University of Toronto
Department of Economics

Working Paper 423

The effects of land transfer taxes on real estate markets: Evidence from a natural experiment in Toronto

By Ben Dachis, Gilles Duranton and Matthew A. Turner

February 14, 2011

The effects of land transfer taxes on real estate markets: Evidence from a natural experiment in Toronto [§]

Ben Dachis[‡]

C.D. Howe Institute

Gilles Duranton[*]

University of Toronto

Matthew A. Turner[†]

University of Toronto

This draft: 4 February 2011

ABSTRACT: Taxes levied on the sale or purchase of real estate are pervasive but little studied. By exploiting a natural experiment arising from Toronto's imposition of a Land Transfer Tax (LTT) in early 2008, we estimate the impact of real estate transfer taxes on the market for single family homes. Our data show that Toronto's 1.1% tax caused a 15% decline in the number of sales and a decline in housing prices about equal to the tax. Relative to an equivalent property tax, the associated welfare loss is substantial, about $ 1 for every $ 8 in tax revenue. The magnitude of this welfare loss is comparable to those associated with better known interventions in the housing market. Unlike many possible tax reforms, eliminating existing LTTs in favour of revenue equivalent property taxes appears straightforward.

Key words: Land transfer tax, property tax, land regulation

JEL classification: R21, R51

[‡]C.D. Howe Institute, 67 Yonge Street, Suite 300, Toronto, Ontario, Ontario M5E 1J8, Canada (e-mail: Ben@cdhowe.org.

[*]Department of Economics, University of Toronto, 150 Saint George Street, Toronto, Ontario M5S 3G7, Canada (e-mail: gilles.duranton@utoronto.ca; website: http://individual.utoronto.ca/gilles/default.html). Also affiliated with the Center for Economic Policy Research, and the Center for Economic Performance at the London School of Economics.

[†]Department of Economics, University of Toronto, 150 Saint George Street, Toronto, Ontario M5S 3G7, Canada (e-mail: mturner@chass.utoronto.ca; website: http://www.economics.utoronto.ca/mturner/index.htm).

1. Introduction

Taxes levied on the sale or purchase of real estate are pervasive but little studied. Such taxes increase the cost of buying or selling a house and to avoid paying these taxes households will stay in houses that are too big, too small, or too far from their workplace. Hence, we expect these taxes to reduce the volume of real estate transactions and to entail a welfare loss. By exploiting a natural experiment arising from Toronto's imposition of a Land Transfer Tax (LTT) on real estate purchases in early 2008, we estimate the impact of real estate transfer taxes on the market for single family homes. Our data show that Toronto's 1.1% tax caused a 15% decline in the number of sales and a decline in housing prices about equal to the tax. Relative to an equivalent property tax, the associated welfare loss is substantial, about $ 1 for every $ 8 in tax revenue according to our calculation. While the magnitude of the welfare loss from Toronto's LTT partly reflects a pre-existing provincial land transfer tax, our estimate is comparable to those associated with other well recognised interventions in the housing market. For instance, the results of Glaeser and Luttmer (2003) suggest a welfare loss from rent control in New York above m$ 200 per year. We obtain about the same figure by extrapolating our finding for Toronto to New York.[1] In short, our estimates suggest that scrapping current LTTs in favour of revenue equivalent property taxes could result large annual welfare gains.

Most jurisdictions in most developed countries impose land transfer taxes. In the United States in 2004, 34 states and the District of Columbia imposed some form of property transfer tax, while for a number of US cities and states, property transfer taxes account for a non-trivial share of revenue. In Canada most provinces impose a land transfer tax. Australian states levy a similar tax, as do the UK and France. In spite of the evident importance of land transfer taxes to public finance, their effects are little studied.

Many other public policies also directly affect the cost of buying and selling property. Most importantly, a key part of the 2009 US Housing Stimulus Bill is a provision allowing homebuyers to reduce their taxes by 10 percent of the price of their new home, up to $ 8,000, for a limited period of time. In effect, this is the opposite of the Toronto LTT (though capped, limited in time, and dependent on buyer's income). Also related, California's Proposition 13 freezes property taxes when a house changes hands. In an environment where real estate is appreciating, this means that tax burdens increase when the house changes hands, and thereby drives up transactions costs. Our results provide a basis for assessing the impact that these sorts of policies will have on real estate markets and mobility.

Our identification strategy is a variant of the regression discontinuity design. Figures 1 and 2 provide a heuristic introduction to this strategy. The top panel of figure 1 illustrates the relationship between the LTT and the change in transaction volume between the two six month periods beginning February of 2007 and February 2008. The heavy line in the top panel describes the change in the number of transactions in an average postal code as the distance from the Toronto border increases. The dotted lines give 5% and 95% bootstrapped pointwise confidence bounds. As we

---

[1]The calculation for the welfare loss of rent control in New York for Glaeser and Luttmer is in their working paper (Glaeser and Luttmer, 1997).

move left from zero the figure describes suburban postal codes further outside the Toronto border, as we move right the figure describes downtown postal codes further inside the border. This figure shows that Toronto residents are less mobile relative to their nearby suburban neighbours after the imposition of the LTT than before. Details about the construction of figure 1 (and figure 2) are available in a technical appendix.

The bottom panel of figure 1 shows time series variation in cross-border changes in transaction volumes. Specifically, this figure considers the a five kilometre band on either side of the Toronto border and calculates the difference in the number of monthly transactions for an average postal code in each region each region as a percentage of the Toronto mean. The dots in this figure show the 44 resulting data points. In this figure we see a spike in transactions volume before the LTT and sharp decrease afterwards.

