Sales Tax, E-commerce, and Amazon's Fulfillment Center Network*

Jean-Francois Houde
The Wharton School
University of Pennsylvania

Peter Newberry
Department of Economics
The Pennsylvania State University

Katja Seim
The Wharton School
University of Pennsylvania

February 29, 2016

Abstract

We estimate the cost savings associated with the expansion of Amazon's fulfillment center network from 2006-2018. We first demonstrate that, in placing a fulfillment center in a new state, Amazon faces a trade-off between the revenue considerations from exposing local customers to sales tax and the cost savings from reducing the shipping distance to those customers. Using detailed data on online transactions, we estimate a model of demand for retail goods and show that consumers' online shopping is sensitive to being charged sales tax. We then use the demand estimates and the spatial distribution of demand relative to Amazon's fulfillment centers to produce predicted revenues and shipping distances under the observed fulfillment center roll-out and under counterfactual roll-outs over this time period. Using a moment inequalities approach, we infer the cost savings associated with being closer to customers that render the observed network roll-out optimal. We find that Amazon saves between \$0.40 and \$1.30 for every 100 miles of \$100 of goods shipped.

Keywords: e-commerce, sales tax, Amazon, distribution, logistics, shipping
JEL Codes: H71, L11, L81

*All correspondence may be addressed to the authors via e-mail at `houde@wharton.upenn.edu`, `pwnewberry@psu.edu`, or `kseim@wharton.upenn.edu`.

1 Introduction

Online retail has grown substantially over the last decade and now makes up nearly 7% of all retail sales.[1] Amazon.com is a key contributor of this growth, with net sales increasing from $8.5 billion in 2005 to $89 billion in 2014.[2] At the same time, Amazon's rise has been accompanied by increasing concentration in online retail, despite initial expectations that the internet would facilitate highly competitive markets where buyers have the option to compare prices and buy from many different platforms and sellers. As the online retail market matures, this no longer appears to be the case; the market today features a few powerful firms. Figure 1 uses data from the comScore Web Behavior data base to illustrate this trend, focusing on product categories where Amazon.com competes. Over the period 2006-2013, the Herfindahl-Hirshman Index increased seven to ten-fold, depending on whether market shares are based on consumer spending or on transaction counts. In this paper, we explore one possible explanation for this increasing concentration: the competitive advantage that comes with cost savings of a large-scale distribution network.

We investigate the role of economies of scale in distribution using Amazon as a case study. As Amazon's revenue has grown, so has its network of distribution centers (called Fulfillment Centers, denoted as FCs going forward). The number of FCs has grown from 8 centrally located centers in 2006 to 46 FCs spread out across the US by the end of 2015.[3] Amazon has plans to further expand its network to over 90 facilities by 2018, including "Amazon Now" hubs that provide same-day delivery to local customers. The main goal of this paper is to investigate the primary drivers of these expansion decisions.

A larger distribution network has both positive and negative implications for firm profitability. A distinction of online markets lies in the fact that an online retailer is only required to charge its customers sales tax if the firm has a physical presence (or "nexus"), such as a warehouse, in the customer's state of residence.[4] As earlier literature (Einav et al. (2014),Baugh et al. (2014))

[1] `www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf`.

[2] `www.statista.com/statistics/266282/annual-net-revenue-of-amazoncom/`.

[3] These figures do not include Amazon fresh FCs or sortation centers. See section 2 for details on distribution center types.

[4] In these states, consumers are required to file a use tax return for all of their online purchases. In most cases, consumers do not do this.

shows, demand is sensitive to sales taxes as they raise the effective price of an online transaction. To limit the revenue implications of exposing customers to new sales tax liabilities, a supply-side response would thus be not to establish a presence in a high-tax state. This incentive would be particularly high in populous states, where in addition, the fixed costs from operating a warehouse are large. Such costs of building a warehouse in a new state are potentially offset by two benefits from proximity to the consumer: first, the firm may be able to make faster deliveries; consumers could derive independent value from such convenience. Second, by cutting the distance between the FC and the consumer, the firm may also be able to derive significant cost saving from shorter shipping routes.

Shipping comprises a significant share of Amazon's operating costs: the financial statements report the costs (net of shipping revenue) are between 3-5% of net sales. These costs totaled nearly $4.3 billion in 2014.[5] We posit that by expanding its network of FCs, Amazon saves on "outbound" shipping cost, or the cost to get the package from the FC to the consumer, by reducing the initial distance a package travels from the FC to a contracted shipper's (i.e., UPS, FedEx, US Postal Service) local sorting center, even if the delivery distance from the sorting center to the customer is unaffected. This reduces the shippers' cost, providing Amazon with bargaining power in negotiations with shippers over the per-package rates they charge Amazon. Expansion also affects the distance of "inbound" shipments from the supplier to Amazon's FC. However, because the suppliers are likely shipping goods to brick and mortar retail outlets all around the country, we predict that these costs do not change significantly as Amazon expands its network and that as a result, Amazon's wholesale cost is unaffected. Therefore, our goal is to estimate the relationship between the outbound shipping distance and outbound shipping cost.

We quantify empirically the trade-off between higher tax-inclusive prices and fixed cost, on the one hand, and additional convenience and shipping cost savings on the other. To accomplish this, we use detailed information on internet purchase behavior at Amazon and other online competitors by a large representative consumer sample, combined with detailed information on the location and characteristics of Amazon's FCs over time. We combine these data with a revealed preference

[5] See annual reports at `http://phx.corporate-ir.net/phoenix.zhtml?c=97664&p=irol-reportsannual`.

approach similar to that of Holmes (2011). Specifically, we use the fact that entrance of a FC into some states results in lost revenue due to new sales tax liabilities. Therefore, assuming profit maximizing behavior on the part of Amazon implies that entry into a state must also result in either increased willingness-to-pay due to faster deliveries or a reduction of shipping costs. This allows us to form a measure of cost saving by relating the combined change in revenue from tax and convenience effects to the reduction in shipping distance through moment inequalities.

The first step in this approach is to estimate a model of shopping mode demand in order to pin down consumer tax sensitivity and response to shipping times. We specify a CES demand model in which a representative consumer in a county chooses her yearly expenditures on four different modes of shopping: (1) Amazon.com, (2) a taxed online competitor (e.g., walmart.com), (3) a non-taxed online competitor (e.g., overstock.com), and (4) a taxed offline competitor (e.g., Wal-Mart). Importantly, the amount of sales tax charged to the consumer varies both across modes and across time. While some of this variation is a result of the expansion of the FC network, there is also rich variation in tax rates across counties within a state and within counties across time. Therefore, the covariation in the sales tax and expenditures across shopping modes and across time identifies consumers' sensitivity to sales tax.

We estimate the demand model using household level data from the comScore Web Behavior Database, which includes details about the online transactions of between 30,000 and 50,000 households from 2006-2013. We supplement these with survey data from Forrester Research, which provide us with information on demographics of consumers choosing not to buy anything online and with Census data on aggregate retail spending. Using these sources, we calculate the expenditures for a representative consumer in each county across the four modes of shopping in each year from 2006 to 2013. In addition to these expenditures, we observe the sales tax rate for each county and year, and are able to account for county-level consumer demographics and county-level measures of offline retail competition via data from the census. This model is very much in the same spirit of the one found in Einav et al. (2014).

The results of the demand estimation suggest that consumers are sensitive to taxes and that their tax elasticity is around -1.4. This is similar to the estimates from Einav et al. (2014), -1.8, and

Baugh et al. (2014), -1.5. The magnitude of this estimate implies that moving from a non-taxed regime to a taxed regime at the average sales tax rate of 6.5%, all else equal, results in reduction of expenditures on Amazon of around 9%. The reduction is even more pronounced for high tax states such as New York (12% reduction in sales) and Illinois (11.5% reduction in sales).[6] These examples are especially pertinent because both states have witnessed a tax regime shift in the recent years.[7] In general, the results suggest that the tax effect must be taken into consideration when Amazon, or other online firms, choose where and when to build their fulfillment center network.

We find little evidence that the entrance of a FC leads to an increase in demand for Amazon from increased convenience. The lack of effect of shipping time on demand is likely due to the fact that Amazon's shipping options and prices did not drastically change between 2006-2013. That is, if a customer wanted a package to arrive in two days, she could receive it in two days no matter where she lives in relation to a FC.[8]

These estimates imply that there are overall negative revenue effects of opening a fulfillment center in a taxed state. To quantify this effect, we compare Amazon's revenue from 2007 to 2018 under the observed fulfillment center network and tax laws to the revenue if the network and tax laws had remained as they were in 2006. We use our model and the planned locations of future FCs in order to predict revenue for Amazon from 2015 to 2018. We find that expansion has lead to a loss in revenue of nearly $9 billion, or 1.6% of total revenue.

We further explore the tradeoff between the revenue effects of charging sales tax to customers in highly populated states and the shipping distance to these customers. For example, Amazon's FCs in 2006 were an average of 252 miles from the closest highly populated metropolitan area, with 7 out their 8 FCs being in a different state than this population hub. If they would have separately

[6] This is calculated using the average sales tax rate across counties in these states as of January 1, 2014. See `http://taxfoundation.org/article/state-and-local-sales-tax-rates-2014`.

[7] In 2008, New York state passed a law that requires Amazon to charge sales tax to its customers, even though Amazon did not have a 'nexus' in the state. The above results suggest that this had a large impact on sales in New York post 2008. Additionally, Amazon has had a long battle with Illinois regarding sales tax, but finally agreed to charge its customers sales tax in the beginning of 2015. Again, the results suggest that Amazon's sales are negatively affected by this shift. However, Amazon has announced plans to build distribution facilities in Illinois to serve the greater Chicago area. See `http://www.chicagotribune.com/business/breaking/ct-amazon-sales-tax-illinois-0124-biz-2-20150123-story.html#page=1` for details.

[8] Of course this has changed with the arrival of same-day shipping. We plan to estimate this effect in future iterations of the paper by observing the entry of FCs on the outskirts of large metropolitan areas.

moved each FC to the population hub, they would have lost an average of $138 million in revenue due to the effect of charging sales tax.

Next, we quantify the shipping cost savings from the larger FC network size that offset these revenue declines by specifying a profit function where the variable cost of providing goods to consumers in a county is a function the shipping distance from the FC to that county. We assume that the observed network placement is optimal so that the profit under this structure must be at least as large as under any perturbation of the network. A perturbed network configuration affects revenue through the potentially differing sales taxes, fixed costs through the wages and rents paid in the counties where Amazon FCs are located, and variable costs through changes in the total shipping distance across customers. We compute perturbations similar to those of Holmes (2011): we swap the observed opening dates of two FCs.

The set of perturbations allows us to form moment inequalities from which we can estimate the effect of distance on variable cost using the methods in Pakes et al. (2015). We define instruments similar to that of Holmes (2011), but make a couple of important adjustments due to the nature of our problem. For one, grouping perturbations in the same manner as Holmes (2011) would result in a violation of the assumption of exogeneity of the instruments. We therefore define exogenous groupings based on a "first stage" regression of county level revenue on exogenous shifters (e.g., populations). Second, we adjust the groupings based on a smooth transformation, implying that we can use the full set of perturbations rather than arbitrarily dropping the ones which do not fall into one of the discrete groups.

Estimates suggest that Amazon saves between 0.4% and 1.3% of revenue per 100 miles of shipping distance. Relative to holding its network size fixed at the 2006 configuration, our estimates suggest that over 2007-2018, Amazon saves between $3 and $12 billion in shipping costs through its FC network expansion, evidence of significant economies of scale in distribution in online markets.

This paper is related to several strands of the literature. First, there are a number of papers that examine the response of consumers to sales taxes in offline markets. Agarwal et al. (2013) and Asplund et al. (2007) demonstrate that people substitute their shopping dollars across state and country borders when there exists differences in sales tax rates, while Agarwal et al. (2013)

finds that consumers respond to 'tax-holidays' by increasing their spending. Chetty et al. (2009) estimates how the saliency of tax rates affects consumer demand and finds that consumers reduce their purchasing when the taxes are included in the posted price.

There is a large and growing literature examining the effect of sales tax on online purchasing. In early papers, Goolsbee (2000a, 2000b) used variation in tax rates across municipalities and found that consumers were highly sensitive to sales tax. Both Alm and Melnik (2005) and Ballard and Lee (2007) find small but significant effects of sales tax on the decision of whether or not to shop online, while Scanlan (2007) finds that this sensitivity is heterogeneous across the level of tax-rates. Ellison and Ellison (2009), Smith and Brynjolfsson (2001), Anderson et al. (2010), and Goolsbee et al. (2010) all find that online shoppers are sensitive to sales tax, but do so focusing on a single product category.

Two papers that are similar to the demand estimation in the current paper in many respects are Einav et al. (2014) and Baugh et al. (2014). Einav et al. (2014) estimate the response to sales tax using eBay data, exploiting the fact that a buyer has the option to buy from an out-of-state seller who does not charge sales tax. They are also able to estimate a consumer's sensitivity to distance because they observe the location of both buyers and sellers. One limitation of their paper is that they use only one online outlet, eBay, to estimate tax-sensitivity. That is, they do not consider that eBay buyers can substitute to other online stores, which may or may not be taxed. Baugh et al. (2014) uses a differences-in-differences approach to estimate the effect of the 'Amazon tax', or changes in the laws that forced Amazon to charge sales tax in various states during 2013-2015. Similar to Einav et al. (2014), they do not consider substitution to other taxed or non-taxed online outlets. We thus contribute to this literature by expanding our analysis of the tax sensitivity beyond a single product category and/or a single online firm to a large number of product categories and online firms.

A second contribution to demand-side studies of the implications of Amazon's FC network choices lies in our analysis of possible convenience effects associated with a broader distribution network that cuts down on delivery times, which consumers may value independently of possible tax implications that come with Amazon's FC entry into a state. Our detailed information on the distance and delivery time between the closest FC to each consumer and their location allows us to investigate whether expenditures on Amazon respond to shorter shipping distances and faster delivery speeds, in addition to the effect of taxes. Our finding that the latter dominate any willingness-to-pay response due to increased delivery convenience rules out such demand-side benefits from proximity as a consideration for Amazon's network formation.

