trip-level instruments used previously. To this end, we use walking speed over the minute preceding a specific pickup as an instrument in the fixed-effect specification. This instrument allows us to use within-trip variation in speed, but has some shortcoming, which we discussed in section (4). The results from this regression are reported in columns (6) and (7) of Table (3). As a point of reference, we first run a specification without consumer fixed effects using the new speed instrument. In column (7), we then also include consumer fixed effects and find an effect of search time on price paid of -0.0178 (standard error of 0.0067), which is not significantly different from our baseline specification.³¹ This robustness check deals with preference heterogeneity only as long as a consumer’s price sensitivity does not vary across categories and trips but search costs do. If instead consumers are more price sensitive in some categories than in others, for instance, due to a stronger preference for quality relative to price in some categories, then consumer fixed effects might not fully address the issue. However, even category-specific preferences are only problematic if search costs are also category specific in a way that creates a spurious correlation. That is, in order to overestimate the effects, categories in which consumers have stronger preferences over quality would have to be categories for which search costs are higher. Finally, due to the fact that we are mostly using *within-trip* variation to identify the effect of search time in this specification, one might wonder why search costs should vary at all over the course of the trip. Although a more thorough discussion is outside of the scope of this paper, we note that we observe systematically shorter search spells toward the end of most trips, possibly suggesting consumers might be less willing to process information and engage in search.

³¹Note that the number of observations for this robustness check varies slightly relative to the baseline IV regression, because we drop consumers for which only one item pickup is recorded when we include the fixed effects. We re-estimated the baseline model without the single-item trips (not reported) and find that the change in the sample size does not affect our results. For the same reason, the results in column (6) are slightly different from the ones reported for the “speed one minute before pickup” instrument in the text in section (4).

4.2 Price Distribution and Expectations

A model primitive that has a key influence on search behavior is the price distribution $F(p)$. We already discussed endogeneity concerns that arise from the fact that category-specific price distributions vary over time due to the fact that different products go on promotion at different points in time. We now turn to two more issues related to the price distribution. First, we consider the effect of consumers having biased expectations about the price distribution. Second, we investigate the consequences of consumers having information about daily prices, in particular promotions, before engaging in search. The latter is likely to arise in our setting due to the presence of feature advertising and in-store displays that provide price information to the consumer before she arrives at the shelf and starts searching.

4.2.1 Incorrect Expectations

A dimension in which consumers' behavior might differ from the stylized model is in the way they form expectations about prices. As in any search model, expectations play a crucial role because they determine the marginal benefit of searching and therefore the optimal amount of search activity.³² In our search model, a deviation from rational expectations can be captured by the fact that the optimal stopping rule would be based on an incorrect price distribution. In other words, the optimal price threshold λ would solve

$$\int_{\underline{p}}^{\lambda} (\lambda - p) d\tilde{F}(p) = c_{product}, \quad (8)$$

where $\tilde{F}(p)$ represents the price distribution used to form expectations. In the case of non-rational expectations, $\tilde{F}(p)$ will be different from the actual price distribution $F(p)$. Note that when the consumer engages in search, prices are still drawn from the true price distribution $F(p)$; however, the stopping threshold might differ from that of a rational consumer. $\tilde{F}(p)$ therefore only affects search time and price via its impact on λ :

$$\begin{aligned} E(p) &= \frac{1}{F(\lambda(\tilde{F}(p)))} \int_{\underline{p}}^{\lambda(\tilde{F}(p))} p dF(p) \\ E(SearchTime) &= \frac{TimePerSearch}{F(\lambda(\tilde{F}(p)))}. \end{aligned}$$

It is easy to see that more pessimistic expectations will lead to shorter search spells as well as a higher expected price paid. The negative correlation between search time and price that our estimation captures could therefore be in part due to heterogeneity in expectations across consumers. However, this variation is in fact not problematic for the interpretation of our estimate. We reiterate that expectations only influence search time and price paid through their influence on the stopping threshold λ . Moreover, the behavior of an overly optimistic

³²In virtually all structural models of search, consumers are assumed to know the true price distribution. Indeed, imposing the expectation process is usually necessary for identification in any dynamic model, including models of search.

consumer is observationally equivalent to the search behavior of a rational consumer with lower search costs. In other words, offsetting an increase in a consumer’s marginal benefit (due to over-optimism for instance) with an increase in the marginal cost in such a way that the stopping threshold remains unchanged is always possible. The offsetting increase in search costs would lead to the same outcome in terms of expected price and search duration, because price expectations and search costs only influence both price paid and search time via λ . For our estimation, whether movement in λ originates from variation in search costs or expectations does not matter.³³ More generally, any factor that influences the search process via equation (2), that is, by altering the stopping threshold, but not equations (4) and (3) is unproblematic for our estimation strategy. On the contrary, variation in the stopping threshold (for whatever reason) is precisely what we want to capture.