Figure 2 is similar to figure 1, but illustrates the relationship between the LTT and prices. The top panel of figure 2 illustrates the relationship between the LTT and the change in transaction price between the two six month periods beginning February of 2007 and February 2008. The heavy line in the top panel describes the change in the price of a transaction in an average postal code as a percentage of the pre-LTT 2007 price, as the distance from the Toronto border varies. The dotted lines give 5% and 95% bootstrapped pointwise confidence bounds. As we move left from zero the figure describes areas farther outside the Toronto border, as we move right the figure describes areas farther inside. This figure shows little systematic cross-border variation in prices around the time of the LTT. The bottom panel of figure 2 shows time series variation in cross-border changes in transaction prices. Specifically, this figure considers the five kilometre band on either side of the Toronto border and calculates the difference in monthly transaction price for an average postal code in each region each region as a percentage of the Toronto mean. The dots in this figure show the 44 resulting data points. In this figure we see a spike in transactions price before the LTT, and a small decrease afterwards.

Figures 1 and 2 suggest that the imposition of the LTT caused two discrete changes in Greater Toronto real estate market, one on the date the LTT is imposed, and the other at the Toronto-suburban municipal boundary. With this said, the shape of the curves illustrated in these figures is somewhat sensitive to econometric technique. Our econometric strategy relies on an elaboration of the regression discontinuity design (Hahn, Todd, and Van der Klaauw, 2001) that exploits both discontinuities to arrive at robust estimates of the effect of the LTT.

2. Background

A. *The Toronto Land Transfer Tax*

Expenditures by the Toronto municipal government increased by about 9% between 2006 and 2007. This increase resulted partly from the 'downloading' of provincially mandated programs and partly from a rapid and sustained increase in operating expenditures. The increase in expenditure was offset by a revenue increase of about 3%. (City of Toronto, 2008). The resulting budget shortfall was about half a billion dollars.
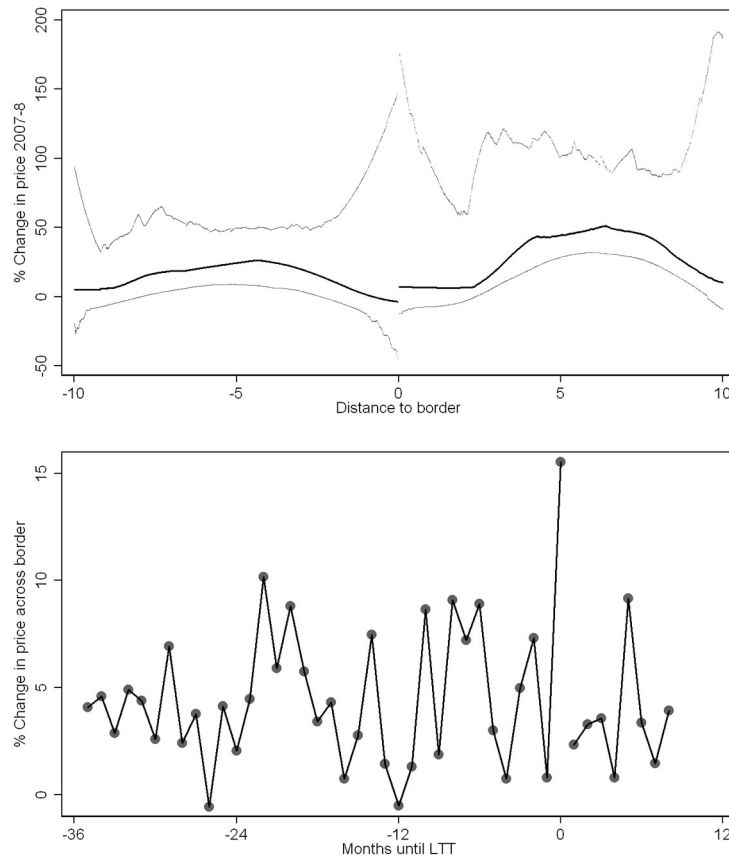
Figure 1. Volume of residential real estate transactions across the Toronto border.
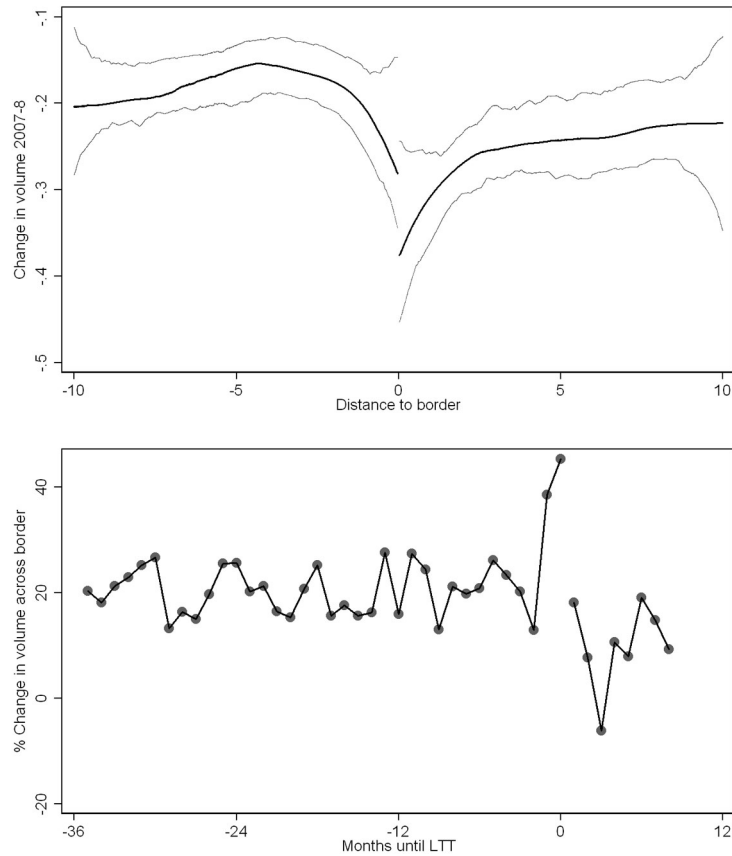
Top panel: Vertical axis is change in number of units transacted in an average postal code after the imposition of the LTT. Horizontal axis is distance from the Toronto municipal border, negative distances are suburban, positive distances are Toronto. Solid line gives mean percentage change in volume for the mean postal code as a function of distance. Dotted lines are 95% and 5% bootstrapped pointwise confidence bounds.
Bottom panel: Vertical axis is percentage change in mean postal code transaction volume from crossing the Toronto border in the months before and after the imposition of the LTT. Horizontal axis counts months from the imposition of the LTT, with negative numbers indicating earlier months and positive numbers indicating later months.