Our analysis is also related to recent economics literature examining the interplay of brick-and-mortar retail store locations choices and the retailer's distribution network. Zheng (2014) relates the proximity of a rival's fulfillment center to the chain's expected future entry, while Holmes (2011) estimates the savings in distribution costs associated with clustering stores near a fulfillment center. Apart from a small operations research literature that looks at the management of distribution networks for online firms (see Agatz et al. (2008) for an overview of this literature), little work to date has studied such classic industrial organization questions in the context of online markets.

The remainder of the paper is organized as follows. The next section discusses the background of the sales tax laws and Amazon's FC network. Section 3 introduces the main data sources and Section 4 presents the demand side analysis and results. Section 5 provides preliminary estimates of distribution cost savings and Section 6 concludes.

2 Amazon's FC Network and Sales Tax

We obtain information about Amazon's fulfillment centers from the supply-chain consulting company MWPVL, International.[9] MWPVL provides information on the location, size, opening date, and closing date of each FC. Also observable for a subset of the locations is the fulfillment center type. The type of FC is usually defined by the size of the item being shipped and/or the speed of delivery. The primary types are centers that focus on large items that cannot be sent combined with any other products (non-sortable), small items that can be combined in one package (small-sortable) and large items that can be combined into one package (large sortable). In the case of Amazon, other types of distribution centers include Amazon Fresh FCs, which supply Amazon's grocery delivery service, return centers, redistribution centers for third-party distribution, and cen-

[9] `http://www.mwpvl.com/`.

ters for select specialty items such as jewelry. Starting in 2014 Amazon started to build 'sortation' centers, which are used to sort packages by zip code after they have shipped from a FC, and 'Prime Now Hubs', which handle same-day delivery for the local markets. For the majority of the analysis, we focus on the three primary types of FCs and the Prime Now Hubs, as these are the centers that ship non-grocery items directly to consumers.

For these types, Amazon has expanded from 8 FCs in 2006 to 46 by the end of 2015 and has plans to expand further to over 90 FCs by 2018. Table 1 and figures 2 and 3 demonstrate this expansion for all types of distribution centers. There are a few things to note to about the expansion until 2010. First, Amazon placed FCs in relatively low population states that were close to highly populated areas. For example by 2010, Amazon had two FCs in Nevada, both of which were on the California border close to that state's major cities. Second, they also placed FCs in states with relatively low sales tax, but close to highly populated areas. For example, there were FCs in New Hampshire and Delaware, which are both close to major East Coast cities and have zero sales tax. This reflects that sales tax rates are positively correlated with population (across states, correlation of between 0.35 and 0.4 across years), so that entry into a small state near a large state has tax implications for only a small population, while allowing the firm to serve both states' populations more efficiently. Third, when Amazon did expand to highly populated states, they focused on states with a relatively low sales tax rates (e.g., Pennsylvania with an average tax rate of 6.14% in 2014, compared to 6.94% across the top 20 US states in terms of population). For comparison purposes, we include a map of Wal-Mart's distribution centers, which are spread across the country more evenly than Amazon's FCs, at least in the initial stages of Amazon's distribution network.

Given these patterns, it is clear that the early strategy of this expansion depended on tax laws that allow e-commerce firms to avoid charging their customers sales tax, which can be significant, reaching close to 10% in some states. In 1992, the United States Supreme court ruled that the Commerce Clause in the US Constitution "prohibits a State from imposing the duty of use tax collection and payment upon a seller whose only connection with customers in the State is by common carrier or by mail." Essentially, the law states that an online retailer does not have to charge a consumer sales tax unless it has a physical presence, or a 'nexus', in that consumer's state of residence. It is the duty of consumers to file a 'use-tax' return every year, which includes purchases from out-of-state vendors. However, very few individuals actually comply with this rule.[10]

For most online firms, a physical presence would come in the form of an office headquarters or a fulfillment center, implying that these firms likely do not have to charge sales in many states. As the popularity of e-commerce has grown, policy makers have started to suggest that these laws may be giving online firms an unfair advantage over their brick-and-mortar competitors. Further, states are likely losing out on millions of dollars of tax-revenue (see Bruce et al. (2009)). Because of their success and growing market share, Amazon.com has become the focus of politicians' complaints about these sales tax laws.

In the late 2000s, states began to introduce legislation that involved expanding the definition of a nexus to include the existence of 'affiliates', or websites that allow retailers, such as Amazon, to advertise on their site. For example, if a blogger based in Illinois has a link to Amazon on her site, then Amazon must charge sales tax to all the residents of Illinois. Not surprisingly, Amazon and other big retailers responded to this by shutting down their affiliate programs in states that passed these laws.[11,12] This certainly suggests that Amazon wanted to avoid charging sales tax to its customers. They even say as much in their 2008 annual report[13]:

> *A successful assertion by one or more states or foreign countries that we should collect sales or other taxes on the sale of merchandise or services could result in substantial tax liabilities for past sales, decrease our ability to compete with traditional retailers, and otherwise harm our business.*

However, it is also clear that as Amazon revenue grew in scale, the network of FCs expanded, presumably to be closer to population hubs despite sales tax implications and higher fixed costs of warehousing in densely populated areas. For example, by 2014, we see entry into highly populated

[10] Baugh et al. (2014) quote that 0.2% of people living in Rhode Island, 0.3% of people living in California and New Jersey, 7.9% of people living in Vermont, and 9.8% of people living in Maine report filing use tax returns.

[11] `http://techcrunch.com/2011/06/10/amazon-shuts-down-associates-affiliate-program-in-connecticut-over-online-sales-tax/`.

[12] `http://www.kansascity.com/news/local/article325412/Amazon-shuts-down-Missouri-associates-program-over-sales-tax-dispute.html`.

[13] `http://media.corporate-ir.net/media_files/irol/97/97664/2007AR.pdf`.

states such as California and Virginia and high tax states such as Tennessee with a tax rate of 9.46% in 2013 compared to a nationwide average of 6.62%. Finally, by 2018, Amazon plans to have operating FCs in Illinois, Georgia, Ohio and North Carolina. Table 1 displays the number of FCs, the numbers of states that have a FC and the number of taxed states by year. Again, we see that as the network expands, Amazon is entering new states and being taxed in a greater number of states. Overall, the pattern of expansion, both in location and in the dynamics, imply that there exists a trade-off between being close to customers and charging them sales tax. Our main goal is to empirically investigate this trade-off.

It is important to note that the entry of a fulfillment center does not necessarily mean that Amazon charges sales tax immediately. For example, Amazon first built a FC in Pennsylvania in 2006, but did not begin to charge sales tax until 2011. This is often due to legal battles with the state government as to what constitutes a nexus. On the other hand, sometimes Amazon charges sales tax even when they do not have a FC in that state. This can be due to changes in state laws (e.g., New York) or because of legal agreements with the state to begin charging sales tax ahead of entry (e.g., Connecticut). The latter explain the, at times, significant discrepancy between the number of states where Amazon's customers pay sales tax and the number of states where Amazon has a FC. Given this, we make various assumptions on how Amazon perceives the relationship between entry and sales tax when making its network location decisions.

3 Data

Consumer Online Purchases

The primary data source for the estimation of the demand model is the comScore Web Behavior Database. ComScore tracks the online purchasing and browsing activity of a sample of between 50,000 and 100,000 internet users (households) per year. The users give comScore explicit permission to monitor their activity. There are two primary databases, one that records each browsing session regardless of whether or not a purchase was made, and another that records transactions.

Because our focus is on buying behavior, we focus on the latter.[14]

For each transaction, we observe a unique household identifier, the time of the purchase, the product category and price for each individual item in the basket, the name of the domain where the transaction occurred, and a 'basket total', which is the total price of the transaction including shipping and taxes. In addition, we observe demographic characteristics for each household such as income, age of head of household, and racial background. Table 2 displays information about the reach of the sample in each year. The sample has gradually shrunk over the years from 86,000 households in 2006 to just 46,000 in 2013. Nevertheless, at least 75% of US counties and every US state plus the District of Columbia (minus Hawaii and Alaska) are represented in the data each year.

We find the comScore sample of households to be generally representative of the Unites States population according the 2010 census, with three exceptions: (1) the head of the household is younger, (2) a higher percentage of the households are white and (3) the household income is higher. All of these facts are likely because the sample is drawn from internet users, who are not perfectly representative of the US population. De los Santos et al. (2012) compare the sample of comScore users in 2002 and 2004 to the Computer Use Supplement of the Current Population Survey and find that the sample generally compares well with the population of online shoppers. However, to account for the possibility that comScore may be over or under sampling certain demographic groups, we adjust our data using sampling weights from the census. We bin each comScore household into categories based on income, age, and racial makeup, and calculate relative sampling weights based on the prevalence of each category in the comScore data relative to its prevalence in the 2010 census at the county level.

Table 3 is a first look at the purchasing patterns in the comScore sample. The first and second columns display the yearly average online expenditures and transactions per household, and the third column shows the percentage of households with zero transactions. Note that we have limited the data to only transactions in product categories that Amazon sells. Therefore, we omit purchases in categories such as travel and online dating. We further adjust yearly expenditures for the number

[14] ComScore is an internet analytics firm that provides data to Fortune 500 companies and large media organizations. See De los Santos et al. (2012) for a deeper discussion of comScore's services.

of weeks we observe the household in browsing data.[15]

Interestingly, we see decreasing average expenditures and transactions over time. Column 3 indicates that this is mostly due to an increasing percentage of households with zero online purchases. This is not in line with anecdotal evidence about the take-up of online shopping.[16] It is possible that the sampling procedure is the cause for this, but a more plausible explanation is that comScore is not recording the transactions for some households. This could be due to the household deactivating the comScore behavior monitor or using a second computer (e.g., at work) or a mobile device for their online shopping. Because we do not know which households truly have zero expenditures, rather than using other devices for purchases, we choose to exclude all households with zero transactions from the sample. In order to account for the extensive margin, we supplement the comScore data with survey data from Forrester Research.

Forrester Research conducts annual surveys of the online shopping behavior of a representative sample of the US population, including whether or not the respondents made any online purchases in the last three months, with the survey being conducted between March and July, depending on the year.[17] In the final column on Table 3, we see that there is an increasing number of households shopping online, according to the survey. We match the Forrester and comScore data based on demographic groupings, and calculate expected expenditures for a household in the comScore sample that accounts for the propensity of not making any online purchases in their demographic category in the Forrester data. See Appendix A for details. Therefore, after adjusting the comScore data with Forrester's extensive margin, we see a pattern that is more in line with the anecdotal evidence: expenditures and transactions are both increasing over time. Figure 4 further illustrates Amazon's growing market share in expected expenditures over time, growing from approximately 10% in 2006 to 40% in 2013. Note that since we focus only on expenditures at retailers that compete in categories Amazon covers, this market share overstates Amazon's market share in online retail in total, which is estimated at approximately 20% in 2013.

[15]We calculate the average weekly expenditures over the number of weeks we observe the household with any browsing activity and then multiply this by 52 weeks.

[16]`http://www.statista.com/statistics/183755/number-of-us-internet-shoppers-since-2009/`.

[17]Note that the survey was only available for 2006-2007 and 2010-2013. We interpolate linearly for the intervening years to construct predicted propensities of buying online.

A second limitation of the comScore data is that they do not record Amazon Marketplace purchases separately from Amazon purchases. A worry may be that consumers find a tax-free seller on Amazon Marketplace in a state in which Amazon itself must charge sales tax. This implies that we may underestimate the effect of sales tax because we would attribute a lack of response in states where Amazon charges taxes to a low sensitivity, when it is actually coming from consumers purchasing from a non-taxed vendor on Amazon Marketplace. However, most laws dictate that even marketplace vendors have to charge sales tax to customers in the states where their goods are housed.[18] To the extent that the vendors employ the 'Fulfilled by Amazon' program, then it is reasonable to assume that the sales tax implications of the network are the same for these vendors and Amazon. We perform a robustness check of the demand model that accounts for this possibility.

Amazon Fulfillment Center Network

We acquire information about each of the FCs from MWPVL, International. From the opening date and the location of the FC, we calculate the straight-line distance between each FC's street address and the population weighted centroid of every county in the US, and take the minimum of this by county and year. The distance serves as a measure of how far Amazon must ship the good between its FCs and the consumers. It is apparent from Figures 2 and 3 that Amazon often opens FCs that are close to an existing FC, which may be to increase capacity or to have FCs of different types near one another. Because of this, we define a FC cluster as a group of FCs that are in the same state and within 20 miles of each other. In the analysis below, we assume that Amazon decides where and when to build a cluster of FCs rather than each individual FC in the cluster, and the observed opening date and shipping distance are associated with the first FC built in the cluster. Note that this does not affect the demand analysis because the tax implications are associated with the first FC built with or without clustering.

Because straight-line distance may not be a perfect measure of shipping speeds, we also obtain the US Postal Service's shipping times between each three digit US zip code and find the minimum

[18] `http://www.avalara.com/learn/whitepapers/fba-sellers-guide-sales-tax/`.

shipping time between each county and each FC three-digit zip code.[19] The shipping times are for four different classes of mail: first, priority, standard, and package. These classes differ slightly in the size and number of packages allowed in the delivery. For a detailed description, see usps.com.[20] We were unable to find shipping times for the other shipping services used by Amazon, UPS and FedEx, but they are likely very similar to USPS.

MWPVL also provides information about the size of every FC in square feet and the number of employees for about half of them. To fill in the missing information about the number of employees, we assume that the number of employees per square foot is the same for FCs of the same type. The last two columns of table 1 provide the average number of employees and size for FCs which are open in a given year. There is a general pattern of an increase in size (and employees) early in the sample period, followed by a decrease as Amazon began to build the smaller Prime Now hubs starting in 2014. Note that we do not observe any information about the FCs for Amazon's competitors, so we cannot speak to competition among online firms in terms of their FC network choices.

