American Economic Review 2014, 104(1): 1–26
http://dx.doi.org/10.1257/aer.104.1.1

Sales Taxes and Internet Commerce†

By Liran Einav, Dan Knoepfle, Jonathan Levin, and Neel Sundaresan*

We estimate the sensitivity of Internet retail purchasing to sales taxes using eBay data. Our first approach exploits the fact that a seller's location—and therefore the applicable tax rate—is revealed only after a buyer has expressed interest in an item. We document how adverse tax "surprises" reduce the likelihood of purchase and shift subsequent purchases toward out-of-state sellers. We then use more aggregated data to estimate that every one percentage point increase in a state's sales tax increases online purchases by state residents by almost 2 percent, while decreasing their online purchases from state retailers by 3–4 percent. (*JEL* H71, L81, L86)

Internet retail amounts to well over a hundred billion dollars annually in the United States and accounts for a growing share of overall retail commerce (US Census Bureau 2011).[1] The majority of Internet transactions occur across state lines, with striking tax consequences. While online sellers located in a particular state must collect sales tax on in-state sales, states currently cannot compel out-of-state sellers to collect tax on sales to state residents. Instead, resident consumers are obligated to pay an equivalent use tax, but enforcement is sufficiently lax that cross-state Internet sales generally go untaxed.[2] As a result, even conservative guesses about purchasing elasticities suggest that taxes may play a significant role in shaping the geography and dynamics of online retail trade.

Recently, the tax treatment of Internet commerce has generated considerable attention.[3] Sales and use taxes account for more than 30 percent of state tax revenues.

*Einav: Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305-6072, and NBER (e-mail: leinav@stanford.edu); Knoepfle: Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305-6072 (e-mail: knoepfle@stanford.edu); Levin: Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305-6072, and NBER (e-mail: jdlevin@stanford.edu); Sundaresan: eBay Research Labs, 2065 East Hamilton Avenue, San Jose, CA 95125 (e-mail: nsundaresan@ebay.com). We are grateful to Glenn Ellison, Amy Finkelstein, Austan Goolsbee, three anonymous referees, and many seminar participants for useful comments. We appreciate financial support from the Alfred P. Sloan Foundation, the National Science Foundation, the Stanford Institute for Economic Policy Research, and the Toulouse Network on Information Technology. Data access for this study was obtained under a consulting agreement between the Stanford authors (Einav, Knoepfle, and Levin) and eBay Research Labs.

† Go to http://dx.doi.org/10.1257/aer.104.1.1 to visit the article page for additional materials and author disclosure statement(s).

[1] For this and any other number mentioned in the text, online Appendix B provides a detailed reference of the source or a detailed explanation of the calculation that gives rise to the number.

[2] Varian (2000) provides useful background on the tax treatment of Internet commerce. A key Supreme Court decision in 1992 found that, absent explicit federal legislation, the Commerce Clause does not allow states to compel sellers without presence (or "nexus") in the state to collect use tax on sales to state residents (*Quill Corp. v. North Dakota,* 504 U.S. 298 (1992)). About half of the states with use taxes ask taxpayers to report use tax obligations on individual income tax returns, but this effort is largely unsuccessful. Less than 2 percent of taxpayers report any use tax in states with this type of self-reporting (Manzi 2010).

[3] As an indication of popular interest, a Google News search for Internet OR online OR e-commerce "sales tax" returned more than a thousand articles published in the first two months of 2012. Many of these articles discuss Internet sales taxes in relation to state budgets.

Foregone taxes on Internet sales could amount to $10 billion a year, and this number is likely to grow (Maguire 2013). Nevertheless, empirical evidence that might inform a discussion about Internet taxation remains rather limited, despite some notable efforts that we discuss below.

In this paper, we provide estimates of how consumers shift their purchasing across states and between offline and online retail in response to state sales taxes. These estimates are based on large-scale and recent data, aggregated to provide measures of online purchasing at the zip code level, as well as state-to-state online trade flows, over time. To complement this analysis, we also use granular browsing data to show how individuals react to sales tax at the item level, how tax sensitivity varies across items and types of consumers, and how consumers substitute to other items when faced with the application of online sales tax. Taken together, these two analyses provide aggregate estimates that may be relevant for policy, as well as micro-level results that illuminate the underlying individual behavior.

Our study uses data from eBay's online marketplace. In the United States, eBay's marketplace accounts for roughly 11–13 percent of Internet retail commerce, or around $30 billion annually. The marketplace is large and diverse, with a huge array of sellers and product categories and millions of buyers. We take advantage of this size and diversity to observe buyers choosing across sellers located in different states, with correspondingly different tax treatments and changes in those treatments over time, in order to estimate the effect of sales taxes on purchasing behavior. Although our data is limited to a single platform, its overall market share is sufficiently robust and its population of shoppers fairly representative that our analysis hopefully provides insight extending more broadly across online retail.[4]

Our estimates rely on three sources of sales tax variation. The first is the difference, for online buyers, between in-state purchases that are taxed and out-of-state purchases that are not. Of course, a direct comparison of intrastate and interstate purchase propensities may understate the effect of taxes if consumers have preference for their "home state" goods or sellers. One way to address this is to use variation across states in sales tax rates, and compare the relative intrastate purchase propensity across low tax and high tax states. There is considerable rate variation: as of January, 2010, state sales taxes ranged from zero (in Alaska, Delaware, Montana, New Hampshire, and Oregon) to 7 percent or more (California, Indiana, Mississippi, New Jersey, Rhode Island, and Tennessee). The variation becomes even greater after accounting for county and local sales taxes. Figure 1 shows the cross-sectional variation in sales and use tax rates by state and county. Finally, a third source of identifying variation comes from the frequent changes in state and local tax rates. Figure 2 shows states and counties in which there was a change in the sales tax rate during our observation period, 2008–2010.

We start by using detailed browsing data to document whether and how consumers respond to sales taxes at the item level. Our approach exploits the fact that most consumers shopping on eBay only observe a seller's location, and hence the relevant sales tax treatment, after they click on a listed item. We use data from millions of such "surprises" to estimate the tax sensitivity of purchasing conditional on

[4] Firms such as Alexa and Quantcast provide demographic information on the users of large websites. See online Appendix A for further details as to the size and representativeness of eBay users.

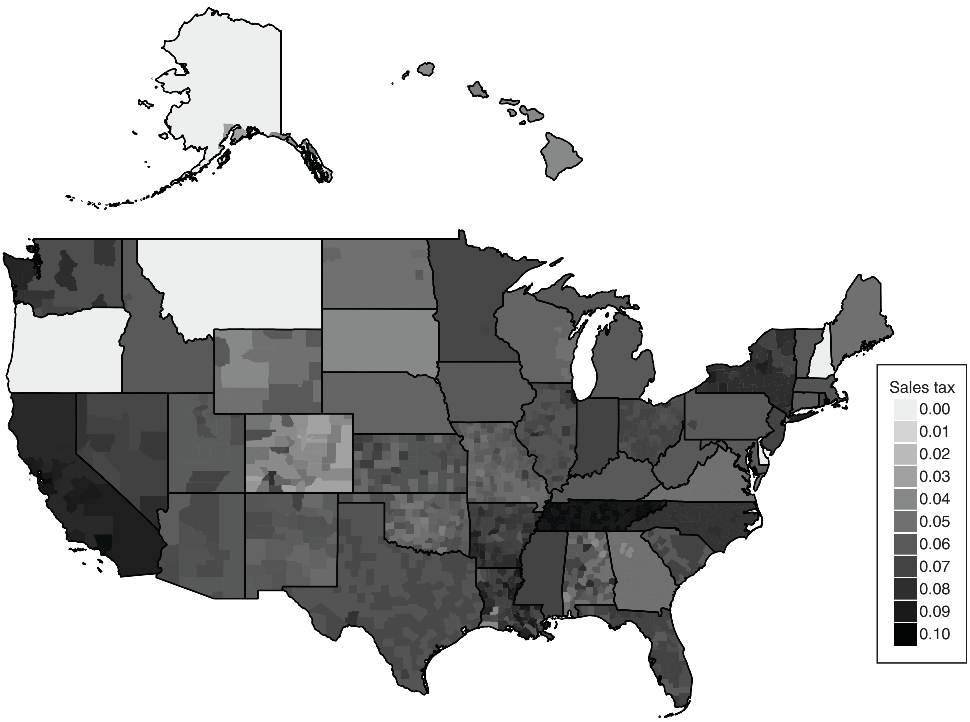

Figure 1. Cross Section Variation in Sales Tax Rates

Notes: Map shows the sales tax rate in each county as of January 1, 2010 (in the middle of our observation period of 2008–2010). The (population weighted) average tax rate in the United States that day was 7.25 percent with a (population weighted) standard deviation of 1.74 percent (for population, we use the 2000 census).

being interested in a given item. This approach allows us to control tightly for the preferences of buyers and the desirability of items located in different states. We estimate that on average, the application of a 10 percent sales tax reduces purchases by 15 percent among buyers who have clicked on an item. The effect is greater for more "commodity" type product categories, such as electronics, cell phones, and computers. We also use the browsing histories to show how consumers substitute toward alternative items following an adverse tax surprise.