4.2.2 Information obtained before searching

Prices at the daily level are likely to be at least partially observed by some set of consumers due to feature advertising and in-store displays. Prior knowledge about prices affects behavior in two ways. Consumers with prior knowledge about daily prices will base their expectations on this information, whereas other consumers form expectations based on the distribution of prices over time and across products. This issue is very similar to the case of consumers having biased expectations. As discussed above, any type of variation in expectation formation does not cause any problems in terms of inference.

Apart from promotional activity having an impact on the *set* of prices being available and consumers’ expectations, it could also affect the probability with which a particular price is drawn. This issue is specific to our setup because all product prices are visually ”accessible” on the shelf. Promotions might therefore provide visual cues that draw the consumer’s attention to the promoted product, either because the consumer knows about the promotion and specifically tries to find the particular product or because promotional signs on the shelf capture her attention. Formally, such an effect would be captured by a shift in the CDF from which prices are drawn, which would now assign more probability weight to products that are promoted on the particular day.³⁴ This type of effect would lead to a negative correlation of promotional activity with search time, similar to the effect of variation in $F(p)$ over time discussed earlier.

Our instrument is valid as long as prior knowledge of prices does not alter the consumer’s walking speed or basket size. Especially with regards to speed, we think prior price knowledge in individual product categories is unlikely to affect our instruments.

³³Note that if any such variation in expectations is present in the data, our IV-strategy will most likely not make use of it. Consumers’ category-specific price expectations seem unlikely to be correlated with their walking speed or basket size.

³⁴In a pure price-search model, the probabilities of drawing a particular price are typically uniformly distributed across products.

4.3 Differences in search-efficiency

The final model primitive whose influence on our analysis we need to look at is *TimePerSearch*, the efficiency of the consumer’s search process. Variation likely occurs across consumers in how much time they need in order to resolve uncertainty regarding a specific number of options. The first-order effect of a decrease in *TimePerSearch* is that it lowers the consumer’s search cost and therefore leads to a lower stopping threshold λ . In other words, consumers who search more efficiently are willing to wait for a lower price draw because it is less costly for them to evaluate additional options in the search process. Search efficiency only affects price via this channel. However, the impact on search time is more complicated. On the one hand, search time will be longer due to the fact that a more efficient consumer is pickier, that is, has a lower λ . At the same time, search time is lower simply because evaluating an additional option takes less time, which we can easily see from equation (4), where *TimePerSearch* enters in the numerator and λ (which is a function of *TimePerSearch*) in the denominator. The consequences of variation in search efficiency for our estimation is similar in nature to a measurement-error problem. Ideally, we would like to measure variation in the extent of search activity in terms of the number of options evaluated, but we only observe search effort in real-time. The total search duration can be decomposed into two components: the number of options evaluated and the time it takes to evaluate each option. The former has an impact on price paid, but the latter does not. Variation in search time due to differences in search efficiency therefore cause attenuation bias in our estimate.

Because search efficiency is a latent concept, assessing how much this issue could affect estimation is hard. We are less sure, in this case, that our speed instrument is able to purge the problematic variation in search efficiency. Conceivably, consumers who are less efficient when searching also generally walk at a lower speed. However, *TimePerSearch* is unlikely to be correlated with the number of purchased items and the basket dummy, both of which we use as alternative instruments. Our results, as shown in Table (2) as well as Table (B2) in the appendix, are robust to using those instruments instead of walking speed.

5 Effect Magnitude

We find returns from searching that are fairly large: roughly \$2.10 per minute. However, because our measure of search time is distributed with a mean of 10 seconds and a standard deviation of 8 seconds, a minute constitutes a strong linear extrapolation relative to the typical search time. In this section, we provide some guidance on how to interpret the magnitude of the effect.