We also obtain information on the extent to which a given fulfillment center can be used to satisfy same day shipping orders. There are two sources for this information. First, we found the date of implementation for Amazon's early version of same day shipping, called 'Local Express Delivery', from various news sources. Amazon began this service, for which all customers are charged a fee, as early as October, 2009. Second, using a tool on Amazon.com, we find the markets which have free same day shipping available to Prime users as of the Fall of 2015.[21] This information allows us to construct a measure of the radius from a FC in which local customers have access to same day delivery. In future work, we plan to use this data in estimating the demand-side preference for same-day shipping as a shifter of county-level expenditures at Amazon.

[19] We choose one three digit zip code per county after verifying that shipping times do not vary much within a given county.

[20] `http://pe.usps.com/businessmail101/classes/welcome.htm`.

[21] `http://www.amazon.com/b?node=8729023011`.

Other Sources of Data

In addition to the above sources, we obtained state, county and local sales tax rates from Tax Data Systems, now part of Thomas Reuters. For each year and county we calculate the average tax rate, as local taxes can vary within a county and may change mid year. Table 4 displays the average and standard deviation of the tax rate across counties, as well as the average standard deviation of tax rates across counties within a state. The table demonstrates that there is not only variation in tax rates across states, but also substantial variation in rates across counties within a state. In addition, tax rates vary across time, as between 30 and 65% of counties experienced tax changes each year and nearly 78% of the counties experienced at least one tax change over our sample period.

We collected county-level wage information from the BLS's Quarterly Census of Employment and Wages and used the annual wage of a 'Retail Trade' employee. While this is not the exact industry in which the Amazon FC employees work, we believe the wages are likely to be similar in a given county. This is evidenced by the fact that the overall average hourly wage of a 'Retail Trade' employee is close to that of an Amazon Associate as reported by glassdoor.com. The land rents were computed by converting the average residential property value in a county to a yearly payment. The property values were obtained from the United State census via `http://metrocosm.com/get-the-data/`. Finally, we collected county level demographics from the 2010 Census, along with annual information on the number of small and large retailers from County Business Patterns.

Table 5 shows the average demographic attributes of the counties and states where Amazon has a FC relative to the average across the US. Where noted, the averages are weighted by population. Before 2006 the FCs where in relatively low populated states, but starting in 2007 the average number of households in states with a FC exceeds the average overall. This is due the entry into states with above average population such as Pennsylvania (2007) and Indiana (2008) in the early stages and highly populated states such as California (2014) and Florida (2015) thereafter. Additionally, before 2016, the average property values for counties in which there is a FC are far lower than the national average, suggesting that Amazon is choosing to put FCs in relatively cheap locations. However, the annual wage for retail workers is higher in these counties, possibly offsetting these costs. Note that the large jump in property values late in the time frame is due to the construction of Prime Now facilities built outside of Los Angeles and San Francisco. Interestingly, pre 2016, Amazon seemed to be locating FCs in higher populated states but within those states, in counties with lower costs for both land and labor. Obviously this changes substantially as they locate near highly populated areas.

Finally, we use Amazon's annual reports to provide additional information about Amazon's finances. That is, for each year, we obtain the aggregate sales in relevant categories, the cost of goods sold, which includes inbound and outbound shipping costs, and the aggregate outbound shipping cost.[22] This allows us to formulate the gross profit margin for Amazon, net of outbound shipping cost, which we discuss in more detail in section 5.

4 Demand Analysis

We begin the demand analysis by providing reduced form evidence that consumers are sensitive to sales tax. First, we estimate the transaction level effect of taxes on the likelihood a consumer purchases from Amazon. That is, we run the following linear probability regression:

$$Pr(A_{hijt} = 1) = \beta_0 + \alpha \ln(1 + \tau_{it}) + \lambda_i + \lambda_j + \lambda_t + \epsilon_{hijt}$$

where each observation is a purchase occasion h, from a consumer in county i, in year t, for a product in category j. The dependent variable is a dummy variable equaling one if the consumer purchased the product from Amazon.com and zero otherwise. The tax rate is given by τ_{it}; it varies depending on the county tax rate and whether or not Amazon has to charge sales tax to customers in that county in a given year. We do not include the item's price or the item's price at Amazon relative to the price charged by an outside seller because we observe the price only on the platform where the transaction occurs, but not on alternative platforms.[23] However, we include product category and time fixed effects to account for average price differences across online retailers at the

[22] `http://phx.corporate-ir.net/phoenix.zhtml?c=97664&p=irol-reportsannual`.

[23] Nevertheless, including a variable that equals to Amazon's price for Amazon purchases and zero otherwise results in similar estimates.

category level and across time.

The result in the first column of table 6 indicates that consumer are sensitive to taxes and that they reduce their purchases from Amazon by nearly 16% for a one percentage point increase in sales tax. In columns 2 through 7, we include various measures of shipping speeds to account for the fact that consumers may value faster shipping times from Amazon. The measures are the log of the distance between the centroid of the consumer's county and the closest FC and the shipping times from USPS for the four different classes of mail: priority, first, standard, and package. Since the variation in shipping days is limited, we aggregate shipping times into two categories, fast and slow shipping speed between two three-digit zip codes, excluding the slow shipping speed indicator from our specifications. We find that the tax effect is robust to including the variety of proxies for shipping time. The log of the shipping distance is statistically significant but has a near-zero economic effect on the propensity to purchase from Amazon. This is also the case for most mail classes, where the probability of purchasing from Amazon does not change by a significant amount compared to customers in zip code locations with a slower shipping speeds, all else equal. The only parameter estimates that indicate that people significantly prefer faster shipping times is comparing fast priority shipping (1 or 2 day delivery time) to slow (3 day) priority shipping.

There are a number of explanations for the mixed results on the shipping speed coefficients. It could be that our measures do not accurately reflect shipping times. Another explanation is that consumers simply do not have tastes for faster shipping, or have very nonlinear preferences for shipping times given Amazon's recent expansion to same-day shipping. Finally, it could be that the expansion of the network of fulfillment centers did not result in faster shipping times to consumers, or if it did, it was a small number of consumers who were affected. We believe this to be the most plausible explanation.

Next, we perform a differences-in-differences analysis of the response in Amazon expenditures to a change in the tax status of a county. We summarize a county's tax status in an indicator variable, $\mathbf{1}_{it}^{taxable}$, that is one if a purchase from Amazon.com by a household in county i in year t is subject to sales tax. The tax status variable changes over time due to the entry of a FC and/or due to the future entry of a FC where the state insisted on collecting sales taxes immediately following the initial agreement. The regression model is given by:

$$e_{it} = \beta_0 + \beta_1 \mathbf{1}_{it}^{taxable} + \lambda_t + \lambda_i + \epsilon_{it}$$

where e_{it} is the log of aggregate expenditures on Amazon from county i in year t. Aggregate expenditures are the sum of expenditures on Amazon in year t for the comScore households who live in county i. Results in table 7 suggest that a change in the tax status of a county results in somewhere between a 10 and 12% reduction in expenditures on Amazon, with this effect being significant at the 10% level in all specifications and at the 5% level for specifications (4) and (6). This exercise is essentially the same as the one performed in Baugh et al. (2014) with different data, and they estimate a tax effect of about 9.5%. We once again include measures of shipping speeds in this demand estimation and find little evidence that they significantly affect the purchasing behavior of consumers.

A weakness shared by both of the above exercises is that we do not consider the substitution between Amazon and other taxed and non-taxed shopping options. Because of this, any reduction in Amazon's transactions due to higher taxes does not equate directly to a tax elasticity since these transactions may be substituted towards other taxed outlets (e.g., walmart.com or Wal-Mart). Therefore, we specify and estimate a model of demand for retail goods across all modes of online and offline shopping in the following sections that allows us to determine the effect on overall transactions of higher sales taxes.

Empirical Model

We specify a model of demand where a consumer chooses how much money to spend on retail goods in each of four different modes of shopping. The modes are Amazon, $j = 1$, taxed online competitors, $j = 2$, non-taxed online competitors, $j = 3$, and offline competitors, $j = 0$.

We follow Einav et al. (2014) and specify a CES utility function. Specifically, a representative consumer from county i solves the following problem in year t:

$$\max_{q_{i0t},..q_{i3t}} \left(\sum_{j=0}^{3} \left(\frac{q_{ijt}}{\zeta_{ijt}} \right)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$$

$$s.t. \sum_{j=0}^{3} p_{ijt} q_{ijt} \leq w_{it}$$

where q_{ijt} represents the quantity purchased via shopping mode j in time period t, p_{ijt} is the price of purchasing one unit of goods via shopping mode j, ζ_{ijt} is the taste for mode j, and w_{it} is total dollars spent on retail goods. The elasticity of substitution between the four modes is given by σ. Solving for the optimal amount of expenditures in each online mode results in:

$$e_{ijt} = \frac{(p_{ijt}\zeta_{ijt})^{1-\sigma}}{P_{it}^{1-\sigma}} w_{it},$$

where P_{it} denotes a weighted-average price index across all four modes. Dividing this by the expenditures for $j = 0$ and taking logs gives:

$$\ln(e_{ijt}) - \ln(e_{i0t}) = (1-\sigma)(\ln(p_{ijt}) - \ln(p_{i0})) + (1-\sigma)(\ln(\zeta_{ijt}) - \ln(\zeta_{i0t})) \tag{1}$$

where the price of each mode can be written as:

$$p_{ijt} = (1 + \tau_{it} \mathbf{1}_{ijt}^{taxable}) \tilde{p}_{ijt}$$

The tax rate in county i and year t is given by τ_{it}, meaning $\tilde{p}_{ijt}$ is the tax-exclusive price of buying goods, which can include shipping charges. Note that the tax liability may vary across j as the non-taxed competitor never charges sales tax and Amazon does not charge sales tax in a number of states and years. We indicate whether mode j entails a sales tax liability for the customer at time t through the indicator $\mathbf{1}_{ijt}^{taxable}$; purchasing from an offline competitor always implies that the customer has to pay sales tax. With this, equation 1 becomes:

$$\begin{aligned}\ln(e_{ijt}) - \ln(e_{i0t}) =& (1-\sigma)(\ln(1+\tau_{it}\mathbf{1}_{ijt}^{taxable}) - \ln(1+\tau_{it})) \\ &+ (1-\sigma)(\ln(\tilde{p}_{ijt}) - \ln(\tilde{p}_{i0t})) \\ &+ (1-\sigma)(\ln(\zeta_{ijt}) - \ln(\zeta_{i0t}))\end{aligned} \tag{2}$$

We model the taste for the each mode $j \in \{1,2,3\}$ as:

$$\zeta_{ijt} = \exp(\xi_{it}^{o} + \xi_{jt} + \beta_j Z_{it} + \lambda_j C_{it} + \epsilon_{ijt})^{\frac{1}{1-\sigma}}$$

and assume that the preference for the offline shopping mode equals $\zeta_{i0t} = \exp(\epsilon_{i0t})^{\frac{1}{1-\sigma}}$. The term ξ_{it}^{o} represents the preference for online shopping in county i at time t. We decompose ξ_{it}^{o} into a county-level fixed taste for online shopping, relative to offline, and a time trend for online shopping that we allow to vary at the Census district level. That is, $\xi_{it}^{o} = \xi_{i}^{o} + \xi_{rt}^{o}$, where r denotes the Census district, of which there are nine in the continental US. Similarly, ξ_{jt} captures the time varying preference for shopping mode j; Z_{it} is a vector of county level demographics; and C_{it} is a vector of variables measuring the level of offline competition.

Notice that we allow the effects of demographics and competition to vary across shopping modes. For example, having a brick-and-mortar Best Buy in a county may affect the purchases through taxed online competitors (including bestbuy.com) differently than purchases from Amazon. Motivated by our descriptive evidence above that suggested that distance to the closest FC – or shipping time – did not affect the propensity of purchasing from Amazon, we do not include it in the base model, but do so in various alternative specifications.

The demand shocks are at the county, time, and mode level, and are assumed to be iid across these delineations. Under these assumptions, equation 2 becomes:

$$\begin{aligned}\ln(e_{ijt}) - \ln(e_{i0t}) =& (1-\sigma)(\ln(1+\tau_{it}\mathbf{1}_{ijt}^{taxable}) - \ln(1+\tau_{it})) \\ &+ (1-\sigma)(\ln(\tilde{p}_{ijt}) - \ln(\tilde{p}_{i0t})) \\ &+ \beta_j Z_{it} + \lambda_j C_{it} + \xi_{i}^{o} + \xi_{rt}^{o} + \xi_{jt} + \epsilon_{ijt} - \epsilon_{i0t}\end{aligned} \tag{3}$$

We make the assumption that the prices for the online shopping modes do not vary across counties, or that $\tilde{p}_{ijt} = \tilde{p}_{jt}$. While some online firms appear to price discriminate based on a consumer's location, there is limited evidence that this is widespread.[24] Further, Amazon attempted to implement price discrimination in 2000, and quickly abolished it after a backlash from customers.[25,26] Finally, we assume that the price of the offline shopping option remains constant over time, or $\tilde{p}_{i0t} = \tilde{p}_{i0}$. With this, we can re-write equation 3:

$$\begin{aligned}\underbrace{\ln(e_{ijt}) - \ln(e_{i0t})}_{\tilde{e}_{ijt}} =& (1-\sigma)\underbrace{\left(\ln(1+\tau_{it}\mathbf{1}_{ijt}^{taxable}) - \ln(1+\tau_{it})\right)}_{\tilde{\tau}_{ijt}} \\ &+ \beta_j Z_{it} + \lambda_j C_{it} + \xi_{rt}^o + \underbrace{\xi_{jt} + (1-\sigma)\ln(\tilde{p}_{jt})}_{\bar{\xi}_{jt}} \\ &+ \underbrace{\xi_i^o - (1-\sigma)\ln(\tilde{p}_{i0})}_{a_i} + \underbrace{\epsilon_{ijt} - \epsilon_{i0t}}_{\tilde{\epsilon}_{ijt}}\end{aligned} \tag{4}$$

resulting in the following equation for the differences in expenditures on mode j and the offline mode:

$$\tilde{e}_{ijt} = (1-\sigma)\tilde{\tau}_{ijt} + \lambda_j C_{it} + \beta_j Z_{it} + \xi_{rt}^o + \bar{\xi}_{jt} + a_i + \tilde{\epsilon}_{ijt} \tag{5}$$

This equation says that the difference in expenditures between mode j and the offline mode is a function of the difference in sales tax between mode j and the offline mode, a time varying mode level effect that includes the price level of mode j, the relative effect of demographics and local competition on mode j, a region-specific time trend, a county effect, and an iid demand shock.