This granular analysis provides a strong indication that taxes play an important role in affecting consumer decisions at the item level. However, changes in state tax rates, or in national sales tax policy, would affect the pricing of many items simultaneously. To get closer to Internet-wide elasticities, we use more aggregated data to estimate how state and local sales taxes affect both the total amount of online purchasing and propensity for online shoppers to purchase out-of-state to avoid sales tax. This exercise is closer to earlier work on Internet tax sensitivity, so that our main contribution (in this part of the paper) is improved data and the use of tax changes as well as cross-sectional differences in tax levels. We consider an econometric specification based on a constant elasticity (CES) model of online purchasing that allows us to map tax sensitivities into substitution parameters governing choices between online goods, and the choice to purchase online. Using different sources of tax variation, we estimate that a 1 percentage point increase in a state's sales tax leads to an

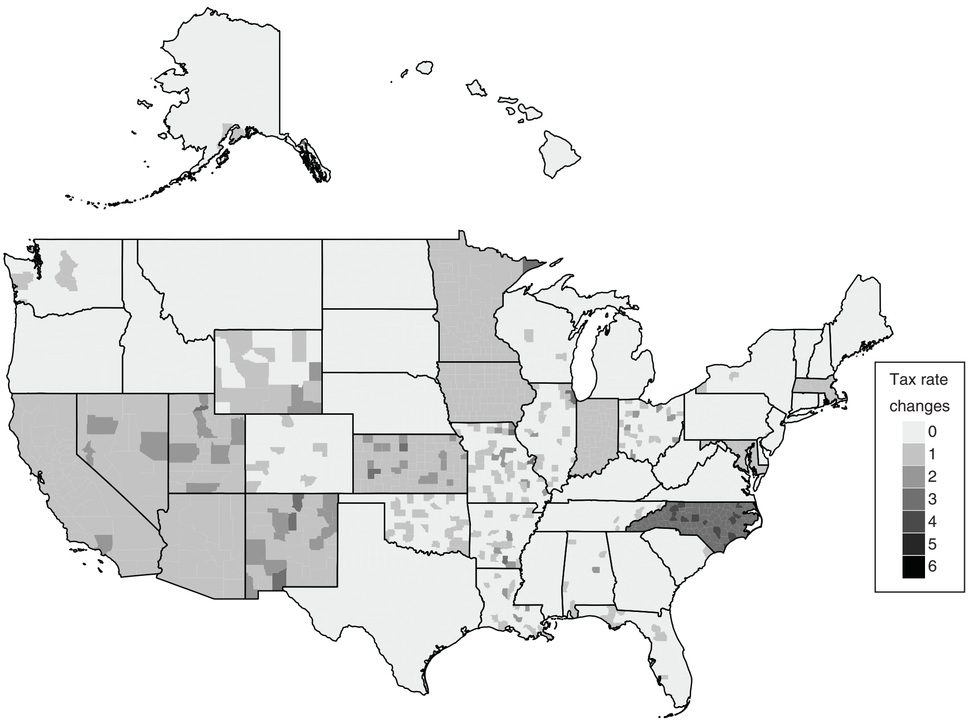

Figure 2. Time Series Variation in Sales Tax Rates

Notes: Map shows the number changes in sales tax rate in each county during our observation period (2008–2010). During this period, 35.6 percent of the United States population has been exposed to at least one change in tax rate (31.6 percent to at least one tax rate increase; 4.1 percent to at least one tax rate decrease). Conditional on a change, the (population weighted) average change (in absolute values) was 0.73 percent with a (population weighted) standard deviation of 0.38 percent (for population, we use the 2000 census).

increase of just under 2 percent in online purchasing, and a 3–4 percent decrease in the volume of online purchases from home-state sellers. We discuss the implications of these estimates for current debates about Internet sales tax policy.

Existing work on sales taxes and Internet commerce dates back to the influential work of Goolsbee (2000a, b). Using data from a 1997 Forrester Research survey, Goolsbee looked at whether respondents in high tax areas were more likely to have made an online purchase. His main estimate, that up to 24 percent of online purchasers would not have purchased online if interstate transactions were taxed, is about twice as large as ours. Later studies by Alm and Melnik (2005), Ballard and Lee (2007), and Scanlan (2007) performed a similar exercise using questions in the 1997 and 2001 Current Population Surveys. The former two studies estimate tax sensitivities at most a fourth as large as that of Goolsbee, while the latter suggests there is minimal tax sensitivity in low tax jurisdictions but very substantial sensitivity in high tax areas. Apart from using surveys that pre-date the widespread use of the Internet, one limitation of these studies is that their key dependent variable is very coarse; the authors effectively project a yes/no indicator of e-commerce participation on home-state sales tax and household characteristics.

Other studies have taken a more targeted approach using data for a particular retailer or product. Ellison and Ellison (2009) examine detailed data on the sale of computer memory modules by a retailer located in California. Using price search data, they estimate that consumers searching for certain memory modules are highly price-sensitive, with price elasticities on the order of −50 and tax-price elasticities on the order of −10. They also use data on the retailer's distribution of sales across states to estimate how sales vary with (offline) tax rates. They find that states with a one percentage point higher tax rate have about 6 percent more purchases from the retailer, but caution that their controls may not adequately isolate tax effects from other cross-state differences. Smith and Brynjolfsson (2001), Anderson et al. (2010), and Goolsbee, Lovenheim, and Slemrod (2010) also find relatively high tax sensitivities for specific types of products, namely online books, clothing, and cigarettes.

Finally, in an interesting paper that relates closely to the first half of our analysis in Section II, Hortaçsu, Martínez-Jerez, and Douglas (2009) use a sample of eBay transactions collected between February and May 2004 to estimate a gravity model of cross-state trade flows. They focus mainly on the relationship between trade volume and distance, but one of their specifications accounts for the sales tax on home-state transactions. Their results indicate that, holding online expenditures fixed, a one percentage point increase in state sales tax decreases same-state online purchases by 10 percent or more, about twice the magnitude of the effect we estimate. We discuss their estimates in more detail below.

I. Individual Responses to Tax "Surprises"

Our first approach to estimating consumer sensitivity to online sales tax uses an item-level empirical strategy. We exploit a particular feature of the search process on eBay, namely that buyers observe seller locations and the sales tax they will be charged only after they click through to an item page. Prior to clicking on a specific search result, buyers may have an expectation as to whether the seller is located in their same state, in which case sales tax is due, or not, in which case the transaction is effectively tax-free. Once a buyer reaches an item page, he or she can observe the seller location and eventually the exact effective sales tax. In what follows, we use data on consumer browsing sessions to identify millions of these "surprises" and estimate an average item-level sensitivity to sales tax.

A. *Research Design*

Consumers shopping on eBay see listings displayed on search results pages, which they can reach by browsing the site or entering search queries. Figure 3A displays a typical search results page. Each result contains a thumbnail picture of the item, a short description, its pre-tax price (or the current high bid if the sale is by auction), and the time until the listing expires. By default, listings in search results are ranked by relevance (determined by eBay's "best match" algorithm); users also can sort results by pre-tax price or listing expiration date. Seller location (and hence sales tax) is not displayed and is not factored into the sort order unless buyers explicitly specify a local search, which is very uncommon. Indeed, the only information about sellers on the default search results page are flags indicating that particular sellers are "top rated."

Potential buyers click on a search result to learn more or make a purchase. A click reveals an item page (Figure 3B) that contains more details, including the seller's

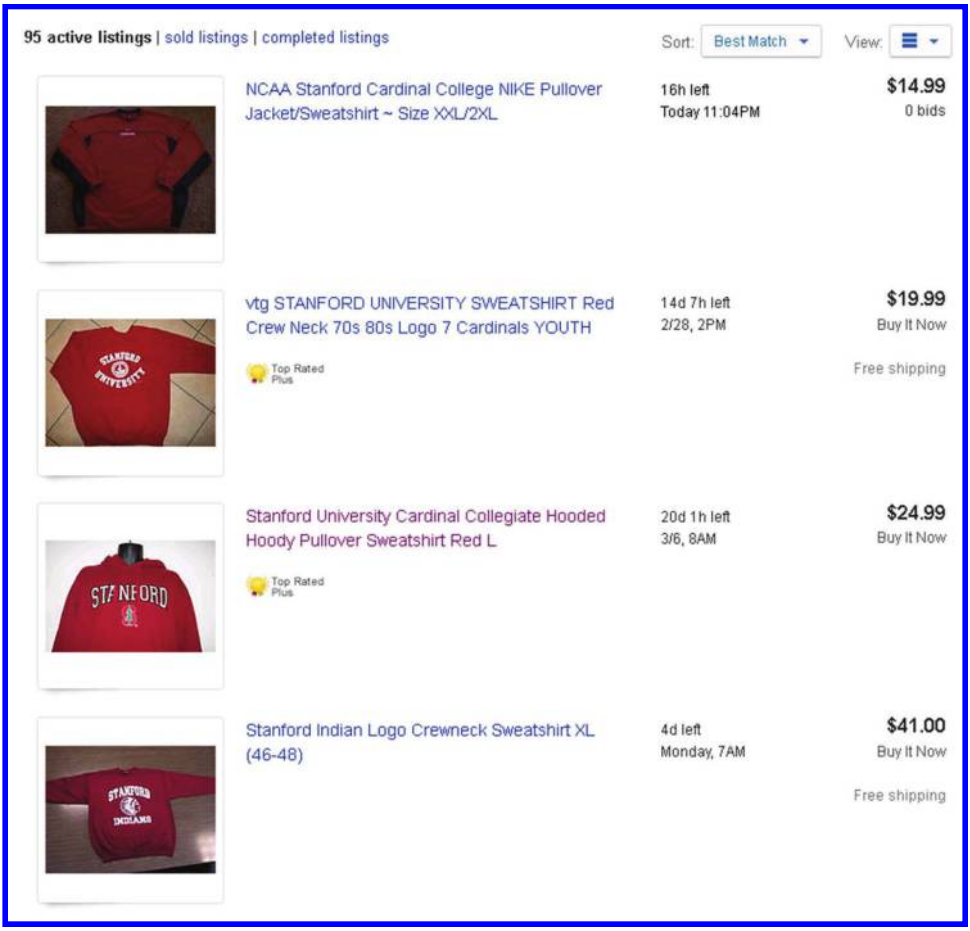

FIGURE 3A. SCREENSHOT OF A TYPICAL EBAY SEARCH RESULTS PAGE

Notes: This is a screenshot of an eBay search results page (for a query that searched for "stanford sweatshirt"). The key thing to notice is that the details about the seller, and especially about the seller's location, are not provided on this page.

location (shown beneath the price and shipping in the middle of Figure 3B). In principle, this is often enough to determine if tax will be charged, but buyers can also click on the "Shipping and payment" tab, shown in Figure 3C, where many sellers list more detailed tax information. Tax information is also displayed after the buyer initiates a purchase and before it is confirmed.