By the nature of the search problem, the relationship between search time and price paid is a non-linear one. Figure (4) illustrates this relationship when varying consumers’ search costs. The graph traces out how lowering search costs leads to a lower stopping threshold (λ), which in turn increases expected search time and decreases expected price (see equations (2) to (4)). As the graph shows, extensions in search time from a lower level are associated with larger gains in terms of finding a lower price. Moreover, the potential gains within a category

are bounded from below by the minimum price within the distribution. Our linear estimate allows us to recover the slope of the curve in Figure (4) for the average consumer in the sample, the magnitude of which depends on whether consumers in our data search relatively little (represented by the red scatter plot) or a lot (the blue scatter plot). In the latter case, the average consumer realizes more of the potential gains from search, and the incremental benefit at the margin is smaller. Due to the local nature of the effect and non-linear shape of the relationship, we have to be careful not to extrapolate out linearly “too far.”³⁵

With this caveat in mind, we use some back-of-the-envelope calculations to compute how large the gains from search can be within a given trip. Extending search time by one standard deviation, that is, by 8 seconds, lowers price by 28 cents. The average consumer purchases from seven categories on a typical trip and could therefore save about \$1.90 in total when extending search time by one standard deviation in each product category. These savings constitute roughly 7% of the average total shopping basket size of \$27. Another way to quantify potential savings from search is to put them into the broader context of the total time budget allocated to the shopping trip rather than just the time spent searching. Consumers spent on average 23 minutes in the store and only about 70 seconds, that is, 5% of their trip, searching. Extending search time by one standard deviation in each category, that is, by 56 seconds, corresponds to a 4% increase in total shopping time and lowers expenditure by \$1.90. Relative to the average trip-level expenditure of \$27, this amount of savings translates into an elasticity of expenditure with respect to shopping time of about -1.7 at the trip level.

6 In-store Search and Product Location

In this section, we explore how the estimates of search benefits derived above can be used to inform product-location decisions as well as pricing as a function of product location. We think in-store product-location choice is a natural area to explore using a model of consumer search. In fact, if consumers did have perfect information about product characteristics and price, production location should matter very little. If, instead, consumers have to engage in costly search, the store environment can be a tool to influence this process. Optimal store design is a complex problem and in its entirety is outside of the scope of this paper. Nevertheless, we think our estimates can shed light on some aspects of product placement. Furthermore, this area is relatively under-researched. Drèze, Hoch, and Purk (1994) is one of the few empirical papers we are aware of on the effect of store layout.³⁶ The advent of in-store tracking data such as the path-tracking data used here and in Hui, Fader, and Bradlow (2009), Hui, Bradlow, and Fader (2009a), Hui, Bradlow, and Fader (2009b) as well as video tracking in Jain, Misra, and Rudi (2014) should allow researchers to re-visit those important questions.

³⁵We did try to estimate the curvature of the relationship presented in Figures (1) and (4). When also including search time squared, we find a negative coefficient on the linear and a positive one on the squared term, which is consistent with the graph. However, neither coefficient is significantly different from zero.

³⁶Another example is Hui, Inman, Huang, and Suher (2013), who run a counterfactual in which they alter the location of certain product categories and analyze the effect of unplanned spending.

To gauge the potential for the store layout to influence search behavior, we regress search time on a set of dummies for different areas of the store. More specifically, we partition the store into 31 regions, which include aisles in the middle of the store as well as wall segments (of similar length as the aisles) along the perimeter of the store. Furthermore, we also partition each aisle in five roughly equally spaced segments. We use a set of dummies for the broad regions as well as a separate set for the within-aisle segments. Results from this regression are reported in Table (4). Due to the large number of regions, we do not report the full set of coefficients, but only some aggregate statistics on the coefficient values. We find differences in search time across regions of up to 9.8 seconds as well as a maximum difference of 5.9 seconds between segments of an aisle. Search-time tends to be longer in the middle / bottom part of an aisle as well as in aisles toward the center of the store. Also, aisles generally see longer search-spells than walls at the perimeter of the store.

The question is whether we can attribute these difference in search time to the physical location. Product categories are, of course, not located randomly throughout the store, and we might thus pick up across-category differences with the location dummies. Second, measurement error in the search time metric might vary across different areas of the store, most likely because the probability of consumers leaving their carts behind might be higher in some areas of the store than in others. We address the first issue by including a set of category fixed effects alongside the location dummies. The two sets of dummies can be identified due to the fact that many categories are stocked in different areas of the store. Note that this control is not ideal, because different locations for the same category are usually characterized by differences in product assortment. The first best would be to use panel data with changes in category location, but unfortunately our data do not contain such variation.³⁷ To address the issue of measurement error, we re-run the regression using only trips during which the consumer used a basket rather than a shopping cart. Using this sub-sample mitigates concerns of locational differences in measurement error, because consumers are presumably less likely to leave their basket behind relative to a cart. Regressions using category fixed effects as additional controls as well as results for a restricted sample of trips with baskets are reported in columns (2) and (3) of Table (4). Both specifications yield similar results to the specification in column (1).³⁸ We note that both robustness checks have their limitations and we see the analysis in this section as more exploratory and suggestive. Ideally, one would randomly vary category locations over time and study the impact on search and purchase behavior. We leave such analysis to future research.