Figure 2. Price of residential real estate transactions across the Toronto border.

Top panel: Vertical axis is percentage change in the price units transacted in an average postal code after the imposition of the LTT. Horizontal axis is distance from the Toronto municipal border, negative distances are suburban, positive distances are Toronto. Solid line gives mean percentage change in price for the mean postal code as a function of distance. Dotted lines are 95% and 5% bootstrapped pointwise confidence bounds.

Bottom panel: Vertical axis is percentage change in mean postal code transaction price from crossing the Toronto border in the months before and after the imposition of the LTT. Horizontal axis counts months from the imposition of the LTT, with negative numbers indicating earlier months and positive numbers indicating later months.

4

Prior to 2007, Toronto's budget was usually balanced or even slightly in surplus. The mayoral election fought late 2006 made barely any mention of budgetary issues. Important issues at the time were a rising wave of gang-related crime, garbage collection, and waterfront renewal.[2] The incumbent mayor, David Miller, even promised to keep residential property tax increases in line with inflation. He was ultimately reelected. Therefore, although the city's budget had been tight for some time, 2007's budget deficit surprised most Toronto residents. Furthermore, we note that this budget shortfall was unique to Toronto because of its unusually generous pay settlement for city workers the previous year and because the downloading of social expenditure by the Province affected the city of Toronto more than surrounding municipalities.

Shortly after the election the newly elected Toronto City Council began to address the budget, and, to keep the mayor's campaign promise, considered alternative forms of taxation. This was possible because the 'City of Toronto Act' (38th Legislature of Ontario, 2006), which took effect on January 1, 2007, granted the city of Toronto the authority to impose a number of additional taxes. Among them was a 'municipal land transfer tax'. On July 16, 2007 the Toronto City Council narrowly defeated the implementation of such a tax. Rather, it voted to defer a decision until October 2007. In response the Mayor announced emergency cuts to municipal services until additional revenue could be found. The Toronto City Council ultimately approved a land transfer tax on October 22, 2007.

The land transfer tax affected all real estate transactions which occurred after January 1, 2008, and closed after February 1, 2008. The final rate schedule for Toronto's LTT is given in Table 1, along with the corresponding provincial tax rates. The LTT is paid by the buyer. First-time homebuyers receive a partial exemption.

Table 1. City of Toronto and Province of Ontario Land Transfer Tax Rates

City of Toronto LTT Tax Rate by Value ($) (Effective Feb. 1, 2008)		Province of Ontario LTT Tax Rate by Value ($) (Effective May 7, 1997)	
0 − 55,000	0.5%	0 − 55,000	0.5%
55,000 − 400,000	1.0%	55,000 − 250,000	1.0%
400,000+	2.0%	250,000 − 400,000	1.5%
		400,000+	2.0%

Sources: Municipal Land Transfer Tax, City of Toronto, http://www.toronto.ca/taxes/mltt.htm; Provincial Land Transfer Tax, Historical Land Transfer Tax Rates, Province of Ontario. Note: For the Municipal LTT, exemptions are given to first-time homebuyers for the value of a purchase under $ 400,000 and for the provincial LTT exemptions are given to first-time home buyers for the value of a purchase under $ 227,500.

We will ultimately estimate the effect of the LTT by comparing the changes to the Greater Toronto real estate market before and after the imposition of the LTT across the Toronto border. The validity of this strategy rests on two assumptions. First, that the real estate market did not

---

[2]See for instance the summary and extracts of the main mayoral debate at `http://www.citynews.ca/news/news_4634.aspx`.

anticipate the tax, and second that no other policy change differentially affected the Toronto and suburban real estate market at the same time as the LTT was imposed.

The evidence suggests that the real estate market could not have anticipated the LTT. The Toronto City Council considered a number of different taxes as responses to the city's budget problems during the fall of 2007, not only a land transfer tax. In addition, the relevant Council votes were close and uncertain, even on the day of the vote. With this history, it is unlikely that markets anticipated the LTT.

In spite of this, it is worth asking how our results would be affected if the market did anticipate the LTT. If participants in the Toronto property market anticipated the LTT before the council vote in late October 2007, we would expect to see a spike in transactions prior to that date along with a small decline in prices. In fact, while we observe a significant spike in transactions volume between the vote of the tax and its the imposition, we do not observe changes in the real estate market before the vote. This suggests that the market did not anticipate the LTT and hence, that our empirical strategy is reasonable.[3]

While we cannot prove that no other event, contemporaneous with the LTT differentially affected Toronto and suburban real estate markets, we can find no evidence for such an event. In Toronto, the City Council approved the LTT together with a $ 60 municipal supplement to the vehicle registration fee. This fee is small and only indirectly related to the property market. We do not think it could plausibly affect it. Property taxes increased by about the same amount in Toronto and surrounding municipalities.[4] Although we note that our theoretical analysis in section 8 suggests that ordinary property tax rates should not affect real estate transaction volumes.

In summary, we are looking at the effect of the unexpected introduction of a new tax affecting Toronto. The revenue generated by this tax was not used to increase service provision but instead to fill the gap left by an unusually generous pay settlement with city employees and a rise in expenditure induced by the downloading of provincial expenditure that affected Toronto disproportionately. Municipalities surrounding Toronto did not face the same need to increase their revenue, nor were their governments permitted to impose a LTT.

B. *Land Transfer Taxes outside Toronto*

Most jurisdictions in most developed countries impose land transfer taxes. According to the Federation of Tax Administrators (Federation of Tax Administrators, 2006), in the United States in 2004, 34 states and the District of Columbia impose some form of property transfer tax. Among these, New Hampshire, Delaware, and the District of Colombia have the highest rates, at 1.5%, 2%, and 2.2%, respectively. Many US municipalities also impose property transfer taxes, often at rates equal to or higher than those imposed by the states. For example, Chicago's land transfer tax is

---

[3]Alternatively, it could be that residents knew well before October 2007 that local taxation would increase significantly (through, say, property taxation or user fees) but that the choice of the LTT as an instrument to address the budget shortfall came as a surprise. Although we view this possibility as remote, it would not affect our analysis on volumes of transactions as made clear below. It would nonetheless affect the interpretation of our price findings. They would need to be thought as the effect of a land transfer tax relative to another form of taxation instead of a pure LTT effect.