The idea of our research design is to compare buyers who arrive at the same item page, some of whom are located in the same state as the seller (and would be charged tax) and some of whom are not. In this way we compare like-minded buyers considering the same item, only with different tax-inclusive prices. Of course, for sales tax to matter, at least some consumers must take note of it. Chetty, Looney, and Kroft (2009) and Ellison and Ellison (2009) have made the point that sales taxes often may not be as salient as retail prices, and we have seen in our own research (Einav et al. 2011) that eBay consumers appear not to fully internalize shipping fees, which if anything are displayed more prominently than taxes. For this reason, tax price sensitivity may understate retail price sensitivity.

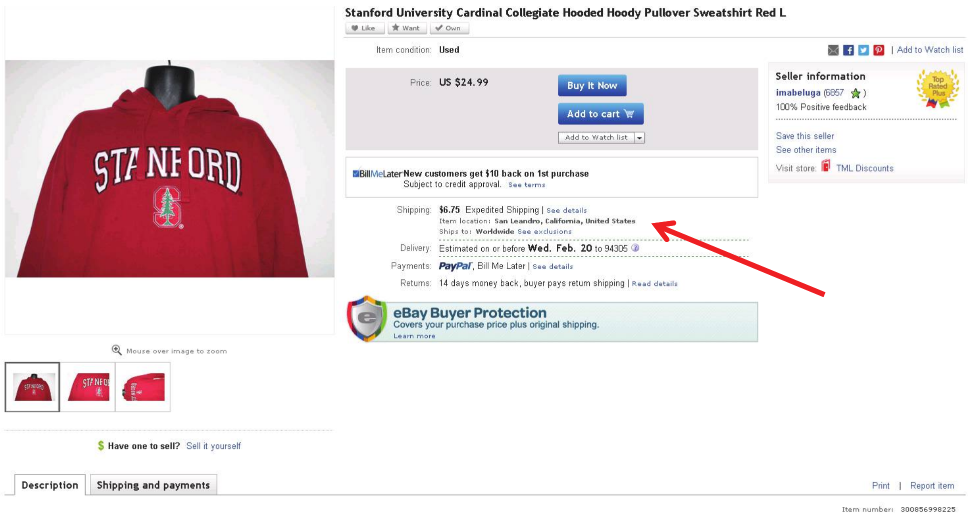

FIGURE 3B. SCREENSHOT OF AN ITEM PAGE

Notes: This is a screenshot of an eBay search results item page. The seller location is now presented below the price in the middle of the page.

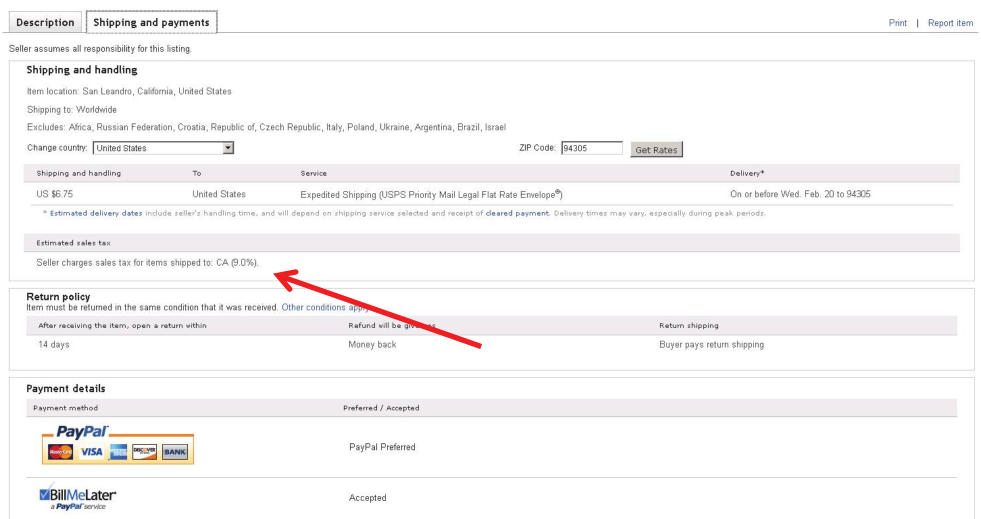

FIGURE 3C. SCREENSHOT OF THE "SHIPPING AND PAYMENTS" TAB

Notes: This is a screenshot of the "Shipping and payments" tab (clickable from the item page). The applicable tax rate is now presented explicitly.

B. *Data*

We assemble detailed browsing and purchasing data for several hundred thousand items available on eBay. We start with the set of all items listed between January 1 and December 31, 2010. From this universe, we select all items that were offered at a posted price, with at least ten available units, by sellers who use eBay's "tax table" application. We focus on items with a relatively large available quantity so that we can observe multiple purchases for each item (almost all

Table 1—Item-Level Data: Summary Statistics

	Observations	Mean	SD	p25	p50	p75
Item list price ($)	275,020	36.95	164.98	6.22	12.99	29.99
Item sales tax[a]	275,020	7.83%	1.71%	7.00%	8.20%	8.88%
Logged-in users viewing item	275,020	24.7	55.6	4	9	23
In-state users viewing item	275,020	1.8	5	0	0	2
Purchase rate (purchases/views)	275,020	0.21	0.17	0.08	0.17	0.33
Average viewer distance[b]	275,020	1,939	732	1,409	1,842	2,469

Notes: Table shows summary statistics for 275,020 items listed on eBay by 10,347 distinct sellers between January 1, 2010 and December 31, 2010. The data cover 6,796,691 page views, each by a different user.

[a] The item sales tax is the combined state and local tax in the seller's zip code as of April 1, 2010. Tax-exempt clothing items are given a rate of zero.

[b] Throughout, distances are measured as the great-circle distance between the centroids of the user zip code and the seller zip code, in kilometers.

transactions on eBay are single unit purchases) and avoid any potential issues that might arise from listings selling out after one or a few purchases. We focus on sellers who use the tax table so we can be confident of their tax collection practices. The tax table is used by retailers who list significant numbers of items: it allows a seller to enter the tax rate it wishes to charge buyers in states where it has nexus, and the seller can apply this rate easily to all its items.

We sort through trillions of user interaction events to identify, for each item, all item page views by logged-in eBay users during the observation period.[5] We restrict attention to users located in the United States. We use the respective zip codes of the buyer and seller to determine the applicable sales tax. We assume that for an in-state sale, the seller charges the combined state and local sales tax in its own zip code, while for out-of-state sales no tax is charged.[6] We also calculate the great-circle distance between the centroids of the buyer and seller zip codes. Finally, for each item page view we determine if the user subsequently purchased the item during the browsing session.

Table 1 presents summary statistics for the items in the data. We report statistics only for those items that had at least one qualifying purchase because in the fixed-effects specifications we use below, items with zero purchases provide no identifying information and hence are dropped from the analysis. The resulting data consist of 275,020 listed items posted by 10,347 different sellers. The average pre-tax price of these items is $37. The average combined state and local sales tax in the seller's zip code is just under 8 percent. We observe an average of around 25 user page views for each item. This gives us a total of 6,796,691 page views. Overall, for the average item in our sample, about one in five of these page views results in a purchase.

[5] We focus on logged-in users so we can reliably identify each consumer's location, and discard observations with incomplete or ambiguous location information. Note that we require that the user logged in prior to having viewed the item in that browsing session, to eliminate the concern that users might log in specifically to complete a particular purchase. Also, if a user viewed a sample item page in a given browsing session and then viewed that or some other sample item page in a subsequent session, we only use the data from the first session. We do this to simplify the analysis, as it allows us to consider a cross section of encounters for each item.

[6] In the event that some sellers charge tax in more than one state, or adjust the local tax depending on the location of in-state buyers, this will introduce some measurement error in the "effective tax rate" we use below in our regressions.

C. *Consumer Tax Sensitivity*

We estimate consumer sensitivity to sales tax using a fixed-effects logit model of the purchase decision. Let k index the items, and i index the viewers of each item. The after-tax price for item k faced by consumer i is $(1 + \tau_{ik})p_k$, the tax multiple times the pre-tax price. We assume that

$$\Pr(i \text{ buys } k \,|\, i \text{ views } k) = \frac{\exp(u_{ik})}{1 + \exp(u_{ik})}, \tag{1}$$

where

$$u_{ik} = \alpha_k + \beta \log(1 + \tau_{ik}) + g(d_{ik}) + \gamma \mathbf{1}\{\text{state } i = \text{ state } k\}. \tag{2}$$

Here the first term α_k is a fixed effect that captures each item's general desirability including its pre-tax price (which is constant across viewers).[7] The second term is the effect of the relevant effective tax rate τ_{ik}, which is equal to the combined sales tax in the item's zip code if i is a same-state buyer, and zero otherwise. We include the distance between the buyer and seller, denoted d_{ik}, as a control to account for the possibility that buyers may prefer nearby items, for instance because they expect faster shipping or have more trust in local sellers.

Our first specification includes only these first three terms. In this specification, the primary source of variation in τ_{ik} is between in-state buyers who are taxed and out-of-state buyers who are not, holding fixed their physical distance to the item. One concern, however, is that buyers may prefer in-state items even controlling for distance. Such "border effects" are common in the international trade literature (Anderson 2011), and appear in Hortaçsu, Martínez-Jerez, and Douglas' (2009) eBay study. Focusing on consumers who already have clicked on an item should rule out many obvious examples of home-state preference (e.g., Nebraska residents preferring Cornhuskers T-shirts), but any residual preference might bias an estimate of β toward zero. In our preferred specification, we include a dummy variable indicating whether the buyer is located in the same state as the item. With this control, the tax parameter β is identified from differences in the same-state "avoidance" of buyers in high and low tax states.

The estimates are reported in Table 2. The first column reports an initial specification with no home-state preference dummy. The second column is our preferred specification. To translate the reported estimate of the tax coefficient β into a price elasticity, one needs to multiply it by one minus the purchase rate, or by approximately 0.79.[8] With that in mind, our preferred specification yields a tax-price elasticity of -1.7. That is, for every one percentage point increase in the sales tax (1 percent increase in the post-tax price), purchasing decreases by about 1.7 percent. A viewer charged a 5 percent sales tax is about 5 percent more likely to purchase

[7] We have the after-tax price $p^*_{ik} = (1+\tau_{ik})p_k$ enter logarithmically as $\beta \log p^*_{ik}$, which yields $\beta \log p_k + \beta \log(1+\tau_{ik})$. As the pre-tax price for item k does not vary across viewers, its effect $\beta \log p_k$ is subsumed by the item fixed-effect α in equation (2).