What do the differences in search behavior reported in Table (4) imply in terms of purchase outcomes? Using the regression results from our baseline regression, moving a product from the lowest search time to the highest search time location, both across and within aisles, implies a difference of almost 16 seconds, which leads to a difference of $-0.54 = -0.0344$

³⁷We also tried including the number of UPCs within each category/location pair and found the results are qualitatively similar.

³⁸Note that the specification using category fixed effects yields much larger standard errors. However, if anything, results from this specification indicate an even larger difference in search time across locations (relative to our baseline specification in column (1)).

* 15.7 in terms of lower price, which corresponds to roughly 17% of the median potential savings at the category level (see Table(1)). Or, using the specification in column (6) of Table (2), a 13-percentage-point increase in the probability of purchasing a promoted product.³⁹ Unfortunately, our specification does not allow us to translate the search time differences easily into differences in terms of price elasticities. Nevertheless, the results suggest that locational factors do have a large impact on price sensitivity. This finding is relevant for manufacturers who pay slotting allowances to place their products in certain locations inside the store. Based on search time differences, some locations do engender closer competition with other brands due to consumers engaging in more search. Similarly, pricing decisions should arguably be a function of product location as well: in high-search locations, running a promotion will be more effective than in areas of the store where consumers' search effort is lower.

7 Conclusion

We estimate the effect of search intensity on the price a consumer pays within a particular category, using data from RFID tags on supermarket shopping carts. Recording search in a physical store environment is generally challenging, and even our detailed data are only able to capture total search time but not which options the consumer evaluated. The technology does, however, have the advantage of not interfering in any way with the consumer's natural shopping experience, and might be the best possible way to gain insights into consumer search in a brick-and-mortar store. To the best of our knowledge, this paper, together with Jain, Misra, and Rudi (2014), is the first to use direct data on search effort to analyze consumer search within a brick-and-mortar environment.

We employ a reduced-form approach to estimate the effect of search intensity on the price a consumer pays within a particular category. We find that an additional minute of search lowers expenditure by about \$2.10. The gains from search are substantial: increasing search time by one standard deviation in each purchased category leads to a 7% reduction in total shopping-basket expenditure. This result is robust to a host of sensitivity checks that deal with possible confounds such as variation in prices over time, measurement error, and correlation between price sensitivity and search costs. Due to the limited amount of observations per category in the data, our evidence comes from regressions that are pooled across categories. Going forward, with path data over a longer time horizon for only one category, modeling the search process in more detail (possibly by means of a structural model) should be possible. In particular, our approach only looks at the effect on price paid and does not directly analyze the role of other product characteristics. We are therefore not able to make any statements about the effect of search on consumer utility. However, we believe the effect of search time on

³⁹A deeper question is why certain locations within the store see more search activity than others. We find in our data that consumers' search time varies systematically over the course of a trip, with search time being lower at the very beginning and end of a trip, suggesting that awareness and cognitive ability to process information might vary over the course of the trip. We conjecture that at least part of the locational variation is due to variation in when consumers tend to reach a particular location during their shopping trip. A more detailed exploration is outside of the scope of this paper.

price is a dimension of the search process that is particularly relevant for informing optimal supply-side behavior.

Our findings imply that, due to the limited amount of search, the use of marketing tools such as feature advertising and in-store displays can be very effective. Furthermore, firm behavior that influences consumer search interacts in an interesting way with pricing decisions. Because more search makes finding a lower price or promoted product more likely, firms have an incentive to encourage search when running a promotion.⁴⁰ Finally, we find that product location can greatly influence consumer behavior due to differences in search intensity in different areas of the store. Generally, we think the type of data and approach presented here opens the door to studying issues of product location and store design in more detail.

⁴⁰The data and empirical approach could also be used to study seasonal variation in search behavior which (as posited by Haviv (2013)) might be a source of counter-cyclical pricing.