[4]The property tax rate increased from .0085 to .0087 form 2007 to 2008, an increase of 2.5%. The corresponding rates for Brampton, Markham, Mississauga, Pickering, Richmond Hill and Vaughan were; 3.4%, 2.9% 3.2% 4.6% 3.1% 2.9%.

1.3%, New York's is between 1.4 and 1.9%, and Philadelphia's is 3%. Given these high rates, it is not surprising that for some cities and states, property transfer taxes account for a non trivial share of revenue; 2.9% of total revenue for the State of New Hampshire, 3.4% for the State of Florida, 3.8% for the District of Columbia, 4.8% for New York, 6.2% for Chicago, 11.0% for Philadelphia, and 22.2% for Oakland (CA).

In Canada, the provincial governments of British Columbia, Manitoba, Ontario, Quebec, New Brunswick, Prince Edward Island, and some Nova Scotia municipalities impose a land transfer tax. Australian states levy a similar tax with rates up to 6.75 percent. In the UK, the LTT, there called a 'stamp duty tax,' ranges from 1 to 4 percent, and in 2007, accounted for Bn £6.5, or about 1.25% of the budget of the British government. Countries in Southern Europe also rely on land transfer taxes. In France the rate for residential property was above 10 percent during the 1980s, but was reduced to 4.8 percent in 1999. Revenue from the property transfer tax represented about 3% of the budget of the French government in 2008.

In sum, land transfer taxes are economically important. They are pervasive and provide an important source of revenue for local and national governments.

3. Literature

A. *Other analyses of land transfer taxes*

Given the economic importance of land transfer taxes, there is surprisingly little systematic analysis of their effects. In a first study, Benjamin, Coulson, and Yang (1993) examine the effects of an increase in the Philadelphia Real Estate Transfer Tax. On the basis of around 350 transactions, they find that properties located inside Philadelphia decrease in price relative to properties outside and that this decrease is much larger than the tax increase. Unfortunately, their data do not allow them to examine volumes of transactions.

Ioannides and Kan (1996) propose a model of residential mobility and housing tenure choice in which purchasing a house is subject to a proportional transaction cost. Empirically, after conditioning out much individual heterogeneity, they find no association between property values and the propensity to move either to rental or owner occupied housing. In the context of their model, they interpret this finding as an indication that proportional transactions costs do not affect mobility. Van Ommeren and Van Leuvensteijn (2005) use the same theoretical framework and more detailed data about transaction costs in the Netherlands. Although transaction costs do not vary over time and over space in their data, they use again differences between renters and homeowners. They find that an LTT-equivalent tax with the same rate as implemented in Toronto would decrease mobility by 8 to 19 percent. Finally, Nowlan (2007) provides an ex-ante evaluation of the Toronto LTT. He argues, from theoretical foundations and hypothetical estimates of the demand and supply for housing, that the LTT would have only a minimal impact on the Toronto real estate market.

In contrast with these studies, our estimation is based on a sample that records the large majority of single-family home transactions in the greater Toronto area for 25 months before the imposition of the LTT and for the first seven months afterward. We exploit detailed information

about the timing of transactions and their location to identify the effect of this tax. Thus, we can reasonably hope to arrive at definitive estimates of the LTT's effects.

B. *Other border studies*

The RD design is increasingly popular and is used to investigate, for example, the effect of class sizes on educational attainment (Angrist and Lavy, 1999), the effect of changes in social assistance programs on employment (Lemieux and Milligan, 2008), or the effect of mayoral party affiliation on municipal policies (Ferreira and Gyourko, 2009). The essential identifying assumption is that, absent the change in policy, the dependent variable would change continuously at the cut-off (Hahn *et al.*, 2001). Best practice is described in Imbens and Lemieux (2008).

The method has also been used to investigate the effect of policies which vary over physical space as one crosses from one administrative unit to another. In this case, the cut-off of interest is an administrative boundary. Holmes (1998) looks at the impact of changes in right-to-work laws on manufacturing employment near state borders. Black (1999), Gibbons and Machin (2003), and Bayer, Ferreira, and McMillan (2007) look at the effect of changes in property values near school district boundaries. Duranton, Gobillon, and Overman (2011) look at the effect of changes in municipal taxation across municipal boundaries on the behaviour of firms near these boundaries.

While our identification strategy is clearly related to those used in the papers listed above, it differs in important respects. Because Toronto imposed a land transfer tax and none of the surrounding municipalities did, our data allow us to isolate the effects of a change in a dichotomous variable at a single border; the effect of imposing a land transfer tax on houses in Toronto along the Toronto-Suburban border. On the other hand, in Black (1999), Gibbons and Machin (2003), and Bayer *et al.* (2007) the policy variable of interest is continuous, 'school district quality', for Holmes (1998), it is an index measuring the extent to which a state is 'pro-business', and for Duranton *et al.* (2011) it is local tax rates. Loosely, these authors identify the effect of their chosen policy in a two-step process. In the first, they estimate the discontinuity of interest at many borders, and in the second they examine the correlation between the magnitude of these cross-border discontinuities and the corresponding cross-border change in the policy variable of interest. Thus, most extant border studies uses variation across administrative borders in a way that is less similar than it may appear at first sight. Because our policy variable is dichotomous and we consider only a single border, we are much closer to the idealised description of the RD design described by Hahn *et al.* (2001) and Imbens and Lemieux (2008) or used to the study of the effects of social assistance by Lemieux and Milligan (2008).