[8] The elasticities reported are computed at the margin corresponding to the average purchase probability for items in our sample.

Table 2—Item-Level Estimates of Tax Sensitivity

	Dependent variable: Indicator that is equal to 1 if item purchased		
	All items (1)	All items (2)	By rate type[a] (3)
log(1 + effective tax)	−1.182	−2.131	−1.897
	(0.104)	(0.406)	(0.408)
log(distance)	−0.029	−0.028	−0.025
	(0.002)	(0.002)	(0.003)
Same state dummy		0.081	0.063
		(0.033)	(0.033)
log(distance) × calc. rate dummy			−0.026
			(0.005)
Fixed effects	Item	Item	Item
Number of distinct items	275,020	275,020	275,020
Number of page views	6,796,691	6,796,691	6,796,691
Mean of dep. variable	0.215	0.215	0.215
Implied tax-price elasticity	−0.928	−1.673	−1.489
	(0.082)	(0.319)	(0.320)

Notes: Table shows coefficient estimates from a conditional logit regression where the dependent variable is equal to 1 if the viewing user purchased the item during the browsing session and zero otherwise. Each observation reflects a distinct page view by a distinct user. As indicated in the table, the mean purchase probability is 0.215, and the tax-price elasticity is the estimated coefficient (at the first row) multiplied by (one minus the purchase rate). Standard errors are given in parentheses below estimates.

[a]Items can be listed as "flat shipping rate" or as "calculated shipping rate." In the latter case, the shipping cost of the item (paid by the buyer) is increasing in the shipping distance.

than an equivalent viewer facing an 8 percent sales tax, and 8 percent less likely to purchase than one who is charged no sales tax.[9]

One reaction to this estimate is that, for retail items in a highly competitive marketplace, demand appears to be surprisingly inelastic. We certainly would not be surprised if the items in our sample had retail price elasticities several times larger (i.e., more negative) than our estimates.[10] There are at least two plausible explanations, however. First, buyers may pay less attention to the sales tax than to the retail price. Second, we are focusing on the response of shoppers who already have identified and expressed interest in an item. If the primary effect of a retail price increase is to cause buyers not to click on the item in the first place, the relevant price elasticity for the sellers of these items could be considerably larger.

The results in Table 2 on the effects of distance and home-state preference are also interesting. There is a clear and consistent relationship between distance and the probability of purchase. All else equal, a consumer who is 250 kilometers from an item is about 3 percent more likely to purchase than one who is 1,000 kilometers

[9]A potential concern with our preferred specification is that it assumes that the "same state" effect is the same for all states. To assess how sensitive the results are to this assumption, online Appendix Table A1 presents results that allow the same-state effect to vary across regions, yielding similar tax elasticity estimates (although the point estimates are about 25 percent smaller).

[10]Unfortunately we do not have an obvious way to obtain such estimates with the current data. In other work, however, we have used cross-sectional variation in the prices of video games offered by multiple sellers to estimate retail price elasticities on the order of −5 to −15 (Dinerstein et al. 2013).

from the item. One possible explanation is shipping time: the closer the item, the less delay a buyer may expect. For a small fraction of the items (just under 15 percent), shipping cost may also be a factor because rather than charging a flat shipping fee (typical on eBay), the seller charges a calculated rate based on the distance.

The presence of these variable shipping rate items provides a useful opportunity to look at the salience of "add-on" prices. In column 3, we allow the effect of distance to vary depending on the type of shipping fee. Consumers are twice as sensitive to distance when it affects the shipping fee. To interpret the magnitude of the coefficient, we observe that the average variable rate shipping fee increases by around \$0.56 for every doubling in distance. If we take the distance coefficient for flat shipping rate items to be a base preference for distance and interpret the additional sensitivity for calculated shipping rate items as a price response, this suggests that for a typical item in the sample, the \$0.56 increase in the shipping fee from a doubling of distance reduces the probability of purchase by around 1.4 percent. For a good priced at \$43 (which is our sample average, calculated shipping included), a \$0.56 increase in the shipping fee corresponds to a price elasticity of about -1.1.[11]

Finally, the estimates in Table 2 suggest a substantial home-state preference. Controlling for distance to an item, consumers are about 7 percent more likely to buy if the seller is located in the same state. This is consistent with the results reported in Section II and in Hortaçsu, Martínez-Jerez, and Douglas (2009), but perhaps more surprising given that we are focusing only on interested buyers. In thinking about this effect, it may be useful to consider how it is identified. Note that we can group the second and fourth terms in our logit specification, equation (2), as $\mathbf{1}\{\text{state } i = \text{state } k\}[\gamma + \beta \log(1 + \tau_k)]$, where τ_k is item k's combined state and local sales tax. By comparing buyers located similar distances from an item but on either side of the state border we can identify the combined effect $\gamma + \beta \log(1 + \tau_k)$. We estimate this to be close to zero on average (i.e., $\gamma + \beta \overline{\log(1 + \tau)} \approx 0$). The variation in individual item tax rates then allows us to identify β, so that γ falls out as an intercept.[12] Identifying γ, however, requires some extrapolation because nearly all the items have a combined tax rate between 5 and 10 percent. As a result, our home-state preference estimate varies a bit across specifications, although it appears in all cases to be substantial.

D. *Heterogeneity in Tax Sensitivity*

Our baseline results yield an average tax-price elasticity across a wide range of retail items. We can take advantage of the rich data by splitting the sample and comparing purchasing behavior for goods in different retail categories or at different price points. Such an exercise is interesting in part because not many studies have been able to provide reliable and comparable price or tax-price elasticity estimates

[11] This elasticity estimate is about half of our estimated tax elasticity. This appears reasonable given that calculated shipping rates, unlike flat shipping rates, depend on distance in a non-transparent way.

[12] Online Appendix Figure A1 may provide intuition. To construct it, we estimate a separate same-state effect for each state, and then plot these estimates against the sales tax rate for state residents, allowing for a visual inspection of how the data determine our β and γ estimates.

for large numbers of retail goods.[13] These tax-price elasticities are a function of both the underlying price elasticities of demand and the salience of sales taxes, and differences may be driven by either or both of these factors.

Table 3A reports separate estimates for the six largest product categories in our sample. We estimate the largest elasticity for electronics (-4.3), followed by sporting goods (-3.3). Three other categories (cell phones, computers, and clothing) are estimated to have a tax-price elasticity of -1.5 to -2. The home and garden category is an exception, as we estimate essentially no tax sensitivity. Although the estimates are not sufficiently precise to be definitive, the results generally conform to the intuitive idea that price sensitivity might be greater in more commodity type product categories than in categories with greater product differentiation.[14]

The clothing category is of particular interest, as clothing items are exempt from sales tax in nine states.[15] We can thus use it as a "placebo" test, as a way to verify that our estimates likely capture the effect of taxes rather than some other unobservables features that vary across states. To do this, we run the same regression (for clothing items) reported in column 4 of Table 3A, but counterfactually assume that tax rates are positive even in the tax exempt states. We then interact the tax rate with a dummy variable that is equal to one for tax exempt states. Because the effective tax rate (for clothing items) in these states is zero, the effect of the (counterfactually positive) tax rate should be zero as well, but if the estimates are driven by some other factors, then the tax rate effect should be negative and similar to the one reported in Table 3A. The results, shown in online Appendix Table A2, are consistent with a tax effect; we estimate a *positive* and insignificant coefficient—a coefficient of 0.53 with a standard error of 2.19—on the "incorrect" sales tax rate in the nine states with tax exemptions.

Table 3B splits the sample based on the retail prices of the sample items. The estimated tax coefficient is larger in magnitude for more expensive items, which also have a lower purchase rate. Translated into tax-price elasticities, we find the elasticity of the cheaper items (selling for less than \$6, or for \$6–\$12) to be between -0.6 and -1.1, compared to an elasticity of -2.1 to -2.5 for more expensive items. One hypothesis is that taxes are more salient for the expensive goods because their dollar effect is larger and perhaps noticed by more consumers. The differing estimates could also reflect differences in the retail price elasticities, which also would be interesting because it is not clear a priori that demand should be more elastic for more expensive items.

In addition to exploring differences across items, we also considered the possibility that different buyers would be systematically more or less sensitive to taxes. In particular, we looked separately at experienced and inexperienced buyers, using a segmentation developed by eBay that correlates roughly with the number of past purchases a buyer has made. The results are shown in online Appendix Table A3. To our

[13] One exception is the marketing literature that uses grocery-store scanner data to estimate price elasticities for a variety of goods. For instance, a well cited paper by Hoch et al. (1995) reports average own-price price elasticities for eighteen categories of goods sold at Dominick's grocery stores. They lie in a remarkably narrow range, from -0.79 to -2.59.

[14] One may notice that the estimated home-state effect varies substantially across product categories. Recall, however, our discussion from the previous section, which explains the issues involved in estimating this effect.

[15] As of January 2010, the beginning of our sample period. The nine states are Connecticut, Massachusetts, Minnesota, New Jersey, New York, Pennsylvania, Rhode Island, Texas, and Vermont. In Massachusetts and New York the exemption is only applicable for clothing items that sell below a certain price (\$175 and \$110, respectively, as of January 2010).