References

- ANDREWS, R. L., AND T. C. SRINIVASAN (1995): “Studying Consideration Effects in Empirical Choice Models Using Scanner Panel Data,” *Journal of Marketing Research*, 32(1), 30–41.
- BRONNENBERG, B. J., M. W. KRUGER, AND C. F. MELA (2008): “The IRI Marketing Data Set,” *Marketing Science*, 27(4), 745–748.
- BRONNENBERG, B. J., AND W. R. VANHONACKER (1996): “Limited Choice Sets, Local Price Response and Implied Measures of Price Competition,” *Journal of Marketing Research*, 33(2), 163–173.
- CHEN, Y., AND S. YAO (2014): “Search with Refinement,” unpublished manuscript.
- COBB, C. J., AND W. D. HOYER (1985): “Direct Observation of Search Behavior in the Purchase of Two Nondurable Products,” *Psychology and Marketing*, 2(3), 161–179.
- DE LOS SANTOS, B. I., A. HORTACSU, AND M. WILDENBEEST (2013): “Testing Models of Consumer Search using Data on Web Browsing and Purchasing Behavior,” *American Economic Review*.
- DICKSON, P. R., AND A. G. SAWYER (1990): “The Price Knowledge and Search of Supermarket Shoppers,” *The Journal of Marketing*, 54(3), 42–53.
- DRAGANSKA, M., AND D. KLAPPER (2011): “Choice Set Heterogeneity and the Role of Advertising: An Analysis with Micro and Macro Data,” *Journal of Marketing Research*, 48, 653–669.
- DRÈZE, X., S. J. HOCH, AND M. E. PURK (1994): “Shelf Management and Space Elasticity,” *Journal of Retailing*, 70(4), 301–326.
- HAVIV, A. (2013): “Does Purchase Without Search Explain Counter-Cyclic Pricing?,” unpublished manuscript.
- HONKA, E. (2014): “Quantifying Search and Switching Costs in the U.S. Auto Insurance Industry,” *RAND Journal of Economics*, forthcoming.
- HOYER, W. D. (1984): “An Examination of Consumer Decision Making for a Common Repeat Purchase Product,” *The Journal of Consumer Research*, 11(3), 822–829.
- HUI, S. K., E. T. BRADLOW, AND P. S. FADER (2009a): “Testing Behavioral Hypotheses Using an Integrated Model of Grocery Store Shopping Path and Purchase Behavior,” *Journal of Consumer Research*, 36(3), 478–493.
- (2009b): “The Traveling Salesman Goes Shopping: The Systematic Deviations of Grocery Paths from TSP Optimality,” *Marketing Science*, 28(3), 566–572.

- HUI, S. K., P. S. FADER, AND E. T. BRADLOW (2009): “Path Data in Marketing: An Integrative Framework and Prospectus for Model Building,” *Marketing Science*, 28(2), 320–335.
- HUI, S. K., Y. HUANG, J. SUHER, AND J. INMAN (2013): “Deconstructing the First Moment of Truth: Understanding Unplanned Consideration and Purchase Conversion Using In-Store Video Tracking,” *Journal of Marketing Research*, 50(4), 445–462.
- HUI, S. K., J. INMAN, Y. HUANG, AND J. SUHER (2013): “Estimating the Effect of Travel Distance on Unplanned Spending: Applications to Mobile Promotion Strategies,” *Journal of Marketing*, 77(2), 1–16.
- JAIN, A., S. MISRA, AND N. RUDI (2014): “On Search, Sales Assistance and Purchase Decisions: An Analysis Using Retail Video Data,” unpublished manuscript.
- KIM, J. B., P. ALBUQUERQUE, AND B. J. BRONNENBERG (2010): “Online Demand under Limited Consumer Search,” *Marketing Science*, 29(6), 1001–1023.
- KOULAYEV, S. (2013): “Estimating Demand in Online Search Markets, with Application to Hotel Bookings,” *RAND Journal of Economics*, forthcoming.
- MCCALL, J. J. (1970): “Economics of Information and Job Search,” *The Quarterly Journal of Economics*, 84(1), 221–242.
- MEHTA, N., S. RAJIV, AND K. SRINIVASAN (2003): “Price Uncertainty and Consumer Search: A Structural Model of Consideration Set Formation,” *Marketing Science*, 22(1), 58–84.
- ROBERTS, J. H., AND J. M. LATTIN (1991): “Development and Testing of a Model of Consideration Set Composition,” *Journal of Marketing Research*, 28(4), 429–440.
- SEILER, S. (2013): “The Impact of Search Costs on Consumer Behavior: A Dynamic Approach,” *Quantitative Marketing and Economics*, 11(2), 155–203.
- SORENSEN, H. (2003): “The Science of Shopping,” *Marketing Research*, 15(3), 31–35.
- STÜTTGEN, P., P. BOATWRIGHT, AND R. T. MONROE (2012): “A Satisficing Choice Model,” *Marketing Science*, 31, 878–899.
- WEITZMAN, M. L. (1979): “Optimal Search for the Best Alternative,” *Econometrica*, 47(3), 641–654.