County Level Expenditures

To form the expenditures for the representative consumer ($\tilde{e}$), we calculate the population weighted yearly average expenditures on each shopping mode for each county. We define a taxed online competitor as one that has a large offline presence such as gap.com, walmart.com and target.com

[24] `http://www.wsj.com/articles/SB10001424127887323777204578189391813881534`.
[25] `http://www.bizjournals.com/seattle/stories/2000/09/25/daily21.html`.
[26] `http://news.cnet.com/2100-1017-240700.html`.

and a non-taxed competitor as one without a national offline presence, such as overstock.com. We note that neither of these groups are perfectly defined, as some sites that we classify as having a national offline presence may not have a store in every state and sites that we classify as not having an offline presence surely have a headquarters and/or a fulfillment center located in at least one state. Table 8 displays the top ten online stores in each of the taxed and non-taxed categories. Overall, we classify about 34% of the websites that compete in product categories that Amazon carries as taxed competitors. We also exclude any expenditures in product categories that Amazon does not cover and exclude any websites that sell only in these categories. Examples of excluded sites are dating websites (e.g., Match.com), travel websites (e.g., orbtiz.com) and food delivery sites (e.g., dominos.com).

Under these restrictions, we calculate the yearly expenditures for the three online shopping modes for each individual household. The level of household expenditures is then adjusted using the Forrester Data as discussed above. The average expenditures for a county are calculated using demographic weights from the 2010 census. Finally, expenditures for the offline mode are calculated by subtracting online expenditures from the total amount of expenditures on retail goods, assuming that total retail spending equals 10% of consumer's income.[27] See Appendix A for further details on how expenditures are calculated.

We examine patterns of spending across the shopping modes in Figure 4. The figure displays four different charts, each with the market share of the shopping modes on the left y-axis. Note that this is the market share of online purchases only. Amazon's market share has steadily increased from under 10% in 2006 to around 40% in 2013, while both the taxed competitors' and non-taxed competitors' market share has declined from around 45% to 30%. At the same time, there has been an increase in the number of fulfillment centers, the number of states with fulfillment centers and the number of states that are taxed. Finally, the average distance from the FC to the consumer has declined from around 300 miles to 150 in 2015 and is projected to further decrease to about 100 miles by 2018.

[27] Using data on the relative share of online retail out of total retail from `https://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf` in combination with total online spending in the comScore data, we calculate average total expenditures on retail goods to be approximately 10%.

Estimation

We estimate equation 5 using OLS with county, mode/year, and region/time level fixed effects. Included in C_{it} is the total number of retail establishments and the number of establishments with more than 50 employees (i.e., large retailers) in county i in year t to account for the fact that large offline retailers (e.g., Wal-Mart) may have a different impact on online demand than small local retailers. County level demographic included in Z_{it} are median income and the distribution of ethnicity.

Identification of the tax sensitivity parameter comes from several sources of variation. First, there is variation in tax rates in a county across time due to changes in local laws and/or entry of an Amazon fulfillment center. Therefore, changes in expenditures between taxed and non-taxed modes as a result of these changes helps to identify σ. This is similar to the traditional identification argument in difference-in-differences models. Second, within a county and year, some modes are taxed and some are not. The extent to which people's expenditures vary across modes because of the sales tax charged also helps to identify tax sensitivity. Therefore, in order to estimate the tax sensitivity, we make the reasonable assumption that the changes in tax rates due to either Amazon's expansion and/or local laws are exogenous to the unobserved local demand shocks. This reflects that the Amazon FC entry decisions are likely made at a higher geographic level than the county and not made purely in response to a local annual demand shock in a specific county. To investigate the sensitivity of our results to this assumption and, similarly, to the assumption above that the price of the offline option does not vary over time, we have also estimated the above specification with county/year fixed effects. We find similar results, but choose to proceed with the main specification to avoid identifying the tax effect purely from variation in taxes across modes within a county, since we observe only three modes and at most two tax levels across the modes.

As in the reduced form regressions above, we explore whether consumers are sensitive to shipping speeds. We identify this effect using changes in the distance and shipping times resulting from the expansion of the FC network. Specifically, expansion of the network results in shorter routes and faster shipping times for some consumers. Further, because the tax effects of the expansion are at the state level, there are some consumers who live near state borders who see a decrease in delivery times, but not a change in tax rates. Note that we do not observe the locations and thus distance to Amazon's competitors' fulfillment centers, but we assume that the time-varying preferences for other modes, ξ_{jt}, capture any overall effect of changes in shipping times for a given mode. Therefore, the effects of shipping speeds are Amazon specific.

Finally, a nice feature of the data is that entry of a fulfillment center into a state does not always lead to Amazon having to charge sales tax (e.g., the Pennsylvania example from above). Amazon may also have to charge sales tax in a state without having a fulfillment center. Therefore, even within a state, we sometimes see variation in distances without variation in taxes and vice versa.

Results

Results of 6 different specifications are presented in Table 9, where each specification varies by the measure of shipping speed included. Specification (1) does not include any measure, specifications (2) includes the log of the distance to the shortest FC and specifications (3)-(6) include variables for shipping speed in four different classes of delivery. In the latter, the longest delivery time is the excluded dummy variable. The estimated value of σ, or tax sensitivity, is around 1.35 and is significant at the 1% level in all specifications. These estimates are similar to the ones reported by both Einav et al. (2014) and Baugh et al. (2014).

The estimates imply that going from not charging sales tax to charging the average sales tax of 6.5% would reduce expenditures on the given mode by around 9 percent. That is, if Amazon agrees to charge sales tax in a state through either an agreement with the state government or because they build a fulfillment center, they can expect their revenue to decrease by up to 9%, all else equal. In 2008, New York state passed a law that required Amazon to collect sales tax. Our estimates imply that this reduced Amazon's revenue in New York by around 12%.

Once again, we do not find strong evidence that our measures of shipping times drive substitution between Amazon and other shopping outlets. This is further evidence that the expansion of the network did not significantly change the shipping times from Amazon, or in other words, the quality effect of expanding the network is approximately zero.

The estimates of the time varying mode fixed effects are presented in Table 10, with the excluded effect being mode 3 in 2006. Amazon's quality is increasing over time, which is not surprising considering the increase in market share over this time period. The taxed online competitors' quality increases at first and then decreases, while the quality of the non-taxed competitors is relatively constant. Finally, the estimates of the parameters on C and Z are presented in 11. This table shows that all types of firms are affected by the number of offline retailers and this effect becomes stronger as the number of large offline retailers increases. However, none of these effects are significant. Many of the demographic coefficients are not significant. Note that because we do not have the demographic variables varying over time, we exclude mode 3's effects.

Table 12 provide some robustness for these results. Specification (1) uses weighted OLS, where the weights are analytical weights based on the number of observations used to calculate the average expenditures in a county, specification (2) incorporates the zeros by changing any mode-level expenditures that are equal to \$0 to \$1, specification (3) uses data from 2008 and beyond because the number of households in the comScore sample shifts after 2007, specification (5) does not use the Forrester correction to adjust expenditures, and specification (6) does not use population weights to create the expenditures. Specification (4) estimates the tax sensitive separately for each mode, as a way of accounting for Amazon Marketplace. That is, it allows for an Amazon specific tax effect that may be lower than the effect for offline retailers and mode 3 because not all customers on Amazon are charged sales tax. Overall, the results of these robustness checks are consistent with the base results with the estimated tax sensitivity being between -1.27 and -1.88. This implies that, if anything, we are likely underestimating the tax sensitivity with our base estimates.

Calculating Total Revenue

Using the estimates of the model, we calculate total revenue in the US for Amazon for the years 2006-2013 and also project revenues in 2014-2018. To do this, we first form predicted revenue for the representative household for every county and year that we observe in the comScore data. Then, for the counties that we never observe, we assume that county level fixed effect is equal to that of the closest county, which allows us to form a predicted amount expenditures on Amazon for the representative household for every county in the US from 2006-2013. We multiply the expenditures of the representative household by the number of households in the county and sum this over all counties for each year.

We compare this revenue to the equivalent sales data reported in Amazon's financial statements.[28] The predicted revenue is less than the revenue reported because, as discussed above, we are missing purchases from other computers or devices and missing purchases from people who are able to turn off the monitoring device. Therefore, we calculate a 'multiplier' of the reported sales divided by the predicted revenue and use it to create an adjusted expenditure for each household across all modes.[29] The implied multiplier is in line with reported shares of purchases that online shoppers make via a mobile device (24.6% in 2014) and the incidence of at-work-online shopping (47% in 2014).[30] The assumption here is that if we are missing purchases, we are missing them equally from all modes. We then re-run the model and perform the same procedure as above to form 'predicted' revenue, that now equals the sales in the financial report.[31]

To calculate revenues for 2014-2018, we need to obtain future values of the mode specific year effects and the region time effects. We do this by using the growth (or decline) of these effects from 2006-2013 to predict future effects.[32] Table 13 reports the total predicted revenue from 2006-2018. Because we observe Amazon sales in 2014 from their financial report, we can assess how well this procedure does in predicting sales in 2014. We find that we over predict revenue (52.8 billion versus 50.8 billion in the reports), but that the prediction is reasonably close. We also predict that expenditures on mode 2 fall from \$28 billion to \$23 billion and expenditures on mode 3 decline by approximately \$6 billion over the next 4 years. Also displayed are the average shipping distance from the FCs to the customers, which decreases over time, and the number of counties that are taxed, which increases with the FC rollout into new states.

[28] We use the figures of sales from North America in the "Media" and "Electronics and Other General Merchandise" categories. This excludes the "Other" category, which is revenue from "Non-Retail Activity" such as Amazon Web Services (AWS).

[29] The multipliers are 1.05, 1.44, 1.37, 1.40, 1.54, 1.50, 1.46, and 1.71 for the years 2006-2013. The increasing rate is likely due to the fact that mobile shopping has increased over this time period.

[30] See `https://www.internetretailer.com/2015/08/18/mobile-commerce-now-30-all-us-e-commerce` and `http://www.careerbuilder.com/share/aboutus/pressreleasesdetail.aspx?sd=12%2F1%2F2014&id=pr854&ed=12%2F31%2F2014`.

[31] Re-running the model only shifts the mode level effects but does not change the estimates of tax sensitivity.

[32] Specifically, we run the following regression $\hat{m}_t = \gamma_0 + \gamma_1 \ln(t) + \gamma_2 \ln(t)^2 + \epsilon_t$, where $\hat{m}_t$ represents the estimate of the effect from the 2006-2013 data. Using the estimates of γ we predict future values of m_t.

Evidence of a Tax/Distance Trade Off

The demand estimates suggest that expanding the network has a negative effect on revenues through the tax effect. In order to quantify this effect, we perform a simulation exercise using the estimates of the base model. We calculate counterfactual revenue under a different network of FCs and different tax laws, and assume everything else remains constant. Specifically, we calculate the total expected revenue if the network of fulfillment centers had remained fixed since 2006.

The difference between the counterfactual revenue and observed revenue is presented the first column of Table 14 along with the difference between the average shipping distance. Results indicate that expansion leads to a reduction in expected revenue of around $8.9 billion, or a 1.5% reduction. However, this reduction in revenue is accompanied by a reduction in shipping distance as evidence by the 2nd column of Table 14. This gives a first glimpse at the trade-off that Amazon faces when expanding the network.

To further investigate this trade off, we perform a few additional exercises which can be found in table 15. First, we note that there is correlation between the population of a state and the sales tax rate, implying that moving an FC to a highly populated state results in lower revenue, but likely results in shorter shipping routes.

Second, we analyze the trade off of having an FC near a densely populated area in terms of distance and lost revenue. Specifically, for each FC, we find the county with the highest population density for which it is the closest FC (i.e., the FC which ships to that county), and "move" the FC to that county in the year of its opening. We calculate the distance between the FC and that county, along with the change in revenue from the move. Table 16 provides some examples of these moves for FCs which opened in different years. For the FCs which opened early, the trade-off is clear: not being located in the highly populated counties results in an increase in revenue. However, there are some instances where moving closer to the population actually results in an increase in revenues. This reflects either a disproportionate increase in population when moving the FC into the most densely populated county it serves (e.g., a move from Balimore, MD to DC), and/or a disproportionate decline in taxes from moving to the most densely populated county from the current location (e.g., a move from Hamilton county, Tennessee to DeKalb, GA, which contains

Atlanta). These patterns disregard, of course, that there may be other advantages to the chosen location; Hamilton county, Tennessee, for example, has significantly lower fixed costs than Atlanta. Table 15 summarizes the collection of moves for each year. Note that we are moving a "cluster" of FCs as defined above, rather than each FC of the cluster. Again we see the pattern that, early on, FCs were built in different states than the closest highly populated area. However, over time, it is clear that as more and more of FCs are being built, a move to of the FCs to the most densely populated area it serves no longer entails a move across state borders. As a result, there is no effect of a move; the FC moves within its own state.

5 Quantifying Cost Savings from Distribution Network

We now turn to estimating the cost savings of expanding the network of FCs. We posit that these cost savings come the fact that, after expansion, Amazon shortens the outbound shipping distance (i.e., from the FC to the customer) that is handled by one of the contracted shipping companies. That is, with more localized FCs, the outbound shipping distance from the FC to the shipping service's sorting center becomes shorter. Further, the expansion may even remove the leg going from the FC to the sorting facility all together, allowing the shipping service to deliver the package directly from the FC to the customer. Under either scenario, the service's costs are lower, giving Amazon more bargaining power when negotiating their per package shipping rates. One concern is that the expansion of the network may also lead to longer inbound shipping routes from the supplier to the FC resulting in higher inbound shipping costs. However, the fact that the suppliers already deliver goods to many retailers across the US and/or have their own widespread distribution network alleviates this concern to a certain extent.