Table 3A—Item-Level Estimates of Tax Sensitivity by Category

	Dependent variable: Indicator that is equal to 1 if item purchased					
	Electronics (1)	Cell phones (2)	Computers (3)	Clothing (4)	Home & garden (5)	Sporting goods (6)
log(1 + effective tax)	−5.325	−2.792	−2.733	−1.648	0.273	−3.864
	(2.105)	(1.436)	(1.468)	(0.653)	(1.707)	(2.273)
log(distance)	−0.029	−0.031	−0.042	−0.016	−0.025	−0.032
	(0.006)	(0.006)	(0.006)	(0.008)	(0.007)	(0.009)
Same state dummy	0.306	0.132	0.111	0.058	−0.074	0.215
	(0.183)	(0.119)	(0.126)	(0.047)	(0.129)	(0.176)
Fixed effects	Item	Item	Item	Item	Item	Item
Number of distinct items	24,013	42,188	45,640	16,489	28,034	12,263
Number of page views	733,753	701,155	707,973	677,031	929,767	468,955
Mean of dep. variable	0.200	0.274	0.292	0.132	0.166	0.144
Implied tax-price elasticity	−4.260	−2.027	−1.935	−1.430	0.228	−3.308
	(1.684)	(1.042)	(1.039)	(0.567)	(1.318)	(1.946)

Notes: As in Table 2, the table shows coefficient estimates from a conditional logit regression where the dependent variable is equal to 1 if the viewing user purchased the item during the browsing session and zero otherwise. Each observation reflects a distinct page view by a distinct user. The mean purchase probability is shown at the second-to-last row of the table, and the tax-price elasticity is the estimated coefficient (at the first row) multiplied by (one minus the purchase rate). Standard errors are given in parentheses below estimates. A Wald test fails to reject the equality of tax coefficients across the six categories (p-value $= 0.34$).

Table 3B—Item-Level Estimates of Tax Sensitivity by Price Level

	Dependent variable: Indicator that is equal to 1 if item purchased			
	< \$6 (1)	\$6–12 (2)	\$12–24 (3)	> \$24 (4)
log(1 + effective tax)	−1.502	−0.809	−2.740	−2.979
	(0.962)	(0.928)	(0.698)	(0.687)
log(distance)	−0.034	−0.028	−0.022	−0.029
	(0.005)	(0.004)	(0.004)	(0.004)
Same state dummy	0.028	−0.020	0.149	0.132
	(0.079)	(0.076)	(0.056)	(0.057)
Fixed effects	Item	Item	Item	Item
Number of distinct items	68,339	62,830	58,997	84,854
Number of page views	1,030,448	1,109,980	1,414,700	3,241,563
Mean of dep. variable	0.265	0.243	0.204	0.160
Implied tax-price elasticity	−1.104	−0.612	−2.181	−2.502
	(0.707)	(0.702)	(0.556)	(0.577)

Notes: As in Table 2, the table shows coefficient estimates from a conditional logit regression where the dependent variable is equal to 1 if the viewing user purchased the item during the browsing session and zero otherwise. Each observation reflects a distinct page view by a distinct user. The mean purchase probability is shown at the second-to-last row of the table, and the tax-price elasticity is the estimated coefficient (at the first row) multiplied by (one minus the purchase rate). Standard errors are given in parentheses below estimates. A Wald test fails to reject the equality of tax coefficients across the four price bins (p-value $= 0.21$).

surprise, we found only very minor differences in tax sensitivity across buyers with different amounts of experience, and no differences that were statistically significant.

E. *Substitution Effects*

Our final exercise in this section uses the tracking allowed by the clickstream data to investigate whether buyers who receive a "tax surprise" are likely to substitute to an alternative item. To do this, we rely on the same set of user-item observations as in the analysis above, but for each user, track whether he or she subsequently purchased a different item, and if so, the characteristics of this purchase.[16] We use this expanded data to investigate the generalized response of browsing consumers who receive an adverse price shock. We view this exercise as interesting in its own right, but mostly as a way to validate that the tax effect documented above is capturing a behavioral response and is not merely a statistical anomaly.

The top panel of Table 4 reports the results of a series of logit regressions. Each column corresponds to a different outcome variable, but the regressors are identical and associated with the original page view in Table 1 and Table 2. The positive tax coefficients in columns 1, 2, 4, and 5 suggest significant substitution: individuals who receive a negative tax surprise are noticeably more likely to purchase a different item subsequently in the session, more likely to purchase an item from a different seller, and more likely to purchase an item from a different seller that is in the same product category as the original item. The negative tax coefficient in column 3 is also interesting. The estimate indicates that consumers who receive a large negative tax surprise are less likely to purchase some other home-state item (which should also have a high tax) than consumers who receive smaller tax surprises.

These results are consistent with the idea that when users experience a negative tax surprise, they keep searching and perhaps buy a similar item from a different seller. In the bottom panel of Table 4, we attempt to hone in on this substitution by relating subsequent purchasing to whether a consumer purchased the original item. We report two types of results based on linear probability models. In the top row we regress an indicator for a subsequent purchase on an indicator for whether or not the consumer purchased the original item. There is little raw correlation between the two purchase decisions. In the bottom row, we report the results from an instrumental variables specification in which we use the location regressors from Panel A as instrumental variables to identify a causal effect of purchasing the first item on the decision of whether or not to make a different purchase. Here we find relatively clear and strong substitution effects. Indeed, the estimated effect is surprisingly large (probably too large): purchasing the first good essentially eliminates the chance some later good is purchased.

II. Aggregate Responses to Sales Taxes

Our item-level analysis shows that consumers shopping online seem to react to sales tax in predictable and significant ways. However, although the estimates have

[16] Here, "subsequently" means subsequently in the same browsing session, where a session is defined (by eBay) as a string of events from the same user in the same browser. A session ends if 30 minutes pass without an event.

Table 4—Substitution Patterns

	Dependent variable: Indicator that is equal to 1 if… during subsequent session				
	Bought any other item (1)	Bought from a different seller (2)	Bought from a diff. seller but in the same state (3)	Bought from a diff. seller, in the same broad category (4)	Bought from a diff. seller, in the same narrow category (5)
Panel A. Reduced form effects					
log(1+effective tax)	1.071	1.234	−1.183	1.163	1.153
	(0.314)	(0.338)	(0.881)	(0.374)	(0.494)
log(distance)	0.016	0.021	−0.017	0.024	0.024
	(0.002)	(0.002)	(0.004)	(0.002)	(0.003)
Same state dummy	−0.068	−0.081	0.169	−0.065	−0.067
	(0.026)	(0.028)	(0.076)	(0.031)	(0.041)
Fixed effects	Item	Item	Item	Item	Item
Number of distinct items (estimation sample)	205,314	192,435	54,821	168,638	121,789
Number of page views	6,348,623	6,217,586	2,831,003	5,834,658	4,830,520
Mean of dep. variable (in estimation sample)	0.228	0.211	0.114	0.190	0.165
Mean of dep. variable (in original sample)	0.178	0.153	0.023	0.120	0.075
Fraction bought the original item (in est. sample)	0.181	0.176	0.162	0.173	0.170
Panel B. Substitution estimates (linear probability models)					
Original item was bought (OLS)	0.019	−0.005	−0.005	−0.025	−0.036
	(0.005)	(0.004)	(0.001)	(0.004)	(0.004)
Original item was bought (IV)	−0.662	−0.779	0.014	−0.698	−0.435
	(0.091)	(0.102)	(0.039)	(0.084)	(0.061)

Notes: In panel A we report conditional logit regressions similar to those in Table 2, except that the dependent variable reflects outcomes from the user's browsing session following the original page view that got him into the sample. All the right-hand-side variables apply to the original page view, as in Table 2. Note also that the estimation sample shrinks for some of the narrower outcomes that lead us to drop items for which subsequent outcomes do not vary (they are all zero). In panel B we use linear probability models to estimate the direct effect of whether the original item was bought or not on the same subsequent outcomes used in panel A. We first report an OLS estimate (with item fixed effects), and then report an IV estimate, in which the regressors from panel A are used as instruments (so that the results reported in Table 2 can be thought of as similar to the first stage).

the advantage of coming from a clean, well-controlled research design, they are a step removed from the relevant policy questions. These questions concern changes in tax rates or tax treatment at the state or national level. In this section, we pursue a second approach that brings us closer to a direct estimate of the policy-relevant parameters. We use aggregated data on trade flows to estimate the effect of sales taxes on online purchasing shares, and on the overall volume of online purchases. As we explain below, we rely on a difference-in-differences strategy which exploits the variation in sales tax rates to identify cross-state substitution in online purchases, while making use of tax changes over time to identify the overall effect on online purchases.

A. *Data and Preliminary Evidence*

We construct measures of online trade flows using all eBay.com transactions, including both posted-price and auction listings, during the years 2008–2010, excluding Autos and Real Estate. We aggregate these data in two ways.

Our first dataset consists of annual state-to-state trade flows. Observations in this dataset are at the *ijt* level, where *i* represents the buying state, *j* the selling state, and *t* the year. We define the applicable tax rate for state *i* in year *t* to be the (population weighted) average combined state and local tax rates for state residents, with the average taken across state resident-months.

Our second dataset, which we use to look at overall online purchasing, groups eBay transactions into total monthly purchase counts by county. Observations in this dataset are at the *it* level, where *i* indexes the buying county and *t* the month. In this case, the applicable tax rate for county *i* in month *t* (population weighted) average combined state and local tax rates for county residents.

We use the data on state-to-state trade flows to look at the propensity of state residents to make online purchases out of state, relative to their overall online purchasing and the quantity and general attractiveness of goods available from different locations. To see roughly how this approach works, let s_{ij} denote the share of state *i*'s online purchases that are from sellers located in state *j*. Let s_j denote the overall share of eBay purchases that are from sellers in state *j*. With this notation, the ratio s_{ij}/s_j captures state *i*'s relative preference for state *j* goods, and a natural way to look for tax sensitivity is to relate the relative preference of state *i* buyers for home-state sellers, that is, s_{ii}/s_i, to the state's applicable tax rate.

Figure 4 presents a first-pass analysis. For each state, we calculate the share of state purchases that were home-state purchases and divide this by the state's share of overall eBay sales. We then plot this measure against the state's average sales tax. We construct purchasing and sales shares using sales counts rather than transaction value; the plot looks very similar using value shares. Two points are immediately apparent. First, all fifty states exhibit a home bias in purchasing, i.e., $s_{ii}/s_i > 1$. Second, consumers in high tax states do notably less home-state purchasing, consistent with tax shifting purchases out of state. Of course, this analysis doesn't account for potentially confounding factors such as state size (intrastate distance) or the distance to states with attractive goods, but we will see below that adding more detailed controls leaves the basic relationship intact.