	Mean	S.D.	P25	Median	P75
<u>TRIP-LEVEL VARIABLES</u>					
Speed (Feet per Second)	2.21	0.31	2.00	2.18	2.39
Number of Items Purchased	4.24	3.67	2	3	6
Basket Dummy	0.15				
Trip Duration (Minutes)	23.65	17.18	11.45	18.81	30.52
Trip Distance (100 Feet)	29.83	19.77	15.84	24.54	37.92
<u>PICK-UP LEVEL VARIABLES</u>					
Search Time	10.25	8.38	4.80	7.84	12.67
Speed 60 Seconds Before Pick-up	2.34	0.77	1.82	2.31	2.81
Price Paid	3.33	3.00	2.50	1.59	3.99
<u>PRICE SAVINGS</u>					
Absolute Difference between Daily Min and Max Price	4.21	3.52	1.78	3.21	5.26
Percentage Difference between Daily Min and Max Price	0.66	0.23	0.53	0.71	0.84
Fraction of UPCs Promoted on a Specific Day	0.3	0.17	0.17	0.29	0.42
Fraction of UPCs Promoted during the Sample Period	0.58	0.32	0.41	0.62	0.83

Table 1: **Descriptive Statistics: Prices.** The unit of observation is a trip in the top panel and an item pickup in the middle panel. Our sample consists of 13,112 trips and 34,109 pickups. The unit of observation in the bottom panel is a category. Our data contain 153 categories.

	(1)	(2)	(3)	(4)	(5)	(6)
Type of Regression	OLS	IV 1st Stage	IV 2nd Stage	IV 1st Stage	IV 2nd Stage	IV 2nd Stage
Dependent Variable	Price	Search Time	Price	Search Time	Price	Promotion Dummy
Search-Time	-0.0071*** (0.0016)		-0.0344*** (0.0127)		-0.0528*** (0.0158)	0.0082* (0.0046)
Speed		-4.763*** (0.191)				
Number of Purchased Items Basket Dummy				0.181*** (0.018)		
				-0.588*** (0.164)		
First-stage F-Stat		619.62		63.44		385.61
Category FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	34,109	34,109	34,109	34,109	34,109	23,444
Trips	13,112	13,112	13,112	13,112	13,112	11,031
Consumers	8,318	8,318	8,318	8,318	8,318	7,247

Table 2: **Baseline OLS and IV regressions.** The unit of observation is an item pickup. Standard errors are clustered at the consumer level.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Type of Regression	IV 2nd Stage	IV 1st Stage	IV 2nd Stage	IV 1st Stage	IV 2nd Stage	IV 2nd Stage	IV 2nd Stage
Sample	Full Sample	Repeat Customers	Repeat Customers	Repeat Customers	Repeat Customers	Consumer with >1 Pick-up	Consumer with >1 Pick-up
Dependent Variable	Price	Search Time	Price	Search Time	Price	Price	Price
Search Time	-0.0344*** (0.0127)		-0.0461*** (0.0144)		-0.0436*** (0.0141)	-0.0247*** (0.0054)	-0.0178*** (0.0067)
Trip-Level Speed		-4.816*** (0.243)		-4.815*** (0.243)			
Average Price Percentile				-0.257 (0.504)	2.2342*** (0.1450)		
First Stage F-Stat	619.62	392.33		392.76		1237.45	761.01
Category FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Consumer FEs	No	No	No	No	No	No	Yes
Observations	34,109	25,166	25,166	25,166	25,166	32,164	32,164
Trips	13,112	8,740	8,740	8,740	8,740	11,167	11,167
Consumers	8,318	6,562	6,562	6,562	6,562	6,373	6,373

Table 3: **Robustness Checks: Price Sensitivity Controls and Trip Fixed-Effect Regressions.** The unit of observation is an item pickup. Standard errors are clustered at the consumer level. Sample size changes due to the fact that we exclude trips with only one pickup when including trip fixed effects and price sensitivity is only defined for consumers with repeat observations in our data.