4. Data

Our estimations are based on data that describe the sale of 139,266 single-family houses in the greater Toronto area that were listed with the Multiple Listing Service (MLS) between January 2006 and August 2008. In particular, our study area includes Toronto and all nearby municipalities: Mississauga, Brampton, Vaughn, Richmond Hill, Markham, and Pickering. A map of the study area is provided in figure 3. For each transaction we observe the sale price, the date when the

Figure 3. Study area
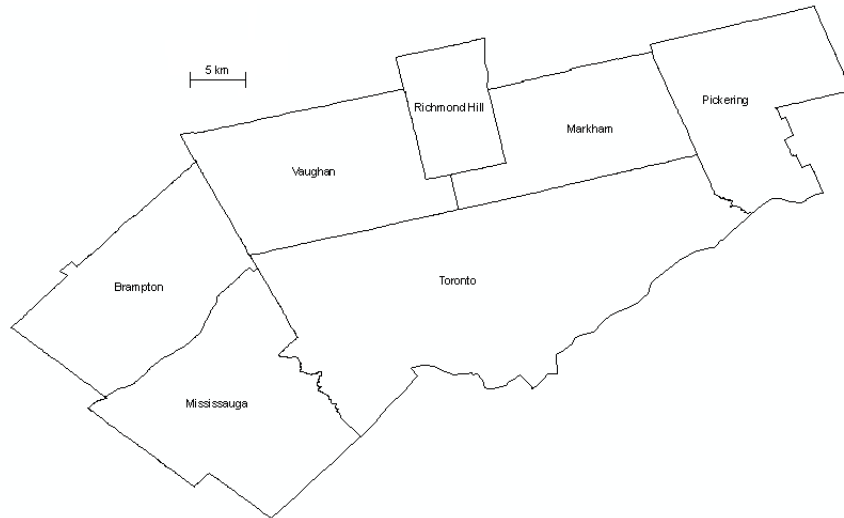
Our study area is the City of Toronto and the immediately surrounding municipalities of Mississauga, Brampton, Vaughn, Richmond Hill, Markham and Pickering. Lake Ontario is to the South of Mississauga, Toronto, and Pickering.

contract was signed, the date the transaction closed, and the property's postal code. We also observe about 40 house characteristics for each sale. Specifically, we have binary variables describing: heating; garage; basement; whether the house is attached; number of stories; construction type; the presence of a family room; and fireplace, along with counts of: bedrooms; bathrooms; kitchens; rooms; parking spaces; square feet and lot size. We exclude condominiums from our analysis.[5]

To determine whether a sale is subject to the LTT, we first verify whether the property is in Toronto or not. We use GIS software to assign to each property the latitude and longitude of the centroid of the property's postal code. We then use these coordinates to check whether the property lies within Toronto municipal boundaries, and also to calculate the distance to the Toronto border. Figure 4 presents a map of a portion of the Toronto border along with outlines of census tracts and postal codes. In this figure, the Toronto border is a wide light gray line. Census tract boundaries are medium-width lines in medium gray. Postal code boundaries are thin black lines. Postal code centroids are black points. In Toronto, postal codes are small, typically one block along one side of a street. As can be seen from the scale bar in figure 4's upper left hand corner, postal code centroids are generally within a few hundred metres of a property, and often much closer. Thus, we can be confident that we are calculating properties' locations accurately.

---

[5]Condominiums are more likely to be purchased by first time buyers exempt from much of the LTT. Moreover, new condominiums are often sold directly by developers who bypass MLS system. Therefore, not only does our data describe the condominium market less well than the market for single family homes, but the market for single-family homes and the market for condominiums appear to be separate: pooling data on single family home sales with data on condominium sales is not appropriate.

Figure 4. . A map showing an area along Toronto's northern border.
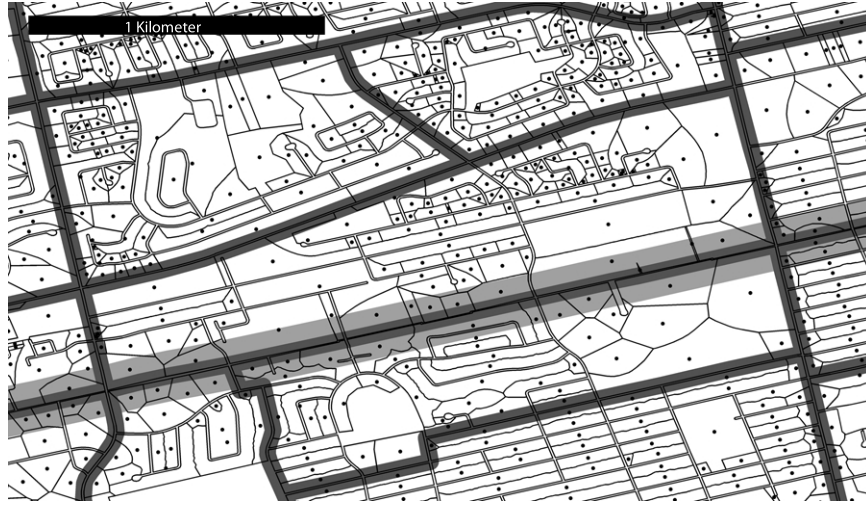
The Toronto border is a wide light gray line. Census Tract boundaries are medium width lines in medium gray. Postal Code boundaries are thin black lines. Postal code centroids are black points. Source: University of Toronto GIS Library.

5. Econometric model

Our econometric model is a hybrid of a regression discontinuity design and a differences-in-differences estimation. Like differences-in-differences, we compare the change in transaction volume for tracts in Toronto that are 'treated' with the LTT to untreated suburban tracts. As in a regression discontinuity design, we calculate our differences as close to the treatment threshold, the Toronto border, as is possible. Because our approach resembles a differences-in-differences estimate, we are concerned about the possibility of differential trends in transaction volume among treated and untreated tracts. For example, if sales rise faster in Toronto than in the suburbs, we will confound this differential trend with the effect of the LTT. Because our approach resembles a regression discontinuity design, we are concerned about the possibility of a gradient in transaction volumes. For example, if the rate of transactions rises as we cross from Toronto into the suburbs, this slope may confound our estimates of the change in transactions volumes caused by the LTT.