The second question of interest is whether sales taxes increase overall online purchases, presumably due to substitution away from taxed offline (local) purchases. This question is more challenging with a purely cross-sectional approach. Intuitively, while the overall share of eBay purchases made from Iowan sellers might be a reasonable proxy for the share of purchases that Iowans should make from these sellers, absent any home-state preference or sales taxes, it is less obvious that the overall online (or eBay) purchasing by residents of other states should be a good proxy for that of Iowans, absent any incentive from sales tax differences. Indeed, Figure 5 provides a simple plot of each state's per capita eBay purchases against the state's average sales tax. The raw correlation is negative, indicating that high tax states generally do less eBay purchasing, a surprising correlation unless other factors apart from taxes are at work.

One way to address this is to control better for cross-state differences. Roughly speaking, this is the approach taken by Goolsbee (2000 a, b), Alm and Melnik (2005), Scanlan (2007), and in the first half of Ellison and Ellison (2009), all of whom regress some statistic of online purchasing on home sales tax and a set of controls. Nevertheless, one may worry that even relatively rich covariates will not suffice

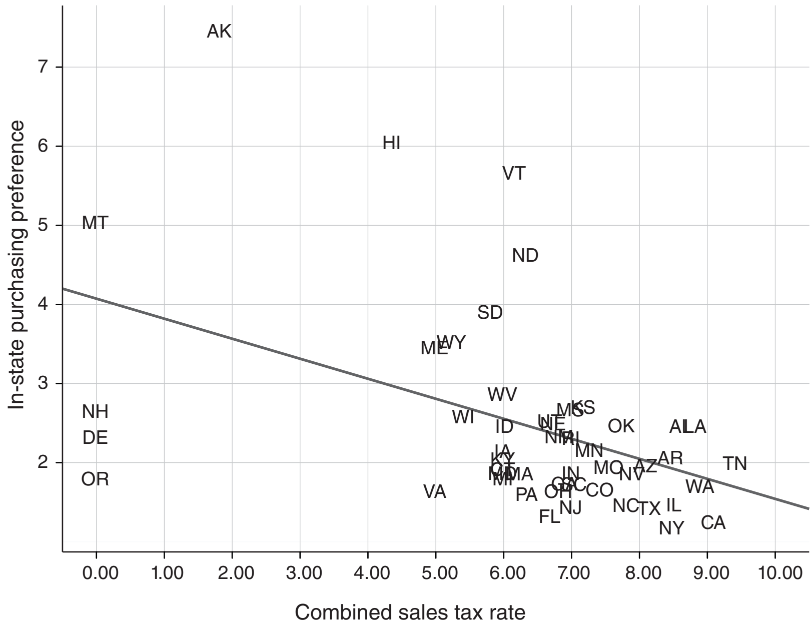

FIGURE 4

Notes: Figure presents the relationship between in-state purchasing rate and the state's (population weighted) sales tax rate. The in-state purchasing rate is the ratio between the state's purchasing share of the state's sales to the state's overall purchasing share. Purchasing and sales are computed as the number of transactions (not their value) on eBay during our observation period (2008–2010).

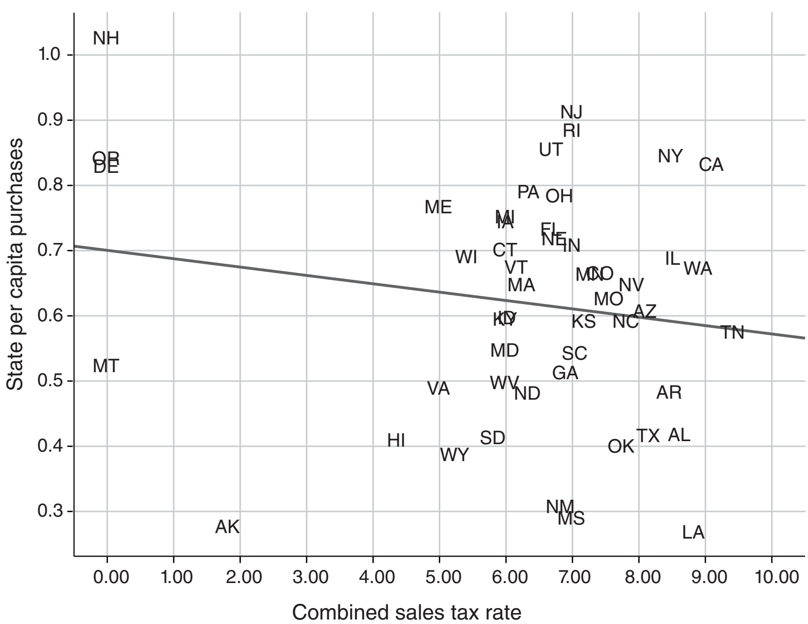

FIGURE 5

Notes: Figure presents the relationship between the state's per capita number of purchases on eBay during our observation period (2008–2010) and the state's (population weighted) average sales tax rate.

to control for underlying heterogeneity in preferences, prices, or patterns of retail behavior or Internet use across states. With this in mind, we also report results that rely on the variation in tax rates caused by changes at the state and local level (shown in Figure 2).

B. *Sales Taxes and Cross-State Substitution*

We start by considering the relationship between taxes and cross-state purchasing patterns. As is common in empirical studies of trade flows, we work with a CES representation of consumer demand (Anderson 2011). We think of each state as having a representative buyer and selling a single composite good. Let i index buyer locations and j index "goods," or equivalently seller locations. Let q_{ij} denote the quantity purchased by state i from state j, and let p_{ij} denote the unit price including any sales tax.

With the CES representation, the quantities q_{ij} solve, for each i,

$$(3) \qquad \max_{q_{i1},\ldots,q_{iJ}} \left(\sum_j (q_{ij}/\zeta_{ij})^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \quad \text{s.t.} \quad \sum_j p_{ij} q_{ij} \leq w_i.$$

Here w_i is i's expenditure on online retail goods, the ζ_{ij} are preference parameters, and σ is the elasticity of substitution. The CES demands are

$$(4) \qquad q_{ij} = \frac{p_{ij}^{-\sigma} \zeta_{ij}^{1-\sigma}}{P_i^{1-\sigma}} w_i,$$

where P_i is the CES price index for online goods.[17] Assuming that this general demand structure applies in each period t, and taking logs, we have:

$$(5) \quad \log q_{ijt} = a_{it} - \sigma \log p_{ijt} + (1 - \sigma) \log \zeta_{ijt} - (1 - \sigma) \log P_{it} + \log w_{it}.$$

This expression will be the basis for our estimates of cross-state substitution in response to the sales tax on in-state purchases. To this end, we express prices as $p_{ijt} = (1 + \tau_{ijt}) p_{jt}$, where p_{jt} is the base price on goods sold from location j, and τ_{ijt} is the applicable sales tax. Suppose that in addition we can write the preference parameter ζ_{ijt} as $\zeta_{ijt} = (h^{\mathbf{1}\{i=j\}} d_{ij}^{\gamma})^{1/(1-\sigma)} \zeta_{jt}$, where h captures same-state purchasing preference, d_{ij} is the distance between location i and j, and ζ_{jt} is the general attractiveness of location j goods. With these assumptions, purchases by state i from state j at time t can be expressed as

$$(6) \qquad \log q_{ijt} = a_{it} + b_{jt} - \sigma \log(1 + \tau_{ijt}) + \gamma \log(d_{ij}) + h\mathbf{1}\{i = j\}.$$

[17] The CES price index is $P_i = \left(\sum_j (\zeta_{ij} p_{ij})^{1-\sigma}\right)^{1/(1-\sigma)}$. The one property of this price index we will use is that $\partial \log P_i / \partial \log p_{ij} = x_{ij}$, where $x_{ij} = p_{ij} q_{ij} / w_i$ is the expenditure share of location i consumers devoted to location j goods. Online Appendix C provides more details on derivations presented in this section.

Table 5—Estimates of Online State-to-State Flows

	Number of state-to-state purchases				
	(1)	(2)	(3)	(4)	(5)
log(1 + effective tax)	−5.556	−5.878	−4.234	−4.743	−3.642
	(1.932)	(2.327)	(2.237)	(3.377)	(1.795)
log(distance)	−0.104	−0.104	−0.105		
	(0.008)	(0.007)	(0.006)		
Same state dummy	0.537	0.560	0.988		
	(0.146)	(0.149)	(0.367)		
log(distance) × same state			−0.105		
			(0.085)		
Fixed effects	Buyer state × year, seller state	Buyer state × year, seller state × year	Buyer state × year, seller state × year	Buyer state × year, seller state × year, buyer-seller state pair	Buyer state × month, seller state × month, buyer-seller state pair
Observations	7,500	7,500	7,500	7,500	90,000

Notes: Table shows results from a Poisson regression where the dependent variable is the number of sales from state i to state j, using a panel of three years (2008–2010); data is aggregated to the yearly level for columns 1–4 and the monthly level for column 5. Standard errors are computed using a state-level block bootstrap with 50 replications. The distance variable is measured at the (i, j) state pair level by computing the average distance over all transactions between a seller zip from state i and a buyer zip from state j.

We estimate the model as a Poisson quasi-maximum likelihood regression using our data on annual state-to-state eBay trade flows.[18] In this specification, the combined term $\sigma \log(1 + \tau_{ijt}) + h\mathbf{1}\{i = j\}$ is identified by the relative propensity of buyers to purchase in-state, after controlling for distance and the attractiveness of each state's products. More narrowly, the tax effect σ is identified by differences in the home bias of states with low and high sales tax rates. One difference with the earlier individual-level approach, however, is that without item-level fixed effects, we control less well for particular idiosyncrasies in the types of goods that buyers in certain states might favor.

Table 5 reports the results from four specifications with progressively tighter controls. In column 1, we allow for buyer state by year fixed effects (a_{it}'s in the above equation) and seller state fixed effects (assuming $b_{jt} = b_j$). In columns 2 and 3, we relax the latter assumption and allow for seller state by year fixed effects. In each of the first two specifications, we use both cross-sectional and time series variation in tax rates to identify the effect of tax rates. In the remaining specifications reported in columns 4 and 5, we replace our distance and same-state controls with fixed effects for each state pair (c_{ij} dummies), and rely solely on the time series variation in tax rates. In column 4 observations remain aggregated to the year level, while in column 5 we disaggregate to monthly purchase counts and tax rates.