	(1)	(2)	(3)
Dependent Variable	Search Time	Search Time	Search Time
Sample	Full Sample	Full Sample	Trips with Baskets
<u> AISLE SEGMENTS </u>			
Top	Omitted Category	Omitted Category	Omitted Category
Middle-Top	2.160*** (0.218)	2.064*** (0.263)	1.803*** (0.542)
Middle	5.006*** (0.249)	5.009*** (0.308)	4.576*** (0.780)
Middle-Bottom	5.946*** (0.286)	5.986*** (0.330)	5.035*** (0.682)
Bottom	3.010*** (0.256)	3.552*** (0.330)	2.068*** (0.630)
<u> STORE REGIONS </u>			
Difference Min - Max Region FE Coefficient	9.753** (4.704)	20.961 (14.317)	7.214*** (0.765)
Difference Top2 - Bottom2 Region FE Coefficient	8.792*** (2.413)	12.895* (7.180)	6.866*** (0.693)
Difference Top3 - Bottom3 Region FE Coefficient	7.619*** (1.594)	9.964** (4.829)	6.274*** (1.445)
Product Category FEs	No	Yes	No
Observations	34,109	34,109	4,005
Number of Store Regions	31	31	31

Table 4: **The Effect of Product Location on Search Time.** The unit of observations is an item pickup. Standard errors are clustered at the consumer level. A full set of store region dummies are included in all specifications. The “store region” panel presents hypothesis tests for differences between averages of groups of fixed-effect coefficients at the top and bottom of the distribution of coefficient values in each specification.




Figure 1: **Data-Structure.** The picture illustrates a consumer traversing an aisle. Consumer location within the aisle is recorded on a grid of traffic points. Products are located at specific locations on the shelf, which are coded up as a grid of product points. Product points are matched to nearby traffic points, allowing us to measure how long a consumer remained near the product when picking it up. The dashed black line denotes the consumer's path when traversing the aisle.




Figure 2: **Search-Time Histogram**




Figure 3: **Relationship between Search Time and Price Paid when Varying Search Costs.** The picture illustrates the relationship between expected price paid and search time for varying levels of search costs. The relationship is non-linear, with extensions in search time from a lower level being associated with larger gains. Moreover, the potential gains within a category are bounded by the lower bound of the price distribution.




Figure 4: **Estimated Local Average Treatment Effect.** The picture illustrates the local nature of the estimated search benefit. The magnitude of our estimate depends on whether consumers in our data search relatively little (red scatter plot) or a lot (blue scatter plot). In the latter case, the average consumer realizes more of the potential gains from search, and the incremental benefit at the margin is therefore smaller.

A Appendix

A.1 Linking Sales and Path Data

One of the important features of our data set is the linkage of sales to trip records. As part of the RFID tracking process, the data report when the consumer arrives at the checkout. Independently, the sales data also have a time stamp for each shopper’s transaction at the checkout. Comparing the time stamp of a particular path with the sales data allows us to define a set of ”candidate” checkout product baskets that occurred at a similar point in time.⁴¹ Matching which trip goes with which specific transaction involves considering the physical location (i.e., longitude = x and latitude = y relative to the store map) of all the UPCs in each candidate basket. Based on how many of those locations lay on the path we are trying to match, a score is created for the baskets and the highest-scoring one is matched to the path.⁴² The matches do not necessarily yield a perfect score, because consumers might occasionally leave the cart and pick up an item. Therefore, we might not see the path of the consumer going past a specific item, even if the item was in her matched purchase basket. In this case, no information on search time will be available for the particular item.

⁴¹The path data time stamp that records the arrival at the checkout can be noisy because the consumer will be stationary when standing in line at the cashier. Therefore, checkout baskets within a certain time window after the consumer became stationary in the check-out area qualify as possible matches.

⁴²The data provider did not disclose the precise algorithm to us.

B Appendix: Tables

	(1)	(2)	(3)
Dependent Variable	Search-Time	Price Paid	Log(Price Paid)
Share of Promoted UPCs within the Category	-1.036* (0.536)	-0.551*** (0.146)	-0.182*** (0.036)
Average of the Dependent Variable	10.25	3.33	0.925

Category FEs	Yes	Yes	Yes
Day FEs	Yes	Yes	Yes
Observations	34,109	34,109	34,109

Table B1: **The Effect of Category-level Pricing on Search.** The unit of observations is an item pickup.