As made clear below by our estimating equations, our approach boils down to a differences-in-differences estimation for observations on a narrow band on both sides of the Toronto border. In the rest of this section, we further develop the econometrics that underlies our approach.

To proceed, let t denote time, with $t < 0$ being before the imposition of the LTT, and $t > 0$ being after. Let x denote distance from the border, with $x > 0$ being Toronto, and $x < 0$ being suburban. Define indicator variables based on x and t,

$$\chi^{\text{LTT}} = \begin{cases} 1 & \text{if} \quad t \geq 0 \\ 0 & \text{else} \end{cases}$$

$$\chi^{\text{TO}} = \begin{cases} 1 & \text{if} \quad x \geq 0 \\ 0 & \text{else.} \end{cases}$$

That is, χ^{TO} is one in Toronto and zero in the suburbs, while χ^{LTT} is one after the imposition of the LTT, and zero before. If both indicators are one then we restrict attention to Toronto after the imposition of the LTT.

Let $v(x,t)$ denote the volume of real estate transactions at a particular location and time. We can decompose this function into five pieces: $g(x,t)$ a continuous latent volume surface describing transactions in the absence of the LTT or other cross-border policy differences; $\alpha_{\text{LTT}}\chi^{\text{LTT}}$ a jump in the volume surface that occurs at the time of the LTT, e.g., a market downturn or bad weather; $\alpha_{\text{TO}}\chi^{\text{TO}}$ a jump in the volume surface that occurs at the Toronto border; $\lambda\chi^{\text{TO}}\chi^{\text{LTT}}$, the effect of the LTT that is the quantity of interest; and finally, $\epsilon(x,t)$, a mean zero error term whose properties we discuss below. To ease exposition, suppose that g is a second order polynomial, $g(x,t) = g_0 + g_x x + g_t t + g_{xx} x^2 + g_{tt} t^2 + g_{xt} xt$.

We now adjust this notation so that it describes our data rather than an abstract surface. To accomplish this, let i index our unit of observation, usually an individual postal code, but sometimes a distribution area or census tract. Let t index months. By analogy to the notation developed above, let v_{it} denote the count of sales in unit i and month t and let x_{it} denote the mean distance from a house in unit i to the Toronto border.[6] Let z_{it} denote a vector describing the mean characteristics of houses sold in unit i in month t. It is also useful to define δ_i as an indicator variable for each postal code i. Finally, we let ϵ_{it} denote the mean zero component of volume determined by unobserved factors, and χ_{it}^{TO} and χ_{it}^{LTT} be indicator variables corresponding to χ^{TO} and χ^{LTT}.

Ideally, we would draw our sample from a region of space and time close enough to $(x,t) = (0,0)$ that we could ignore all terms of $g(x,t)$ except the constant. To obtain a usefully large sample, however, we must sample from a region of space and time large enough that variation of the latent volume surface cannot not be ignored. In this case we have,

$$v_{it} = [g_0 + g_x x + g_t t + g_{xx} x^2 + g_{tt} t^2 + g_{xt} xt] + \alpha_z z_{it} + \alpha_{\text{TO}}\chi_{it}^{\text{TO}} + \alpha_{\text{LTT}}\chi_{it}^{\text{LTT}} + \lambda\chi_{it}^{\text{TO}}\chi_{it}^{\text{LTT}} + \epsilon_{it}. \qquad (1)$$

Inspection of equation (1) makes clear that our method of identifying the effects of the LTT requires that no other contemporaneous change in policy differentially affects transaction volumes in Toronto and the surrounding municipalities. If, for example, Toronto raised property taxes at the same time as it imposed the LTT and the suburbs did not, then our estimate of λ would reflect the total effect of the property tax change and the LTT. As noted in section 2A, we do not appear to face this problem in our data.

Recalling that χ_{it}^{TO} is a function of x and χ_{it}^{LTT} a function of t, inspection of equation (1) makes clear the importance of controlling for spatial and temporal variation in the latent volume surface. If we estimate equation (1) without adequate controls for $g_x x + g_{xx} x^2$, then these terms are part of the residual and χ^{TO} is obviously endogenous. Similarly, if we estimate equation (1) without adequate controls for $g_t t + g_{tt} t^2$, then these terms are part of the residual and χ^{TO} is endogenous.

Equation (2) addresses these problems using postal code fixed effects, δ_i, and monthly time fixed effects, χ^t, to control non-parametrically for $g_0 + g_x x + g_{xx} x^2$ and $g_t t + g_{tt} t^2$. Since postal

---

[6]Note that x_{it} need not be constant over time. We calculate distance by assigning transactions to individual postal codes. Thus, for a (larger) census tract, the mean distance from a transaction to the border may vary if the distribution of transactions across individual postal codes in the census tract is not stationary.

codes are entirely included in a municipality, a Toronto effect, α_{TO}, and postal code fixed effects cannot be separately identified. Similarly, a post LTT effect, α_{LTT}, and monthly indicators cannot be separately identified. Including postal code and monthly indicators and normalising both α_{TO} and α_{LTT} to zero in equation (1) leaves us with,

$$v_{it} = \delta_i + \sum_t \alpha_t \chi^t + \alpha_z z_{it} + \lambda \chi^{\text{TO}}_{it} \chi^{\text{LTT}}_{it} + \mu_{it}. \tag{2}$$

Equation (2), and a variant which excludes house characteristics, will be our main estimating equations. We here note three potential problems with this baseline specification.

First, comparing equations (1) and (2), we see that the error in equation (2) is $\mu_{it} = g_{xt} x_{it} t_{it} + \epsilon_{it}$. μ. Since $g_{xt} x_{it} t_{it}$ is part of the error term, χ^{TO} and χ^{LTT} are endogenous unless $g_{xt} = 0$. In words, after controlling non-parametrically for all purely temporal and purely spatial variation in g, equation (2) requires that there not be different trends in the growth of transaction volume in Toronto and the suburbs. We postpone a discussion of our approach to this problem.