5.1 Profit Function

To quantify cost savings formally, we formulate a profit function for Amazon that depends on the network of FCs it operates in year t, denoted as a_t, in a similar manner as Holmes (2011). In county i, Amazon generates $R_{it}(a_t)$ in revenue, which is a function of a_t due to sales tax implications and/or the shipping speeds. Variable profit in year t from county i is given by:

$$\pi_{it}(a_t;\theta) = \mu R_{it}(a_t) - \theta d_{it}(a_t) R_{it}(a_t)$$

where μ is the percent markup net of shipping costs and variable shipping cost is given by $\theta d_{it}(a_t) R_{it}(a_t)$. The parameter θ measures the shipping cost per dollar of goods sold per mile and $d_{it}(a_t)$ is the distance from consumers in county i to the closest FC. The biggest difference with Holmes (2011) is that we allow variable costs to vary across network location decisions. We assume that this variation comes from differences in the shipping distance from a FC to consumers in county i. Therefore, total profit for Amazon, starting in 2006 is given by:

$$\Pi(a;\theta) = \sum_{t=2006}^{\infty} \beta^t \sum_i \mu R_{it}(a_t) - \theta d_{it}(a_t) R_{it}(a_t) - F_t(a_t) \tag{6}$$

where a represents the 'rollout' of the entire network and β is the discount factor. The fixed costs of the network rollout are given by $F(a_t)$. The fixed cost of the network are assumed to be made up of wages paid to employees of the FCs in county i, called 'Associates' and denoted by A_{it}, the land rents for the total square footage of building space in county i, L_{it}, and the other unobserved fixed costs of building/running a FC, F, which are constant across location and time:

$$F_t(a_t) = \sum_{i \in I(a_t)} w_{it} A_{it} + r_{it} L_{it} + F$$

The sum is over all the counties where the FCs are located under network a, $I(a_t)$.

We assume that Amazon chooses its optimal prices independently of a_t, so the profit maximization problem, in terms of network formation, can be written as:

$$\max_a \Pi(a;\theta) \tag{7}$$

Given the maximization problem in 7, we know that the profit under the observed network rollout

(a^o) must be greater than any deviation from this (a), or:

$$\Pi(a^o;\theta) \geq \Pi(a;\theta) \forall a \neq a^0$$

This allows us to formulate a set of moment inequalities that we use to estimate θ.

5.2 Moment Inequalities

As in Holmes, we construct moment inequalities by comparing the discounted profits of the observed sequence of openings with counter-factual roll-outs scenarios. We focus in particular on counter-factual roll-outs that *swap* the opening dates of two FCs. By swapping two opening dates (e.g. $t < t'$), we ensure that the difference in the discounted stream of profit cancels out beyond period t'. This allows us to impose optimality conditions without observing the complete sequence of openings. Additionally, swapping the opening dates of two FCs implies that the unobserved fixed costs, F, will cancel out, eliminating any issues with selection of the FC location based on unobservables.

Consider two opening sequences: observed sequence a^0, and counter-factual sequence a. If profits were measured without error, profit maximization implies the following linear profit difference inequality:

$$y_a - x_a\theta \geq 0 \tag{8}$$

where $y_a = \sum_{t=2006}^{\infty} \beta^t \left[\mu(R_t(a_t^0) - R_t(a_t)) - (F_t(a_t^0) - F_t(a_t))\right]$ is the discounted profit differential net of shipping costs, $x_a = \sum_{t=2006}^{\infty} \beta^t \sum_i \left[d_{it}(a_t^0)R_{it} - d_{it}(a_t)R_{it}\right]$ is the discounted differences in revenue weighted shipping distance, and θ is the shipping cost per mile per dollar of revenue.

The cost savings associated with the optimal roll-out of fulfillment centers are identified by revealed preference. Intuitively, the magnitude of θ is such that the predicted revenue loss relative to counter-factual sequence a (i.e. $y_a < 0$) is offset by a reduction in shipping distance (i.e. $x_a < 0$). This gives a upper bound on the distance cost. Similarly, θ is such that a counter-factual sequence associated with an increase in shipping distance (i.e. $x_a > 0$) is offset by a revenue gain from tax avoidance (i.e. $y_a > 0$). This provides a lower bound on the distance cost.

Although this tradeoff is clear when comparing y_a and x_a for alternative paths, we measure

Amazon's profits and weighted shipping distance with error. In particular, we observe $\tilde{y}_a$ and $\tilde{x}_a$:

$$\begin{aligned}\tilde{y}_a &= y_a + \eta_a, \\ \tilde{x}_a &= x_a + \nu_a.\end{aligned}$$

These errors may originate from a variety of sources: (i) measurement error in the operating costs of FCs, (ii) estimation error in Amazon's revenue that would flow through to both y_a and x_a, (iii) omitted variables, and (iv) differences in tax implementation (i.e. Amazon's beliefs). We assume that both errors are conditionally independent of a vector of non-negative instruments z_a:

$$E(\eta_a|z_a) = E(\nu_a|z_a) = 0. \tag{9}$$

Applying these two conditional moment restrictions to the "measured" inequality condition (analog from equation 8) gives:

$$E\left(\tilde{y}_a - \theta\tilde{x}_a|z_a\right) = E(y_a - \theta x_a|z_a) + \underbrace{E(\eta_a - \theta\nu_a|z_a)}_{=0} \geq 0 \tag{10}$$

Thus, our estimator is based on a set of unconditional moment conditions consistent with the previous inequality:

$$E\left[z_a(\tilde{y}_a - \theta\tilde{x}_a)\right] \geq 0. \tag{11}$$

This leads to the following sample moment inequalities:

$$\frac{1}{M}\sum_a z_{a,k}(\tilde{y}_a - \theta^0\tilde{x}_a) = \tilde{m}_k(\theta^0) \geq 0, \quad \forall k = 1, \ldots, K. \tag{12}$$

where M denotes the number of inequalities. The objective function therefore becomes:

$$Q(\theta) = \sum_k \min\{0, \tilde{m}_k(\theta)\}^2 \tag{13}$$

and we search for the value(s) of Θ which minimizes equation 13:

$$\hat{\theta} = \arg\min_{\theta} Q(\theta) \tag{14}$$

We take two approaches in constructing the instruments. First, following Holmes, we break the choice of instruments into two parts: (i) basic instruments, and (ii) interactions of the basic instruments with state variables. The basic instruments are indicator variables that identify perturbations that capture the interaction between the effect of sales taxes and the shipping distance. This corresponds to the interaction between cannibalization and economies of density in Holmes. In his case, there is a one-to-one mapping between these two effects: changes in population density affect both firm costs and demand-side cannibalization between stores. As a result, perturbed store networks can be selected based purely on an increase or decrease in exogenous population density relative to the observed roll-out.

In our case, an increase in shipping distance does not directly lead to tax savings or vice versa. We thus cannot identify appropriate perturbations based on changes in a single, exogenous variable. Instead, we select inequalities based on the interaction of both elements, that is based on how changes in shipping distance and/or taxes affect profit net of shipping cost and variable shipping cost (i.e. $\tilde{y}_a$ and $\tilde{x}_a$). An additional complication is that $\tilde{y}_a$ and $\tilde{x}_a$ are measured with error. To get around this problem, we construct "exogenous" predicted responses to network perturbations, $\hat{y}_a$ and $\hat{x}_a$, that are uncorrelated with this measurement error. We identify relevant perturbations for our basic instruments using these alternative predicted responses.

Consider for instance a counter-factual rollout plan a. We calculate predicted exogenous responses in shipping cost ($\hat{x}_a$) and remaining profit ($\hat{y}_a$) as:

$$\hat{x}_a = \sum_t \sum_i \beta^t \left[d_i(a_t^0)\hat{R}_{it}(a_t^0) - d_i(a_t)\hat{R}_{it}(a_t) \right] \tag{15}$$

and

$$\hat{y}_a = \sum_t \sum_i \beta^t \left[\mu \left(\hat{R}_{it}(a_t^0) - \hat{R}_{it}(a_t) \right) - \left(F_t(a_t^0) - F_t(a_t) \right) \right] \tag{16}$$

where $\hat{R}_{it}(a_t)$ is predicted revenue in county i estimated based on pre-determined exogenous variables. We formulate $\hat{R}_{it}(a_t)$ by running the following "first-stage" regression of revenue predicted by the CES demand model on county-level observables (X_{it}, population, income, offline retail structure, etc.) that do not depend on the network rollout and the applicable tax rate:

$$R_{it}(a_t^0) = \beta_0 + \beta_1 X_{it} + \beta_2(\ln(1 + \tau_{it}\mathbf{1}_{it}^{taxable}(a_t^0))) + \epsilon_{it} \tag{17}$$

We use the parameter estimates from this regression to calculate $\hat{R}_{it}(a)$ for the a given roll out a_t in the following manner:

$$\hat{R}_{it}(a_t) = \hat{\beta}_0 + \hat{\beta}_1 X_{it} + \hat{\beta}_2(\ln(1 + \tau_{it}\mathbf{1}_{it}^{taxable}(a_t))) \tag{18}$$

An alternative to representing exogenous changes in $\tilde{x}$ and $\tilde{y}$ is to simply substitute an exogenous shifter of revenue R_{it} in its place in equations 15 and 16. For example, we use population in county i as a correlate of Amazon's revenue in that county and define $\hat{x}^2$ and $\hat{y}^2$ based on the difference in the population weighted distance and sales tax of roll out a^t:

$$\hat{x}_a^2 = \sum_t \sum_i \beta^t \left[d_{it}(a_t^0)pop_{it} - d_{it}(a_t)pop_{it}\right]$$

and

$$\hat{y}_a^2 = \sum_t \sum_i \beta^t \left[(1 - \tau_{it}\mathbf{1}_{it}^{taxable}(a_t^0))pop_{it} - (1 - \tau_{it}\mathbf{1}_{it}^{taxable}(a_t))pop_{it}\right].$$

We then use $\hat{y}_a$ ($\hat{y}_a^2$) and $\hat{x}_a$ ($\hat{x}_a^2$) to define our basic instruments. We identify groups of counter-factual roll-outs that are useful in identifying the lower and upper bound cost-savings. We select two types of "experiments" that reflect the revenue/shipping distance tradeoff. First, observing Amazon choosing to enter relatively early into a high-tax state when a low-tax (but further away) option was available identifies an upper bound of the shipping cost saving parameter. To identify this upper bound, we look for counter-factual roll-outs that *increase the shipping distance* of the network to larger markets and result in higher revenues (Experiment 1). Similarly, counter-factual roll-outs that *decrease the shipping* of the network to large markets identify an lower bound on the cost savings (Experiment 2). We group counter-factual rollouts that qualify for each experiment based on their predicted revenue change into three categories defined by the magnitude of their shipping distance changes relative to the observed rollout. We consider three subgroups, which gives us six instruments for the basic moments.

1. **Experiment: Increase in shipping distance.** Let t_j be the chosen opening date of FC j and denote as S_1 the set of FCs opened at dates $t_{j'} > t_j$ in states that would lead to higher net revenue $\hat{y}_{a(j,j')} > 0$. Define as $a(j, j')$ the counter-factual sequence that swaps j with $j' \in S_1$. Group all $j' \in S_1$ into three categories defined as:

   - Group 1: $j' \in S_1$ such that $\kappa_0 < \hat{x}_{a(j,j')} \leq \kappa_1$
   - Group 2: $j' \in S_1$ such that $\kappa_1 < \hat{x}_{a(j,j')} \leq \kappa_2$
   - Group 3: $j' \in S_1$ such that $\kappa_2 < \hat{x}_{a(j,j')}$

2. **Experiment: Decrease in shipping distance.** Let t_j be the chosen opening date of FC j and denote as S_2 the set of FCs opened at dates $t_{j'} > t_j$ in states that would lead to lower net revenue $\hat{y}_{a(j,j')} < 0$. Define as $a(j, j')$ the counter-factual sequence that swaps j with $j' \in S_2$. Group all $j' \in S_2$ into three categories defined as:

   - Group 1: $j' \in S_2$ such that $-\kappa_0 \geq \hat{x}_{a(j,j')} > -\kappa_1$
   - Group 2: $j' \in S_2$ such that $-\kappa_1 \geq \hat{x}_{a(j,j')} > -\kappa_2$
   - Group 3: $j' \in S_2$ such that $-\kappa_2 > \hat{x}_{a(j,j')}$

These groups define a set of dummy variables that indicate whether a given counterfactual sequence that swaps observed opening date of FC j with that of FC j' qualifies for experiments 1 and 2, and if so, which group. These indicators constitute our six basic instruments z. We also interact these instruments with functions of $\hat{x}_a^+$ and $\hat{y}_a^+$, which are positive transformations of $\hat{x}$ and $\hat{y}$ (i.e., $\hat{x}_a^+ = \hat{x}_a - \min(\hat{x})$).

In practice, we define κ based on the distribution of $\hat{x}$. For experiment 1, κ_0 is the 10th percentile of $\hat{x}$, κ_1 is the 25th percentile and κ_3 is the 75th percentile. For experiment 2, κ_0 is the 90th percentile of $\hat{x}$, κ_1 is the 75th percentile and κ_3 is the 25th percentile.

The second approach we take is an alternative to constructing the instruments as indicator variables. Specifically we use smooth transformations of the indicator variables:

1. **Experiment: Increase in shipping distance**

$$\begin{aligned} z_a^1 &= \Big\{\Phi(\hat{y}_a/\sigma)\times\Phi(\hat{x}_a/\sigma), \Phi(\hat{y}_a/\sigma)\times\Phi(\hat{x}_a/\sigma)\times\hat{x}_a^+, \\ &\quad \Phi(\hat{y}_a/\sigma)\times\Phi(\hat{x}_a/\sigma)\times\hat{y}_a^+, \Phi(\hat{y}_a/\sigma)\times\Phi(\hat{x}_a/\sigma)\times\hat{y}_a^+\times\hat{x}_a^+\Big\} \end{aligned}$$

2. **Experiment: Decrease in shipping distance**

$$\begin{aligned} z_a^2 &= \Big\{(1-\Phi(\hat{y}_a/\sigma))\times(1-\Phi(\hat{x}_a/\sigma)), (1-\Phi(\hat{y}_a/\sigma))\times(1-\Phi(\hat{x}_a/\sigma))\times\hat{x}_a^+, \\ &\quad (1-\Phi(\hat{y}_a/\sigma))\times(1-\Phi(\hat{x}_a/\sigma))\times\hat{y}_a^+, (1-\Phi(\hat{y}_a/\sigma))\times(1-\Phi(\hat{x}_a/\sigma))\times\hat{y}_a^+\times\hat{x}_a^+\Big\} \end{aligned}$$

The parameter σ is a smoothing parameter. As σ approaches zero, the product of the two CDFs approaches the dummy variables above. The advantage of using the smooth instruments is that we use all available FC swaps, rather than arbitrarily dropping those that do not satisfy the conditions.