1st STAGE	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(DV: Search Time)							
Speed	-4.763*** (0.191)			-4.374*** (0.204)			-4.431*** (0.320)
Number of Purchased Items Basket Dummy		0.185*** (0.018)		0.135*** (0.016)			0.133*** (0.020)
Trip Duration (Units: Minutes)			-0.935*** (0.158)	0.504*** (0.166)	0.050*** (0.004)		-0.006 (0.026)
Trip Length (Units: 100 Feet)						0.036*** (0.003)	0.006 (0.022)
F-Stat	619.62	108.32	34.95	198.98	173.72	119.91	134.42
<hr/>							
2nd STAGE							
Coefficient on Search Time (DV: Price)	-0.0344*** (0.0127)	-0.0499*** (0.0159)	-0.0998* (0.0535)	-0.0376*** (0.0112)	-0.0584*** (0.0171)	-0.0645*** (0.0205)	-0.0378*** (0.0112)
<hr/>							
Category FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	34,109	34,109	34,109	34,109	34,109	34,109	34,109
Trips	13,112	13,112	13,112	13,112	13,112	13,112	13,112
Consumers	8,318	8,318	8,318	8,318	8,318	8,318	8,318

Table B2: **Robustness Check: Alternative Instruments.** The unit of observations is an item pickup. Standard errors are clustered at the consumer level. All specifications are identical except for a change in the instrument(s).

CENTRE FOR ECONOMIC PERFORMANCE
Recent Discussion Papers

1295	Matt Dickson Paul Gregg Harriet Robinson	Early, Late or Never? When Does Parental Education Impact Child Outcomes?
1294	Colin Hottman Stephen J. Redding David E. Weinstein	What is 'Firm Heterogeneity' in Trade Models? The Role of Quality, Scope, Markups and Cost
1293	Warn N. Lekfuangfu Francesca Cornaglia Nattavudh Powdthavee Nele Warrinnier	Locus of Control and Its Intergenerational Implications for Early Childhood Skill Formation
1292	Gianmarco Ottaviano	Immigration, Diversity and the Labour Market Outcomes of Native Workers: Some Recent Developments
1291	John Morrow	Benford's Law, Families of Distributions and a Test Basis
1290	Andreas Georgiadis Alan Manning	The Volatility of Earnings: Evidence from High-Frequency Firm-Level Data
1289	Francesco Caselli	The Latin American Efficiency Gap
1288	Thomas Sampson	Dynamic Selection: An Idea Flows Theory of Entry, Trade and Growth
1287	Fabrice Defever Alejandro Riaño	Gone for Good? Subsidies with Export Share Requirements in China: 2002-2013
1286	Paul Dolan Matteo M. Galizzi	Because I'm Worth It: A Lab-Field Experiment on the Spillover Effects of Incentives in Health
1285	Swati Dhingra	Reconciling Observed Tariffs and the Median Voter Model
1284	Brian Bell Anna Bindler Stephen Machin	Crime Scars: Recessions and the Making of Career Criminals

1283	Alex Bryson Arnaud Chevalier	What Happens When Employers are Free to Discriminate? Evidence from the English Barclays Premier Fantasy Football League
1282	Christos Genakos Tommaso Valletti	Evaluating a Decade of Mobile Termination Rate Regulation
1281	Hannes Schwandt	Wealth Shocks and Health Outcomes: Evidence from Stock Market Fluctuations
1280	Stephan E. Maurer Andrei V. Potlogea	Fueling the Gender Gap? Oil and Women's Labor and Marriage Market Outcomes
1279	Petri Böckerman Alex Bryson Jutta Viinikainen Christian Hakulinen Laura Pulkki-Raback Olli Raitakari	Biomarkers and Long-term Labour Market Outcomes: The Case of Creatine
1278	Thiemo Fetzer	Fracking Growth
1277	Stephen J. Redding Matthew A. Turner	Transportation Costs and the Spatial Organization of Economic Activity
1276	Stephen Hansen Michael McMahon Andrea Prat	Transparency and Deliberation within the FOMC: A Computational Linguistics Approach
1275	Aleksi Aaltonen Stephan Seiler	Quantifying Spillovers in Open Source Content Production: Evidence from Wikipedia
1274	Chiara Criscuolo Peter N. Gal Carlo Menon	The Dynamics of Employment Growth: New Evidence from 18 Countries
1273	Pablo Fajgelbaum Stephen J. Redding	External Integration, Structural Transformation and Economic Development: Evidence From Argentina 1870-1914

The Centre for Economic Performance Publications Unit
Tel 020 7955 7673 Fax 020 7404 0612
Email info@cep.lse.ac.uk Web site <http://cep.lse.ac.uk>