Second, equation (1) implicitly requires that the parameters, $\alpha^{\text{TO}}, \alpha^{\text{LTT}}$ and λ, be invariant in x and t. This is probably not true. In particular, figure 1 suggests a run-up in sales volume during the fall of 2007 as buyers and sellers raced to complete their transactions ahead of the imposition date for the LTT.[7] If so, then the effect of the LTT varies over time. To prevent this adjustment period from driving our estimate of λ, define the indicator variables

$$\chi^\tau = \begin{cases} 1 & \text{if} \quad x \geq 0 \text{ and } t = \tau \text{ for } \tau \in T = \{-3, -2, -1, 1, 2, 3\} \\ 0 & \text{else}, \end{cases} \tag{3}$$

In words, the χ^τ are six monthly dummy variables for Toronto for the months immediately before and after the imposition of the LTT. If we include these variables in equation (1) then we no longer use these six months to identify the effect of the LTT.

Third, we are concerned that the effect of the LTT is larger farther from the border, in downtown Toronto, where prices are higher. To allow for this possibility define the indicator variable,

$$\chi^{x>k} = \begin{cases} 1 & \text{if} \quad x \geq k \text{ and } t > 0 \\ 0 & \text{else}. \end{cases}$$

That is, $\chi^{x>k}$ is one when the LTT is in effect and the area in question is in Toronto at least distance k from the municipal border. By including $\chi^{x>k}$ in our estimating equations we are able to test whether the effect of the LTT is larger for Toronto parcels further from the border.

Adding these new indicator variables to equation (2) we have,

$$v_{it} = \delta_i + \sum_t \alpha_t \chi^t + \alpha_z z_{it} + \lambda \chi^{\text{TO}}_{it} \chi^{\text{LTT}}_{it} + \sum_{\tau \in T} \alpha_\tau \chi^\tau_{it} + \alpha_{x>k} \chi^{x>k} + \mu_{it}. \tag{4}$$

Finally, we note the possibility that different segments of the property market may react differently to the LTT. To investigate this possibility, we experiment with separate regressions for different price classes of houses.

---

[7]We note that this spike in transactions just visible in bottom panel of figure 2 in the increasing rate of divergence between Toronto and suburban transaction volumes in the months immediately prior to the imposition of the LTT.

Table 2. Effect of the LTT on the number of transactions: characteristics, time trends, and types of houses

Variable	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
LTT	-0.15	-0.15	-0.13	-0.13	-0.33	-0.15	-0.28	-0.09
	$(0.04)^a$	$(0.04)^a$	$(0.04)^a$	$(0.04)^a$	$(0.07)^a$	$(0.05)^a$	$(0.07)^a$	(0.06)
House charact.	N	Y	N	Y	N	N	N	N
Indicators TO +/-3 m.	N	N	Y	Y	Y	Y	Y	Y
Time trends TO	N	N	N	N	One	Two	N	N
Houses	All	All	All	All	All	All	<k$ 400	>k$ 400
Observations	226995	226990	226995	226990	226995	226995	107645	130642
% Effect of LTT	-16	-16	-14	-14	-39	-16	-32	-9

All regressions are Poisson regressions estimated for postal codes within three kilometres of the Toronto border. They all include fixed effects for each month and postal code fixed effects. House charact. is a vector of house characteristics as discussed above. Indicators TO +/- 3 m are six dummy variables for Toronto postal codes during the last three months of 2007 and the first three of 2008. Time trends TO indicates the presence of one or two time trends for Toronto postal codes. Standard errors in parentheses. *a*, *b*, *c*: corresponding coefficient significant at 1, 5, 10%.

6. Results for transaction volumes

Table 2 reports estimation results for transaction volumes. Because of the discrete nature of our data and the large number of zeroes, Poisson regressions are more appropriate than OLS. Thus, all the regressions in table 2 are Poisson regressions. We experiment with alternative estimation methods below. For all regressions in table 2 the dependent variable is the number of sales within three kilometres of the Toronto border in each postal code in each month between January 2006 and August 2008. We experiment below with alternative geographic units. All regressions contain postal code and month indicators.

Column 1 presents our baseline specification using only month indicators and postal code fixed effects as controls. That is, we estimate the following simplified version of equation (2),

$$v_{it} = \delta_i + \sum_t \alpha_t \chi^t + \lambda \chi_{it}^{TO} \chi_{it}^{LTT} + \mu_{it},$$

which does not include house characteristics z_{it}. The explanatory variable of interest is the LTT indicator variable, $\chi_{it}^{TO} \chi_{it}^{LTT}$, which takes the value one for transactions subject to the LTT and zero otherwise. In column 1, the coefficient of the LTT indicator is negative and significantly different from zero at the 1 percent level.

While the Poisson estimator is well suited to predicting our count data, the resulting coefficients are difficult to interpret. To assess the magnitude of our estimated effect, for all postal code months subject to the LTT we calculate the predicted number of sales when the LTT is in effect and the predicted number of sales in the counterfactual case when the LTT is not in effect. Our estimate of the effect of the LTT is the mean, over all postal code months, of the percentage difference between these two quantities. This estimate of the effect of the LTT is given in the last row of the table. The coefficient of -0.15 in the second row of column 1 corresponds to a decline in transaction volume

13

in Toronto of 16%.[8] This is a large effect. The principal pecuniary costs associated with buying a property are brokers' fees, the pre-existing provincial LTT, and the physical cost of moving. If we take these transaction costs to represent between 6 and 7% of the purchase price, the average LTT of 1.1% paid by Toronto buyers corresponds to an increase of about 15% in mobility costs. Hence, the estimate in column 1 suggests that a 1% increase in transactions costs leads to a decline in transactions volume of approximately 1%.

In column 2, we add house characteristics, z_{it}, to the specification of column 1, but do not report their coefficients. Only a minority of house characteristics are significantly correlated with the propensity to sell. For instance, the association between selling and more bathrooms is positive whereas the association between selling and more kitchens is negative. The coefficient of the LTT indicator remains the same as in column 1. Because we worry that house characteristics might be determined simultaneously with the decision to sell, we do not systematically include this set of controls in our regressions.