[18] Here we follow common practice in the empirical trade literature (Anderson 2011), which is to use a count specification rather than a log-linear regression model. In online Appendix Table A4, we report results from an alternative specification using dollar volumes of trade, rather than counts.

Our main interest is in the parameter $-\sigma$ given by the estimated tax coefficient, which is similar across specifications, ranging from -3.6 to -5.9.[19] The interpretation is that a one percentage point increase in a state's sales tax rate will be associated with a roughly 5 percent decrease in online home-state purchases. This calculation holds fixed the total online expenditure; as we discuss below, the reduction in online same-state purchases will be offset if a sales tax increase shifts purchasing from offline to online. Note that although the point estimates are fairly stable across specifications, the estimates are not terribly precise: taking column 2 as our benchmark specification, the standard error is 2.3, and the 95 percent confidence interval is -1.3 to -10.4.

The other coefficient estimates in Table 5 are also of interest, in part because they are quite similar to those reported in Hortaçsu, Martínez-Jerez, and Douglas (2009). As in their paper, we find that trade drops off with distance: state i's purchases fall by roughly 7 percent as the distance to the selling state doubles. There is also a substantial home-state effect: after controlling for the adverse tax consequences of home-state purchases, intrastate trade is about 75 percent higher than would be expected based on distance alone. As a comparison, Hortaçsu, Martínez-Jerez, and Douglas (Table 3, Model III) reported estimates that imply a doubling of distance reduces trade by about 5 percent and find an almost identical same-state excess trade of 75 percent.[20] Interestingly, the estimated distance effects for eBay purchasing are substantially smaller than what is estimated in many similar gravity-type regressions (including estimates for purchasing on MercadoLibre, a South American platform, also reported in Hortaçsu, Martínez-Jerez, and Douglas 2009).

C. *Sales Taxes and Online-Offline Substitution*

The results reported in the previous section speak to the effect of sales tax on the geographic distribution of online trade, holding fixed total online spending. In this section, we consider the effect of sales tax on the overall propensity to shop online.

We start with a simple log-log representation of consumer demand for online purchases,

$$\log Q_{it} = \xi_{it} - \eta \log(P_{it}/\overline{P}_{it}), \tag{7}$$

where Q_{it} are counts of total online purchases by consumers in location i at time t, ξ_{it} captures local preferences and overall consumption, η is the price elasticity, and P_{it}

[19] Online Appendix Table A5 reports additional specifications that allow the same-state effect to vary across census regions and divisions. The range of the tax elasticity estimates remain similar.

[20] One important difference between our exercise and the one reported by Hortaçsu, Martínez-Jerez, and Douglas (2009) is that they focused only on buyers who were also eBay sellers, so that they could recover user locations through web-scraping. There are also some other differences between specifications. For instance, Hortaçsu, Martínez-Jerez, and Douglas measure interstate distance as the great circle distance between state capitals and intrastate distance as the population weighted distance between the two most populous cities in the state, whereas we measure distance as the average eBay transaction distance with the distance of each transaction computed using the distance between buyer and seller zip codes. As noted in the introduction, their paper also includes state sales tax in one set of regressions (Table 7, Models II and III). Their estimated tax effects are not directly comparable to ours, as they do not account for county and local taxes, use indicators for integer state tax levels instead of a continuous regressor, and interact tax rate with distance. To first approximation, their estimated tax effect is larger than ours, at least -10.

and $\overline{P}_{it}$ are, respectively, online and offline price indices.[21] Making the assumption that "own-location" purchases comprise only a small share of online purchases, but essentially all offline purchases, we can write $P_{it}/\overline{P}_{it} = (1 + \tau_{it})^{-1}R_{it}$, where R_{it} represents the relative online-to-offline prices before sales tax is imposed.

For our econometric model, we further assume that both the general level of online demand ξ_{it} and the pre-tax relative prices R_{it} can be decomposed into a location-specific component, a time component, and effects that are captured by observed covariates $\mathbf{z}_{it}$. So we have

$$\log Q_{it} = a_i + b_t + \mathbf{z}_{it}'\lambda + \eta \log(1 + \tau_{it}). \tag{8}$$

Implicit in this specification is the (conventional) assumption that targeted changes in state or local sales tax are passed through fully to consumers (Poterba 1996; Besley and Rosen 1999).

We start by attempting to use only cross-sectional variation, using county-level counts of eBay purchases during 2010. In panel A of Table 6 we report specifications that use cross-state and within-state variation in county-level tax rates, with and without a rich set of county-level controls (see the notes to Table 6 for details). The estimated tax effect is imprecise and varies greatly across specifications, indicating the difficulty of constructing suitable controls for local purchasing propensities.

Our preferred approach, therefore, is to rely on within-locality tax changes. The results are reported in panel B of Table 6. In our baseline and preferred specification (column 1 of the top part of panel B), we include fixed effects for each locality and each month, so that identification is based on changes in locality-level purchasing following a tax change as compared to the average change over the same time period for localities that did not experience a tax change. The results in columns 2 and 3 permit greater heterogeneity in time trends, with month-by-region and month-by-division fixed effects, respectively. Columns 4–6 repeat the same specifications, but also control for county-level unemployment rates, which slightly decrease the estimated elasticities.

One plausible concern is that the county may be too small of an area for the purpose of defining the applicable tax rate consumers face, as consumers may travel across county boundaries to make purchases, especially if nearby counties have much lower tax rates. In such counties, changes in tax rates may have a smaller impact as the relevant tax rate for many purchases is this associated with the lower-tax neighboring county. To address this possibility, the bottom part of panel B of Table 6 repeats the analysis, but for a sample from which we drop about a fifth of the counties, which border lower-tax counties on the other side of a state boundary. Indeed, consistent with this hypothesis, our estimated elasticities in the bottom panel are higher, by about 25 percent.

Taken together, our preferred estimate of η is around 1.8, meaning that a one percentage point increase in sales tax increases online purchasing by 1.8 percent. In comparison, Goolsbee's (2000a) baseline estimated elasticity using cross-sectional variation in tax rates was 2.3, increasing to 3.5 with the addition of

[21] Note that, for consistency with the previous section, one can think of P_{it} as the CES price index and Q_{it} as the CES aggregator of online consumption. In estimation, however, we will use overall purchase counts as our measure of Q_{it}.

Table 6—The Effect on Overall Online Purchasing

	Number of purchases in county, 2010					
	(1)	(2)	(3)	(4)		
Panel A. Identification off cross-sectional variation						
log(1 + effective tax)	−2.101	0.448	−5.145	−0.547		
	(0.395)	(0.236)	(2.312)	(1.136)		
Fixed effects	None	None	State	State		
Other controls	Population	All[a]	Population	All[a]		
Observations (counties)	3,050	3,037	3,050	3,037		
	Number of monthly purchases in county, 2008–2010					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel B. Identification off within-locality changes in sales tax						
Sample: all counties (109,944 county-month observations)						
log(1 + effective tax)	1.822	1.999	1.788	1.414	1.781	1.562
	(0.855)	(0.658)	(0.687)	(0.816)	(0.674)	(0.650)
Unemployment rate				0.629	0.506	0.518
				(0.188)	(0.188)	(0.189)
Fixed effects	County, month	County, month × region[b]	County, month × division[c]	County, month	County, month × region	County, month × division
Sample: excluding counties on state borders with cross-border tax avoidance opportunities[d] (89,388 county-month observations)						
log(1+effective tax)	2.296	2.531	2.167	1.821	2.278	1.934
	(0.872)	(0.720)	(0.678)	(0.810)	(0.690)	(0.702)
Unemployment rate				0.647	0.484	0.479
				(0.237)	(0.221)	(0.209)
Fixed effects	County, month	County, month × region	County, month × division	County, month	County, month × region	County, month × division

Notes: Table shows results from a Poisson regression where the dependent variable is total of eBay purchases in the county over 2010 (panel A) or every month from January 2008 to December 2010 (panel B). Standard errors are computed by a county-level block bootstrap with 50 replications.

[a] County-level controls include population, average income, gender (percent female), race (percent white, black, Asian), education (percent high school, some college, college, graduate degree), age (percent 0–9. 10–17, 18–29, 30–49, 50–69), and variables associated with internet connectivity (residential broadband connections, percent living in college housing, percent working in info industry, percent institutionalized).

[b] Region refers to the four census regions: Northeast, Midwest, South, and West.

[c] Division refers to the nine census divisions.

[d] In the bottom part of panel B we drop all counties (571 of 3,054) that are adjacent to state borders and for which at least one of the adjacent counties across the state border has a strictly lower sales tax rate.

more sophisticated controls. The elasticity for memory modules reported in Ellison and Ellison (2009), again identified off cross-sectional variation in state tax rates, is even higher, roughly 6 or 7.[22] While our estimate appears to be somewhat small relative to those reported previously, it nonetheless implies substantial effects of sales taxes on online trade. Given an average combined tax rate of about 7 percent, it suggests that sales tax effects might be responsible for boosting online purchasing by 10 percent or more.

[22] The estimates in Ellison and Ellison (2009) concern differences in purchasing from their California retailer by residents of high and low tax states, and hence combine online-offline and cross-state substitution effects. To the extent that each state represents only a small share of online sales, however, their number should reflect mainly online-offline substitution.