Estimation

We set $\beta = 0.95$ and calculate μ based on information obtained from Amazon's financial reports. Specifically, Amazon reports the total amount of revenue in the "Media" and "Electronics and Other General Merchandise" categories in North America, which is roughly equivalent to the revenue we predict from our model.[33] They also report the "Cost of Goods Sold" for all of their sales. We compute the cost of goods sold for North America by multiplying the total cost by the ratio of sales from North America to total sales. This provides us with a value for the "gross margin". However, as Amazon states, the reported cost of goods sold includes both outbound shipping costs and

[33] Revenue from Canada and Mexico is not included in the ComScore data. Therefore, these sales are accounted for through the 'multiplier' that we use to match the predictions of the model in the financial reports. See section 4 for a discussion of these multipliers.

inbound shipping costs (through wholesale prices). Recall that we assume that inbound shipping costs from suppliers to Amazon do not vary with the network of FCs. As we are estimating the outbound shipping costs, we exclude them from the gross margin by adding in the reported "Net Shipping Cost as a Percentage of Revenue" to the gross margin. This value grows over time, but we suspect this is due to an increase in Amazon's non-shipping related activities (i.e. Amazon Web Services and digital goods). We therefore set $\mu = 0.23$.

As mentioned earlier, the perturbations are created by swapping two FC opening dates. Because of FC clustering, each swap consists of moving the opening date of the earliest FC in the cluster. The swapped fixed costs are also assumed to be fixed costs (land and employees) associated with the earliest opening FC. The total number of perturbations is around 1,500 and after removing the swaps which open in the same year, it is 1,415. Of those, how many perturbations we use in estimation varies with the definition of the instruments, as can be seen in last columns of tables 17 and 19.

Finally, because a FC opening does not always coincide perfectly with changes in the tax status of a state, we consider three different assumptions about what Amazon believes entry into a state implies. The first assumption is that Amazon believes they will be forced to collect sales tax immediately upon entry into a state, which we refer to as the "No Lag" assumption. The second is that the lag between entry and a change in tax status follows a given FC. For example, if we swap a FC that opened in Pennsylvania in 2008 and resulted in sales tax liabilities starting in 2011 with one that opened in Nevada in 1999 but resulted in sales tax liabilities only in 2014, we assume that the swapped FC in Pennsylvania would now be opened in 1999 with sales taxes being collected from Pennsylvania residents beginning in 2014. We refer to this as the "FC Lag" assumption. Finally, the third assumption is that the lag is state specific, or the "Stage Lag" assumption. For the swap discussed above, this implies that the FC in Pennsylvania opens in 1999 and sales tax is collected beginning in 2002 (i.e., there is a 3 year lag as there is with the actual FC that opened in Pennsylvania in 2008).

Results

We present the estimates of θ under the "No Lag" assumption for five different definitions of $\hat{y}$ and $\hat{x}$ and for both the discrete and continuous instruments in 17. The value of the parameter can be interpreted as the net shipping cost as a percent of revenue for every 100 miles in shipping distance. Our preferred specification is the one that defines $\hat{y}$ and $\hat{x}$ based on $\hat{R}$, which results in a $\hat{\theta}$ of around -0.40. This implies that it costs Amazon about \$1.20 to ship \$100 of goods 300 miles (the average distance in 2006) and \$0.50 to ship \$100 in goods 125 miles (the average distance in 2018).

Because there is only one parameter to estimate, we are able to plot the objective function over different values of the parameter. Figure 5 displays plots for different definitions of $\hat{y}$ and $\hat{x}$. Focusing on plot (b), we can see that the objective function is relatively flat between $\hat{\theta} = -0.3$ and $\hat{\theta} = -0.6$, which, absent of computing confidence intervals, can loosely be thought of as the estimate of the bounds on the shipping cost.[34] When using definitions of $\hat{y}$ and $\hat{x}$ based on population only, which are plausibly weaker instruments, we see that the objective function is flatter, especially at the top end. Finally, using the continuous instruments appears to give us stronger identification for lower values of θ.

Tables 18 and 19 present the results for the alternative behavioral assumptions. The estimate of $\hat{\theta}$ for our preferred specification is as low as -1.3 and as high as -0.3. The latter implies that it costs Amazon about \$3.90 to ship \$100 of goods 300 miles and \$1.60 to ship \$100 in goods 125 miles.

With these estimates we can compute the total amount of money Amazon has saved (and will save) from the expansion of the network from 2006-2018. Table 14 shows that the total amount of savings is \$3.6 billion for $\hat{\theta} = -0.04$ and \$12 billion for $\hat{\theta} = -1.3$. Note that they also save a significant amount in fixed costs by not opening new FCs, which can explain why total cost savings from shipping may be less than the revenue effects from expansion based on our least conservative estimate.

[34]In future iterations of the paper, we will perform inference on our parameter estimates.

6 Conclusion

We study the trade-off Amazon faces when choosing the location of a new fulfillment center. Amazon benefits from their fulfillment centers being close to consumers for two potential reasons: first, the customer herself may value faster shipping, and second, it saves on delivery costs. At the same time, state laws dictate that Amazon must charge sales tax to consumers in most states where they have a physical presence. By raising the tax-inclusive price the consumer faces, this reduces consumers' willingness-to-pay for Amazon's services. Being close to significant population clusters also raises the fixed cost of operating fulfillment centers.

We find that Amazon indeed faces this tradeoff: consumers dislike paying taxes, meaning there be must be gains of network expansion due to faster shipping and/or reduced delivery costs. Our demand estimates indicate that consumer demand does not respond to our measures of shipping times, most likely due to the fact that expansion did not actually result in faster shipping speeds for the vast majority of consumers apart from the possibility of one-day shipping. Therefore, we find that the network expansion from 2006-2018 resulted in a loss in revenue of around \$8.9 billion dollars for Amazon, but at the same time reduced the average shipping distance from FC to consumer by around 170 miles by the end of 2018. We use a moment inequalities approach, together with the assumption that Amazon's network expansion path is optimal, to infer shipping cost savings from the observed fulfillment center network relative to alternative configurations. Results suggest that Amazon saves between \$0.40 and \$1.30 per 100 miles for every \$100 dollars of goods shipped. Therefore, the expansion of the network has resulted in between \$3 and \$12 billion in savings on shipping costs.

References

Agarwal, S., S. Chomsisengphet, T. Ho, and W. Qian (2013). Cross border shopping: Do consumers respond to taxes or prices? Working Paper.

Agarwal, S., N. Marwell, and L. McGranahan (2013). Consumption responses to temporary tax incentives: Evidence from state sales holidays. Working Paper.

Agatz, N. A., M. Fleischmann, and J. A. Van Nunen (2008). E-fulfillment and multi-channel distribution–a review. *European Journal of Operational Research 187*(2), 339–356.

Alm, J. and M. I. Melnik (2005). Sales taxes and the decision to purchase online. *Public Finance Review 33*(2), 184–212.

Anderson, E. T., N. M. Fong, D. I. Simester, and C. E. Tucker (2010). How sales taxes affect customer and firm behavior: The role of search on the internet. *Journal of Marketing Research 47*(2), 229–239.

Asplund, M., R. Friberg, and F. Wilander (2007). Demand and distance: evidence on cross-border shopping. *Journal of Public Economics 91*(1), 141–157.

Ballard, C. L. and J. Lee (2007). Internet purchases, cross-border shopping, and sales taxes. *National Tax Journal 60*(4), 711–725.

Baugh, B., I. Ben-David, and H. Park (2014). The "Amazon Tax": Empirical Evidence from Amazon and Main Street Retailers. Technical report, National Bureau of Economic Research.

Bruce, D., W. F. Fox, and L. Luna (2009). State and local government sales tax revenue losses from electronic commerce. *State Tax Notes 52*(7), 537–558.

Chetty, R., A. Looney, and K. Kroft (2009). Salience and taxation: Theory and evidence. *The American Economic Review 99*(4), 1145–1177.

De los Santos, B., A. Hortaçsu, and M. R. Wildenbeest (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *The American Economic Review 102*(6), 2955–2980.

Einav, L., J. Levin, and N. Sundaresan (2014). Sales taxes and internet commerce. *The American Economic Review 104*(1), 1–26.

Ellison, G. and S. F. Ellison (2009). Tax sensitivity and home state preferences in internet purchasing. *American Economic Journal: Economic Policy 1*(2), 53–71.

Goolsbee, A., M. F. Lovenheim, and J. Slemrod (2010). Playing with fire: Cigarettes, taxes, and competition from the internet. *American Economic Journal: Economic Policy 2*(1), 131–154.

Holmes, T. J. (2011). The diffusion of Wal-Mart and economies of density. *Econometrica 79*, 253–302.

Pakes, A., J. Porter, K. Ho, and J. Ishii (2015). Moment inequalities and their application. *Econometrica 83*(1), 315–334.

Scanlan, M. A. (2007). Tax Sensitivity in Electronic Commerce. *Fiscal Studies 28*(4), 417–436.

Smith, M. D. and E. Brynjolfsson (2001). Consumer decision-making at an internet shopbot: Brand still matters. *Journal of Industrial Economics 49*(4), 541–558.

Zheng, F. (2014). Spatial Competition and Preemptive Entry in the Discount Retail Industry. Working Paper.

A Data Appendix

Forrester Data

We purchased survey data from Forrester Research, Inc which provides information on the extent of online shopping.[35] The survey, called the "North American Technographics Online Benchmark Survey", was conducted from 2006-2007 by mail and 2010-2014 online and surveyed between 30,000 and 60,000 households in the United States and Canada. The exact content of the questions on the survey varies by year, but generally the most pertinent question for us is the one which asks the user "Have you bought anything online in the past three months?". Additionally, the survey asks for information about the age, income, race, and zip code of the respondent. The documentation of the survey is available from the authors upon request.

Calculating Expenditures

In what follows, we provide a description of the procedure to calculate the expenditures for the representative consumer in county i and year t on the three different modes of shopping.

We first use the Forrester data to estimate the extent of online shopping for a given demographic group which we will use to account for the extensive margin which we believe to be measured poorly in the ComScore data. We run the following linear probability regression:

$$Pr(D_{\iota t} = 1) = \beta_0 + \beta_1 Z_\iota + \gamma_t + \epsilon_{\iota t}$$

where Z is a set of dummy variables indicating the characteristics of respondent ι which includes the race, income, age and the census region of the respondent and γ_t is a year dummy variable. The dependent variable, $D_{\iota t}$, indicates whether or not the consumer answered yes to the question of whether or not they had purchased anything online in the past three months. We then use the estimates of the model to predict the probability that a consumer who belongs to group Z purchased something online in year t. That is, we form $\hat{p}_{zt}$ for each consumer group and year. For the years, 2008 and 2009, we use linear projections based on the predictions in 2007 and 2010.

[35] `https://www.forrester.com/home/`

Using these predicted probabilities, we calculate the Forrester adjusted expected total online expenditures for a household who belongs to group z in county i and year t:

$$e_{zit} = \frac{1}{N_{it}(z)} \sum_{h \in H_{it}(z)} \hat{p}_{zt} \tilde{e}_{ht}$$

Here, $\tilde{e}_{ht}$ is the observed expenditure of household h, $H_i(z)$ is the set of households which belong to group z in county i, and $N(z)_{it}$ is the size of this set. Importantly, we drop any consumer where expenditures are equal to 0, because this is already accounted for through the Forrester correction.

Next, we calculate group level shares for each shopping mode in each county using the unadjusted ComScore data:

$$s_{zitj} = \frac{\bar{e}_{zitj}}{\sum_j \bar{e}_{zitj}}$$

where j indicates the shopping mode and $\bar{e}_{zitj}$ is the average expenditures on mode j in county i for group z:

$$\bar{e}_{zitj} = \frac{1}{N_{it}(z)} \sum_{h \in H_{it}(z)} \tilde{e}_{hitj}$$

Using census weights for demographic group z, denoted w_{zi}, and the adjusted ComScore data, we create the total online expenditures for a representative consumer in a given county:

$$e_{it} = \sum_z w_{zi} e_{zit}$$

and the representative shares:

$$s_{itj} = \sum_z w_{zi} s_{zitj}$$

Finally, to get the expenditures for the representative consumer across each mode, we combine the previous two calculations:

$$e_{itj} = s_{itj} \tilde{e}_{it}$$

B Tables and Figures

B.1 Tables

Table 1: Expansion of fulfillment center network

Year	# of Centers	# States with with Fulfillment Center	# States with Sales Tax Liability	Ave Size of FCs (1,000 Sq ft)	Ave Employees of FCs
2006	8	6	4	544	504
2007	9	7	4	522	484
2008	12	10	5	487	452
2009	17	10	5	570	529
2010	17	10	5	570	529
2011	24	10	5	657	610
2012	32	12	8	718	666
2013	41	14	16	765	709
2014	48	14	23	707	656
2015	54	16	26	665	616
2016	90	27	27	562	521
2017	101	28	28	552	511
2018	104	28	28	558	517

Notes: The number of states where Amazon purchases are subject to sales tax exceeds the number of states with a FC due to states negotiating for sales taxes being collected immediately after agreeing to let Amazon build a FC in the state, even if there is a delay between the time of agreement and the actual opening of the warehouse.

Table 2: comScore Sample

Year	Households	States (%)	Counties (%)
2006	86,405	49 (100)	2,872 (91)
2007	89,412	49 (100)	2,865 (91)
2008	56,722	49 (100)	2,725 (87)
2009	55,585	49 (100)	2,661 (85)
2010	53,788	49 (100)	2,621 (83)
2011	62,849	49 (100)	2,695 (86)
2012	54,794	49 (100)	2,623 (83)
2013	46,216	49 (100)	2,437 (78)

Notes: These figures exclude Alaska and Hawaii.