D. *Combined Effects of Sales Tax Changes*

So far we have considered the two margins of substitution—online-offline and online cross-state—separately. To think about the possible effect of changes in sales taxes, or changes in the current legal regime, it is useful to combine the effects. To do this, we combine our model of overall online purchasing (equation 7) with our model of how online spending is distributed (equation 5), noting that in the latter we can represent overall online expenditure w_i as $P_i Q_i$.[23]

Now, consider the effect of an increase in state i's sales tax τ_i, which under the current legal regime will be applied to both offline and in-state online purchases. To the extent that state i represents a relatively small share of both online demand and sales, we can assume that this will have no direct effect on either online (pre-tax) prices or on i's online price index P_i, and we continue to assume that offline sellers fully pass through the tax to consumers.[24] Then we have

$$(9) \qquad \frac{\partial \log Q_i}{\partial \log(1+\tau_i)} \approx \eta^{25}$$

and, using the fact that $\partial \log w_i / \partial \log(1+\tau_i) \approx \eta$,

$$(10) \qquad \frac{\partial \log q_{ij}}{\partial \log(1+\tau_i)} \approx -\sigma \mathbf{1}\{i=j\} + \eta.$$

So, if we consider a 1 percentage point decrease in state sales tax (such as occurred in California on July 1, 2011), our baseline estimates suggest roughly a 1.5–2 percent decrease in online purchases by state residents, and a corresponding decrease in cross-state online purchases, but a 3–4 percent increase in online purchases by state residents from home-state sellers.

A more sophisticated analysis might relax the "small-share" assumption. To see that it is not particularly important, let $x_{ii} = (p_{ii} q_{ii})/w_i$ denote the share of online expenditure that state i devotes to home-state purchases. With CES demand, $\partial \log P_i / \partial \log(1+\tau_i) = x_{ii}$, so if x_{ii} is not trivial, an increase in τ_i will affect online (post-tax) prices as well as offline prices. Instead of the expressions above, we have $\partial \log Q_i / \partial \log(1+\tau_i) = \eta(1-x_{ii})$, and

$$(11) \qquad \frac{\partial \log q_{ij}}{\partial \log(1+\tau_i)} = -\sigma 1\{i=j\} + \eta + (\sigma - \eta) x_{ii}.$$

[23] Note that for this connection to be tight, then as noted in footnote 21 above, we need to interpret Q_{it} in the overall online demand model as the CES aggregate of online consumption, not as a count of online purchases as we did in our empirical implementation.

[24] Note that more generally, if $p_{ij} = (1 + \tau_i \mathbf{1}_{\{i=j\}}) p_j$, and sellers do not change pre-tax prices in response to a change in τ_i, then under our CES specification $\partial \log P_i / \partial \log(1+\tau_i) = (p_{ii} q_{ii})/w_i = x_{ii}$. The assumption that $x_{ii} \approx 0$ is a reasonable approximation for most states. Using expenditure shares for eBay, the median state has $x_{ii} = 0.03$, and only two states (California and New York) have $x_{ii} > 0.10$ (see online Appendix Table A6, column k).

[25] Note that $\tau_i \approx \log(1+\tau_i)$ for lower tax rates, so the semi-elasticity with respect to the tax rate $\partial \log Q_i / \partial \tau_i$ is approximately equivalent to $\partial \log Q_i / \partial \log(1+\tau_i)$, the elasticity with respect to the tax multiple $(1+\tau_i)$.

To see that this makes little difference, note that for most states $x_{ii} < 5$ percent and even for California x_{ii} is only 0.21, so that $\partial \log Q_i / \partial \log(1 + \tau_i)$ is still $1.8 \cdot 0.79 = 1.4$.

To illustrate the magnitude of the estimate, consider a large structural change such as imposing a requirement that sales tax be collected on all interstate online sales. While considerable caution should be placed on such a large extrapolation from the environment generating our estimates, a back-of-the-envelope calculation is interesting. As of January 1, 2010, the population-weighted average sales tax in the United States was about 7.3 percent. Taken literally, our estimates imply that if that tax rate were applied to all interstate online transactions, and online prices responded in the same way that offline prices do to the tax changes in our data, overall online purchasing would fall by about 12 percent. This decline could be substantially lower if some of the tax change would not be fully passed through into prices, and instead get absorbed into sellers' margins, as may be the case in response to such a major tax change.

III. Conclusions

Internet sales taxes have been the subject of considerable attention since the beginning of Internet commerce. This paper has used detailed data from eBay to offer two complementary pieces of evidence on how sales taxes affect online browsing and purchasing behavior.

Our first set of results show how individual shoppers respond to sales taxes at the item level, using a research design based on individual-level "tax surprises". We found that purchases by interested buyers fall by roughly two percent for every one percentage point increase in the sales tax charged by the seller. The degree of sensitivity appears to vary in natural ways depending on the type of item, and the application of sales tax appears to generate substitution toward untaxed items. Moreover, to the extent that consumers pay less attention to taxes than to base prices, our estimates can be interpreted as providing an informative lower bound on retail price elasticities for interested buyers.

Our second set of results address the effect of sales tax policy at the state and national level. These results are based on the relationship between taxes and aggregate online trade flows (on eBay). Using the considerable cross-state variation in sales tax rates as a source of identification, we estimated that, holding fixed the overall online spending of state residents, a one percentage point increase in a state's sales tax leads to a 3–6 percent decrease in online purchasing from home-state sellers. We also used changes in state and local sales taxes over time to estimate the overall effect of sales taxes on online purchasing. We find an elasticity of online purchasing with respect to sales tax of around 1.8, a substantial sensitivity but only about half the magnitude reported by Goolsbee (2000a). Combining these estimates, a one percentage point increase in a state's sales tax leads to an increase of just under 2 percent in online purchasing from other states, and a 3–4 percent decrease in online purchasing from home-state sellers.

We view the two analyses as complementary but the estimates are not directly comparable, as they attempt to measure conceptually different tax sensitivities. In the working paper version (Einav et al. 2012), we provided a framework in which the estimates could be reconciled, although doing so in practice is complicated by the different data samples underlying the estimates.

Our analysis has focused largely on consumer behavior. An interesting avenue for future research would be to explore how sales tax treatment affects online sellers' decisions about where to locate. Amazon, for instance, assiduously avoided establishing tax presence in California and other large states for many years.[26] More generally, the current structure of sales taxes creates a trade-off. Locating close to demand reduces transportation costs and may boost demand if buyers prefer nearby or "home-state" sellers, but it also means collecting more sales tax. Changes in national sales tax policy would shift this trade-off, and might well affect location decisions by online retailers, as well as consumer behavior.

REFERENCES

Alm, James, and Mikhail I. Melnik. 2005. "Sales Taxes and the Decision to Purchase Online." *Public Finance Review* 33 (2): 184–212.

Anderson, Eric T., Nathan M. Fong, Duncan I. Simester, and Catherine E. Tucker. 2010. "How Sales Taxes Affect Customer and Firm Behavior: The Role of Search on the Internet." *Journal of Marketing Research* 47 (2): 229–39.

Anderson, James E. 2011. "The Gravity Model." *Annual Review of Economics* 3 (1): 133–60.

Ballard, Charles L., and Jaimin Lee. 2007. "Internet Purchases, Cross-Border Shopping, and Sales Taxes." *National Tax Journal* 60 (4): 711–25.

Besley, Timothy J., and Harvey S. Rosen. 1999. "Sales Taxes and Prices: An Empirical Analysis." *National Tax Journal* 52 (2): 157–78.

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review* 99 (4): 1145–77.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan. 2013. "Consumer Price Search and Platform Design in Internet Commerce." Unpublished.

Einav, Liran, Dan Knoepfle, Jonathan D. Levin, and Neel Sundaresan. 2012. "Sales Taxes and Internet Commerce." National Bureau of Economic Research Working Paper 18018.

Einav, Liran, Dan Knoepfle, Jonathan Levin, and Neel Sundaresan. 2014. "Sales Taxes and Internet Commerce: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.104.1.1.

Einav, Liran, Theresa Kuchler, Jonathan D. Levin, and Neel Sundaresan. 2011. "Learning from Seller Experiments in Online Markets." National Bureau of Economic Research Working Paper 17385.

Ellison, Glenn, and Sara Fisher Ellison. 2009. "Tax Sensitivity and Home State Preferences in Internet Purchasing." *American Economic Journal: Economic Policy* 1 (2): 53–71.

Goolsbee, Austan. 2000a. "In a World without Borders: The Impact of Taxes on Internet Commerce." *Quarterly Journal of Economics* 115 (2): 561–76.

Goolsbee, Austan. 2000b. "Internet Commerce, Tax Sensitivity, and the Generation Gap." In *Tax Policy and the Economy,* Vol. 14, edited by James M. Poterba, 45–65. Cambridge, MA: MIT Press.

Goolsbee, Austan, Michael F. Lovenheim, and Joel Slemrod. 2010. "Playing with Fire: Cigarettes, Taxes, and Competition from the Internet." *American Economic Journal: Economic Policy* 2 (1): 131–54.

Hoch, Stephen J., Byung-Do Kim, Alan L. Montgomery, and Peter E. Rossi. 1995. "Determinants of Store-Level Price Elasticity." *Journal of Marketing Research* 32 (1): 17–29.

Hortaçsu, Ali, F. Asís Martínez-Jerez, and Jason Douglas. 2009. "The Geography of Trade in Online Transactions: Evidence from eBay and MercadoLibre." *American Economic Journal: Microeconomics* 1 (1): 53–74.

Maguire, Steven. 2013. "State Taxation of Internet Transactions." Congressional Research Service Report R41853. http://www.fas.org/sgp/crs/misc/R41853.pdf (accessed August 26, 2013).

Manzi, Nina. 2010. "Use Tax Collection on Income Tax Returns in Other States." Policy Brief, Research Department, Minnesota House of Representatives.

[26] As we were writing this paper, Amazon agreed to collect sales taxes on California sales starting in September 2012 (and appears to be reaching similar deals with other states), and consequently plans to open three new distribution centers in California.

Poterba, James M. 1996. "Retail Price Reactions to Changes in State and Local Sales Taxes." *National Tax Journal* 49 (2): 165–76.

Scanlan, Mark A. 2007. "Tax Sensitivity in Electronic Commerce." *Fiscal Studies* 28 (4): 417–36.

Smith, Michael D., and Erik Brynjolfsson. 2001. "Consumer Decision-Making at an Internet Shopbot: Brand Still Matters." *Journal of Industrial Economics* 49 (4): 541–58.

US Census Bureau. 2011. "E-Stats Report." http://www.census.gov/econ/estats (accessed February 26, 2012).

Varian, Hal R. 2000. "Taxation of Electronic Commerce." *Harvard Journal of Law & Technology* 13 (3): 639–52.