Table 3: Household Purchasing

Year	Online Expenditure	Online Transactions	% Zero Expenditure	Adjusted Expenditure	Adjusted Transactions	% Offline Shoppers Only
2006	$239	2.4	51.8%	$242	2.4	55.7%
2007	$254	2.5	52.0%	$242	2.4	60.8%
2008	$196	2.0	60.0%	$258	2.6	-
2009	$141	1.4	67.9%	$268	2.7	-
2010	$125	1.4	68.6%	$275	2.9	32.1%
2011	$131	1.4	69.7%	$327	3.4	23.0%
2012	$151	1.8	64.0%	$311	3.6	23.9%
2013	$120	1.7	65.2%	$294	4.0	15.5%

Notes: The percent offline shoppers only denotes the share of respondents who answered no to the question whether they had shopped online in the previous three months in the Forrester Technographics Survey.

Table 4: Tax Rate Variation

Year	Ave Tax Across All Counties	St. Dev Tax Across All Counties	Ave St. Dev Tax Across Counties within State	% of Counties with Tax Change
2006	6.28%	1.58%	0.33%	32%
2007	6.30%	1.59%	0.33%	28%
2008	6.34%	1.58%	0.32%	61%
2009	6.33%	1.55%	0.31%	65%
2010	6.41%	1.61%	0.31%	57%
2011	6.44%	1.60%	0.33%	52%
2012	6.44%	1.60%	0.34%	52%
2013	6.47%	1.60%	0.34%	33%
2014	6.50%	1.61%	0.35%	30%

Notes: County level tax rates are calculated as a sum of state, county and local sales tax.

Table 5: Characteristics of Counties with FCs

FCs Open by	Households in State (M)	Weighted Average of Median Income in State	Weighted Average of Wage in County annually)	Weighted Average of Property Value in County ($ per Sq Ft)
2006	1.35	21,800	27,600	19
2007	2.59	22,000	27,300	25
2008	2.93	21,600	27,100	25
2009	2.58	21,300	24,700	17
2010	2.58	21,300	24,500	15
2011	2.58	21,300	24,400	15
2012	2.58	21,300	24,000	15
2013	2.51	21,100	23,800	14
2014	3.25	21,300	23,600	12
2015	3.25	21,300	23,400	12
2016	3.25	21,300	25,600	585
2017	3.41	21,200	23,900	232
2018	3.52	21,100	23,500	214
All US	2.25	21.9	21.3	74

Notes: The large change in property value from 2015 to 2016 is due to the opening of Prime Now FCs near Los Angeles and San Francisco.

Table 6: Transaction Level Demand Estimates

Variable name	(1)	(2)	(3)	(4)	(5)	(6)
Tax Elasticity	-0.163**	-0.165**	-0.158**	-0.171**	-0.172**	-0.171**
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Log Distance		0.001				
		(0.000)				
1 or 2 Day Priority			0.020**			
			(0.003)			
1, 2, or 3 Day Package				-0.004**		
				(0.001)		
1 Day First Class					-0.005**	
					(0.001)	
1, 2, or 3 Day Standard						-0.004**
						(0.001)
Obs	2,153,810	2,153,810	2,153,810	2,153,810	2,153,810	2,153,810
R-Sq	0.364	0.364	0.364	0.364	0.364	0.364

** 1% * 5%. Notes: Regressions include dummy variables for product category, year and county. Shipping times are grouped into "long" and "short" with long being the excluded category.

Table 7: Diff-in-Diff Demand Estimates

Variable name	(1)	(2)	(3)	(4)	(5)	(6)
Taxed Dummy	-0.104	-0.106	-0.100	-0.107	-0.113	-0.107
	(0.058)	(0.059)	(0.059)	(0.059)	(0.059)	(0.059)
Log Distance		0.009				
		(0.033)				
1 or 2 Day Priority			0.083			
			(0.102)			
1, 2, or 3 Day Package				-0.010		
				(0.040)		
1 Day First Class					-0.049	
					(0.050)	
1, 2, or 3 Day Standard						-0.010
						(0.040)
Obs	12,825	12,825	12,825	12,825	12,825	12,825
R-Sq	0.618	0.618	0.618	0.618	0.618	0.618

** 1% * 5%.Notes: Regressions include dummy variables for year and county. Shipping times are grouped into “long” and “short” with long being the excluded category.

Table 8: Taxed and Non-Taxed Competitors

Sales Rank	Taxed	Non-Taxed
1	walmart.com	dell.com
2	jcpenney.com	qvc.com
3	staples.com	yahoo.net
4	victoriassecret.com	hsn.com
5	officedepot.com	yahoo.com
6	bestbuy.com	quillcorp.com
7	apple.com	overstock.com
8	target.com	ebay.com
9	sears.com	orientaltrading.com
10	costco.com	zappos.com
Total (%)	192 (34)	375 (66)

Notes: Table displays top domains which we define as taxed and non-taxed.

Table 9: CES Demand Estimates

Variable name	(1)	(2)	(3)	(4)	(5)	(6)
Tax Elasticity	-1.337**	-1.341**	-1.354**	-1.341**	-1.375**	-1.341**
	(0.450)	(0.454)	(0.450)	(0.451)	(0.450)	(0.451)
Log Distance		-0.001				
		(0.015)				
1 or 2 Day Priority			-0.092			
			(0.054)			
1, 2, or 3 Day Package				-0.005		
				(0.029)		
1 Day First Class					-0.083**	
					(0.032)	
1, 2, or 3 Day Standard						-0.005
						(0.029)
Obs	43,117	43,117	43,117	43,117	43,117	43,117
R-Sq	0.154	0.154	0.154	0.154	0.154	0.154

** 1% * 5%. Notes: Presented are the estimates of tax sensitivity and the effect of shipping speeds. Regressions include mode/year dummies along with mode level effects of local demographics. Shipping times are grouped into "long" and "short" with long being the excluded category.

Table 10: Demand Estimates (Mode Fixed Effects)

Year	Amazon	Other Taxed Online	Non-taxed Online
2006	-1.601*	0.535	
	(0.711)	(0.681)	
2007	-1.576*	0.538	-0.069
	(0.714)	(0.684)	(0.078)
2008	-1.043	0.895	0.157
	(0.715)	(0.685)	(0.084)
2009	-0.900	0.848	0.070
	(0.716)	(0.686)	(0.087)
2010	-0.776	0.687	-0.088
	(0.717)	(0.686)	(0.090)
2011	-0.399	0.856	-0.123
	(0.716)	(0.686)	(0.087)
2012	-0.023	0.708	-0.315**
	(0.715)	(0.686)	(0.086)
2013	0.225	0.566	-0.208*
	(0.717)	(0.687)	(0.094)

** 1% * 5%. Notes: Presented are the estimates the mode/year effects from specification (1) in table 9.

Table 11: Demand Estimates (Demographics)

Varaible	Amazon	Mode 2	Mode 3
Total Offline Competitors	-0.002	-0.002	-0.005
	(0.015)	(0.015)	(0.015)
Large Offline Competitors	-0.008	-0.012	-0.002
	(0.007)	(0.007)	(0.007)
Income	-0.002	-0.096	
	(0.066)	(0.063)	
% Pop Black	-0.342	0.348	
	(0.247)	(0.221)	
% Pop White	0.056	0.239	
	(0.237)	(0.231)	
% Pop Asian	1.735*	0.045	
	(0.784)	(0.771)	

** 1% * 5%. Notes: Presented are the estimates the mode level demographic effects from specification (1) in table 9. The offline competition variables are measured in houndreds of establishments.

Table 12: Alternative Demand Models

Variable name	(1)	(2)	(3)	(4)	(5)	(6)
Tax Elasticity	-1.520**	-1.720**	-1.858**		-1.337**	-1.373**
	(0.267)	(0.576)	(0.536)		(0.450)	(0.390)
Tax Elasticity (Amazon)				-1.243**		
				(0.482)		
Tax Elasticity (Mode 3)				-1.426		
				(0.801)		
Obs	43,117	53,914	29,584	43,117	43,298	43,117
R-Sq	0.302	0.145	0.083	0.153	0.145	0.191
Regression	A-Weights	Zeros	2008-2013	Individual Tax Effect	No Forr Adjustment	No Pop Weights

** 1% * 5%. Notes: Presented are the robustness estimates excluding any shipping speed effect. The results when including these effects are similar to that of the base regressions.

Table 13: Predicted Revenue

Year	Amazon ($B)	Mode 2 ($B)	Mode 3 ($B)	Ave Distance (miles)	# Counties Taxed
2006	5.60	20.40	24.13	296.61	317
2007	7.77	27.74	30.57	292.06	317
2008	9.78	28.77	30.48	236.45	379
2009	12.27	29.90	30.35	226.97	379
2010	17.87	32.88	33.41	242.85	380
2011	25.27	37.91	31.35	236.01	380
2012	32.45	28.94	22.95	222.56	755
2013	40.78	25.09	23.41	207.05	1,194
2014	52.82	28.57	22.02	181.86	1,670
2015	65.08	27.11	20.42	158.24	1,943
2016	79.41	25.71	18.98	133.98	1,989
2017	95.88	24.38	17.69	127.55	2,053
2018	114.73	23.13	16.53	125.88	2,053
Total	559.72	360.54	322.28		

Notes: Presented is the total revenue based on the predictions of the model. Future revenues are calculated based on projects of the mode/year effects and knowledge of the future network of FCs. Average distance is the average distance to households in the US.

Table 14: Revenue and Cost Effects of Expansion

Year	Δ Revenue (\$M)	Δ Ave Distance (miles)	Cost Saving (\$M, $\hat{\theta} = 0.4$)	Cost Saving (\$M,$\hat{\theta} = 1.3$)
2007	6.42	4.55	1.32	4.29
2008	30.94	60.16	27.00	87.73
2009	36.75	69.64	40.51	131.65
2010	57.85	53.76	44.55	144.80
2011	85.71	60.60	70.83	230.19
2012	305.42	74.05	111.18	361.34
2013	558.87	89.56	188.90	613.91
2014	912.33	114.75	303.37	985.94
2015	1,256.75	138.37	450.10	1,462.84
2016	1,552.41	162.63	650.58	2,114.39
2017	1,874.66	169.06	812.43	2,640.40
2018	2,242.83	170.74	982.87	3,194.34
Total	8,925.85	1,167.87	3,683.65	11,971.85

Notes: Revenue is predicted assuming that tax status of a county does not change from 2006 onwards. Additionally, shipping distance remains fixed. The cost saving estimates come from the estimates of the cost side model

Table 15: Effects of Moving FC to Highly Populated Area

Open Year	Correlation of Tax and Population	# of FCs Opened Moved to New State (Tot FCs)	Ave Distance Moved (Miles)	Ave Revenue Change from Moving ($M)
<= 2006	0.35	7 (8)	252	-138
2007	0.34	1 (1)	79	450
2008	0.34	3 (3)	129	-76
2009	0.38	2 (2)	229	-113
2010	0.38	-	-	-
2011	0.35	3 (3)	87	180
2012	0.33	3 (4)	167	74
2013	0.34	2 (5)	80	453
2014	0.34	0 (3)	116	0
2015	-	0 (2)	222	0
2016	-	4 (23)	44	232
2017	-	2 (4)	102	47
2018	-	0 (1)	216	0

Notes: Calculations based on moving a “cluster” of FCs which are within 20 miles of one another in the same state.

Table 16: Examples of Moving FC to Highly Populated Area

Year Built	Observed Location	Highest Populated Area	Distance	Revenue Change ($M)
2006	New Castle, DE	New York, NY	115	-495
2006	Lyon, NV	Orange, CA	412	-318
2006	Fayette, KY	Cook, IL	311	-212
2008	Hillsburough, NH	Suffolk, MA	39	-95
2008	Lake, IN	Cook, IL	23	-124
2011	York, PA	DC	87	411
2012	Lexington, SC	Pinellas, FL	427	-43
2012	Hamilton, TN	DeKalb, GA	102	168
2016	Pierce, WA	Multnomah, OR	110	432
2016	Baltimore City, MD	DC	35	110

Notes: Calculations based on moving a "cluster" of FCs which are within 20 miles of one another in the same state..

Table 17: Estimates of Cost Saving (No Lag)

$\hat{y}$	$\hat{x}$	Instruments	$\hat{\theta}$	S1 Size	S2 Size	Total
$R - F$	$R * d$	Discrete	-0.619	594	225	819
		Continuous	-0.473			1415
$\hat{R} - F$	$\hat{R} * d$	Discrete	-0.366	811	165	976
		Continuous	-0.400			1415
$\hat{R} - F$	$Pop * d$	Discrete	-0.231	734	171	905
		Continuous	-0.276			1415
$(1 - \tau) * Pop$	$\hat{R} * d$	Discrete	-0.309	918	162	1080
		Continuous	-0.332			1415
$(1 - \tau) * Pop$	$Pop * d$	Discrete	-0.158	840	172	1012
		Continuous	-0.208			1415

Notes: Parameter estimates are the shipping cost per 100 miles per $100 of revenue.

Table 18: Estimates of Cost Saving (FC Lag)

$\hat{y}$	$\hat{x}$	Instruments	$\hat{\theta}$	S1 Size	S2 Size	Total
$R - F$	$R * d$	Discrete	-0.596	659	277	936
		Continuous	-0.387			1415
$\hat{R} - F$	$\hat{R} * d$	Discrete	-0.476	779	205	984
		Continuous	-0.307			1415
$\hat{R} - F$	$Pop * d$	Discrete	-0.271	544	201	745
		Continuous	-0.236			1415
$(1 - \tau) * Pop$	$\hat{R} * d$	Discrete	-0.045	1089	217	1306
		Continuous	-0.080			1415
$(1 - \tau) * Pop$	$Pop * d$	Discrete	-0.018	844	187	1031
		Continuous	-0.058			1415

Notes: Parameter estimates are the shipping cost per 100 miles per $100 of revenue.

Table 19: Estimates of Cost Saving (State Lag)

$\hat{y}$	$\hat{x}$	Instruments	$\hat{\theta}$	S1 Size	S2 Size	Total
$R - F$	$R * d$	Discrete	-0.969	651	334	985
		Continuous	-0.692			1415
$\hat{R} - F$	$\hat{R} * d$	Discrete	-1.326	765	369	1134
		Continuous	-0.758			1415
$\hat{R} - F$	$Pop * d$	Discrete	-0.583	426	267	693
		Continuous	-0.510			1415
$(1 - \tau) * Pop$	$\hat{R} * d$	Discrete	-0.309	918	162	1080
		Continuous	-0.332			1415
$(1 - \tau) * Pop$	$Pop * d$	Discrete	-0.158	840	172	1012
		Continuous	-0.208			1415

Notes: Parameter estimates are the shipping cost per 100 miles per $100 of revenue.

B.2 Figures

Figure 1: Hirshman-Herfindahl Indices in Online Retail, 2006-2013

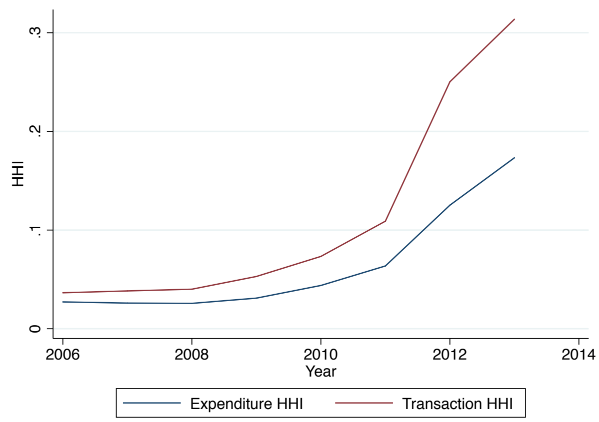

Notes: HHI is based on online sales from product categories which Amazon sells.

Figure 2: Amazon's fulfillment center Network and Number of Households

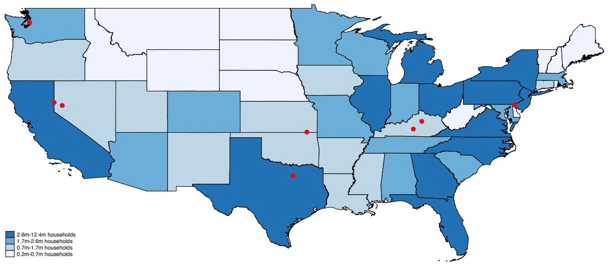

(a) 2006

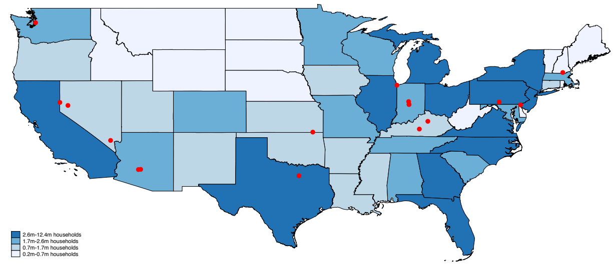

(b) 2010

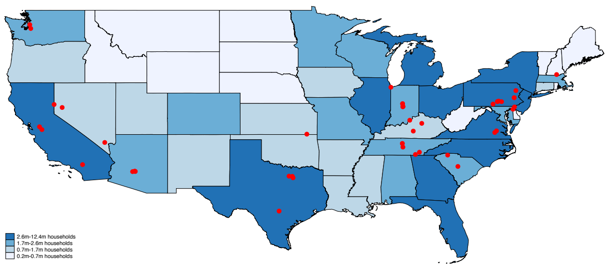

(c) 2014

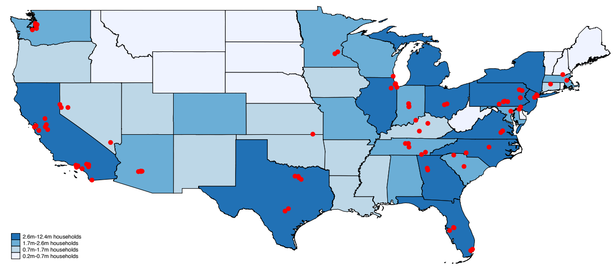

(d) 2018

(e) Wal-Mart

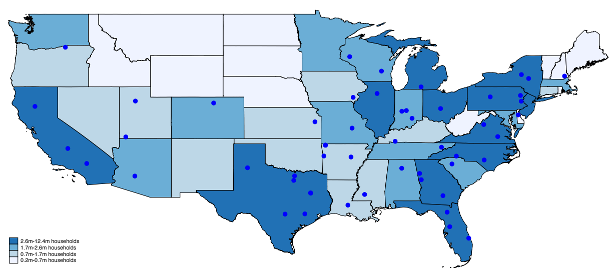

Notes: Figures include all FCs except for Amazon Fresh, returns centers and sortation centers. Shading is based on the number of households in the given state.

Figure 3: Amazon's fulfillment center Network and Average Tax Rates

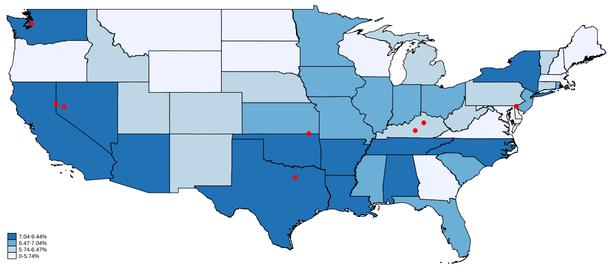

(a) 2006

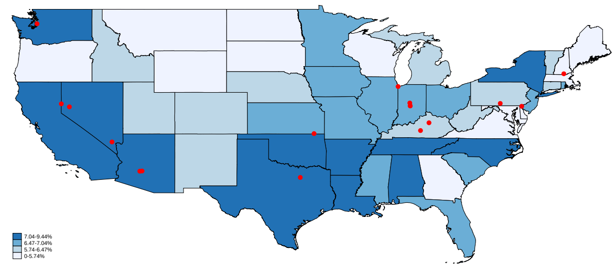

(b) 2010

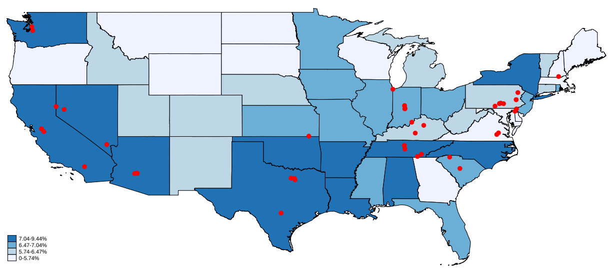

(c) 2014

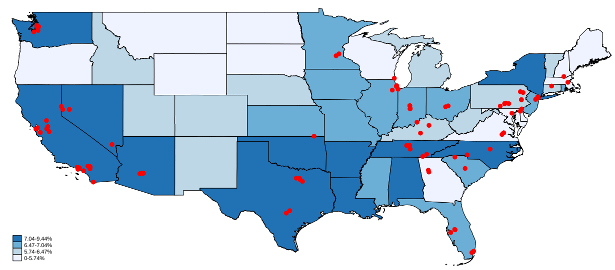

(d) 2018

(e) Wal-Mart

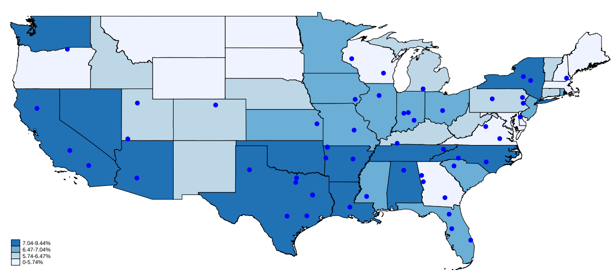

Notes: Figures include all FCs except for Amazon Fresh, returns centers and sortation centers. Shading is based on the state's average tax rate.

Figure 4: Market Share Dynamics

(a) Number of DCs

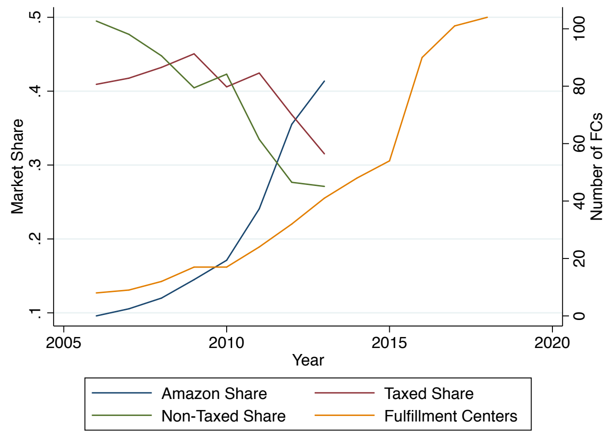

(b) Number of States with DCs

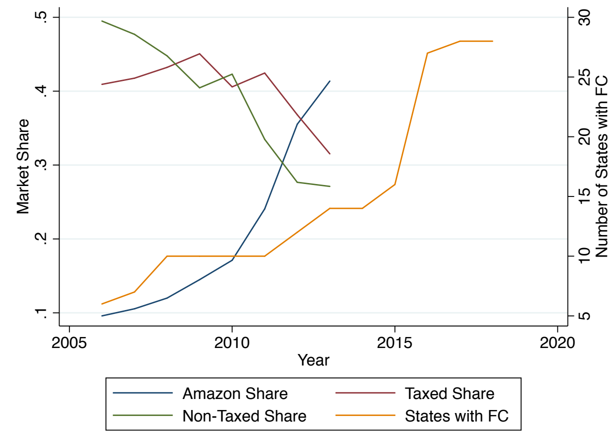

(c) Number of Taxed States

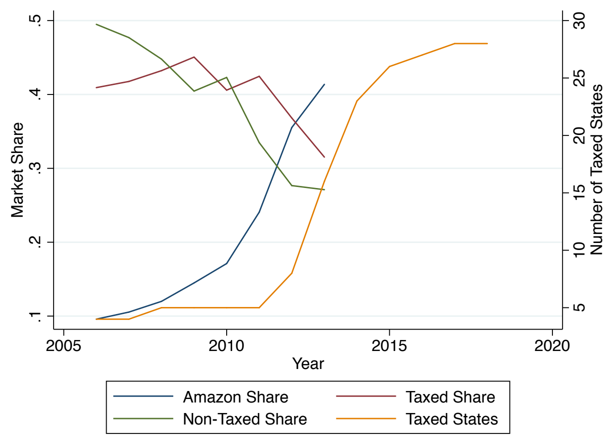

(d) Distances to DC

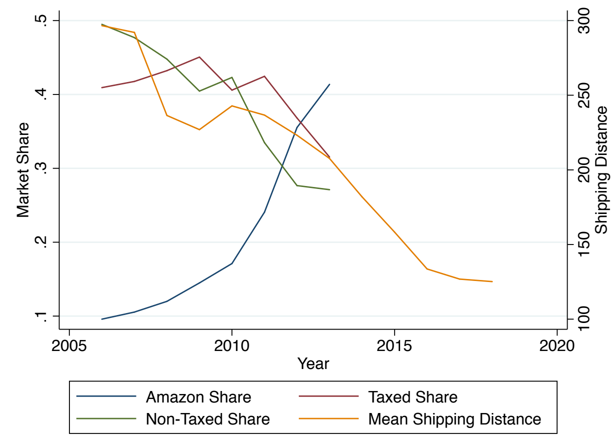

Notes: Shares are based on revenue and are conditionally on shopping online.

Figure 5: Estimation Results (No Lag)

(a) $\hat{y}(R)$ and $\hat{x}(R * d)$

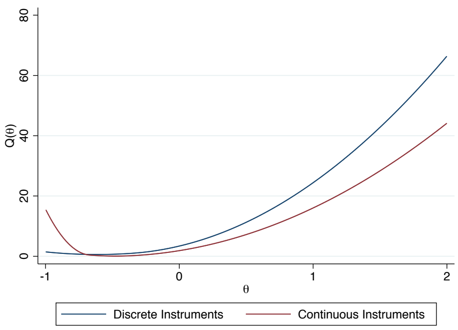

(b) $\hat{y}(\hat{R})$ and $\hat{x}(\hat{R} * d)$

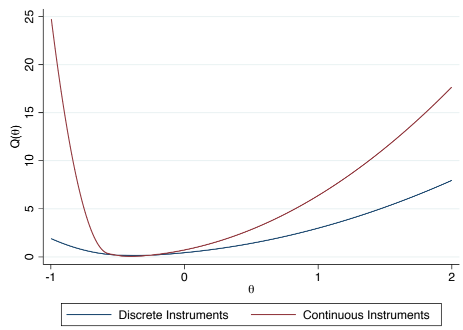

(c) $\hat{y}(\hat{R})$ and $\hat{x}(Pop * d)$

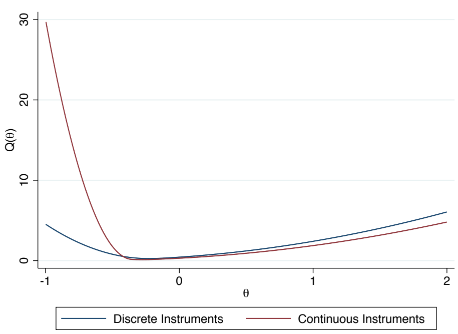

(d) $\hat{y}((1 - \tau) * Pop)$ and $\hat{x}(Pop * d)$

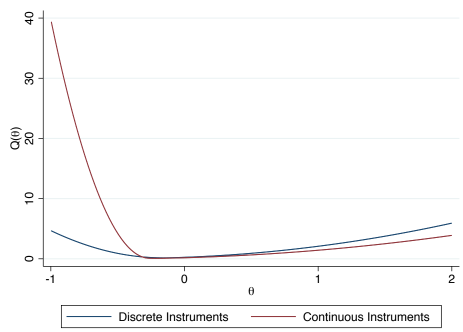

Notes: Presented is the value of the objective function over different values of the parameter.