Column 3 is our preferred specification. In column 3, we include 6 monthly indicator variables for Toronto postal codes from October 2007 to March 2008, the term $\sum_{\tau \in T} \alpha_\tau \chi_{it}^\tau$ in equation (4). Including these controls changes our estimate of the LTT effect from -16 to -14%. In the run-up to the LTT in November and December 2007, Toronto postal codes experience about a 20% increase in transactions volume.[9] Including the six additional monthly indicator variables conditions out this run-up in volume, and also the small decrease that occurred in early 2008. Given that these monthly indicators affect our results (albeit slightly), we retain them from now on. In column 4, we add house characteristics to the specification in column 3. The LTT coefficient remains unchanged.

We now turn our attention to the possibility of spatially differentiated trends. In column 5, we begin in the simplest way by adding a linear time trend for Toronto postal codes. This trend covers our entire study period except for October 2007 to March 2008. The coefficient for the trend variable is large, corresponding to an increase in sales of around 1% per month in Toronto postal codes relative to suburban postal codes. The upward trend in Toronto volumes mechanically implies that the magnitude of the LTT coefficient increases. In fact, including the trend variable changes our estimated LTT effect to -39%. Thus, the results from column 5 suggest that different trends in transactions volume on the Toronto and suburban sides of the border are an important feature of the data and that the Toronto specific trend has confounded our previous estimations.

In column 6, we test the robustness of column 5's result by allowing for a piecewise linear trend for Toronto. That is, we consider a first trend for Toronto from January 2006 to September 2007 and a second trend, also for Toronto, from April to August 2008. The coefficient on the LTT is now much

---

[8]Running the same regression without postal code fixed effects estimates a smaller LTT effect of around -6%. This is unsurprising because this regression fails to account for the fact that postal codes in Toronto experience a greater number of sales, possibly because they are more intensely developed. We do not report these fixed effects for obvious reasons. We do not report months effects either. It is enough to say that they exhibit strong seasonality.

[9]The coefficient for October 2007 is actually negative but insignificant. This is consistent with the notion that the LTT was not anticipated before it was voted at the end of the month. The coefficient for January 2008 is positive and again insignificant. The positive sign for January probably reflects the subtleties of the implementation of the LTT with a rush to sell at the beginning of the month so that the closing date would be before the LTT. This rush slightly dominates the slowdown later in the month. The effects for February and March 2008 are negative, but small and not significant. We experimented with the dates. Considering earlier or later months does not affect our estimate of the LTT.

Figure 5. Time path of post code, month and characteristic corrected transactions across the Toronto border.
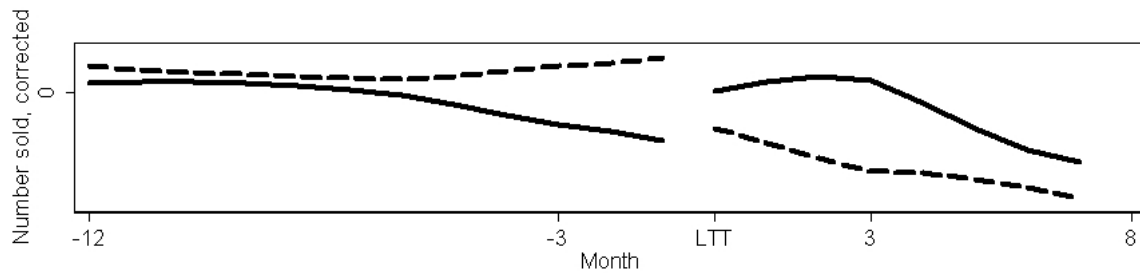
The solid (dashed) line shows the average corrected volume of transactions from January 2006 until August 2008 for a five kilometre wide band just outside (inside) the Toronto border.

closer to our preferred estimate in column 3, with an effect of -16%. Thus, contrary to column 5, column 6 suggests that spatially differentiated trends are not affecting our estimates.

Figure 5 resolves the apparent contradiction. This figure describes monthly transaction volumes in five kilometre bands just inside and just outside the Toronto border during the period around the implementation of the LTT, January 2006 to August 2008 on which our estimations. This is the period and geographic region on which most of our estimations are based. However, rather than report mean transaction volumes, figure 5 reports mean postal code transaction volumes, after these transactions volumes have been corrected for postal code, month and house characteristics. That is, it reports the mean residuals, by month and area, from a regression like one in column 2 of table 2, but excluding the LTT indicator.

Figure 5 shows the presence of spatially differentiated trends in corrected transactions volume. In particular, if we disregard the three months on either side of the implementation date, we see that Toronto volumes are diverging from suburban volumes prior to the LTT and converging afterwards. This suggests that the simple linear trend implemented in column 5 of table 3 is not appropriate and should be ignored in favour of the more complicated specification of column 6. In sum, columns 5 and 6 together with figure 5 indicate the importance of spatially differentiated trends in our data, but also indicate that they do not have much effect on our estimate of λ once we correctly control for them in column 6.

In columns 7 and 8, we allow for the possibility that the LTT may affect different market segments differently. Given the simultaneous determination of property prices and the LTT rate, a complete examination of this problem is beyond the scope of this paper. However, as a beginning, in column 7 we estimate the effect of the LTT for houses with a price below $ 400,000. In column 8, we conduct a similar exercise for houses with a price above $ 400,000. We find that the LTT has a much larger effect on cheaper houses (-32%) than on more expensive houses (-9%, insignificant).[10]

Table 3 reports further results for transaction volumes with a focus on geography and estimation techniques. In column 1, we extend our sample of observations to consider all houses in postal codes within five kilometres of the Toronto border instead of 3. Except for this difference, this

---

[10]We also experimented with splitting our sample along a number of dimensions related to house characteristics but did not find anything conclusive.

Table 3. Effect of the LTT on the number of transactions: distances and estimation techniques

Variable	[1] Poisson	[2] Poisson	[3] Poisson	[4] Poisson	[5] Poisson	[6] Poisson	[7] OLS	[8] Neg. Bin.
LTT	-0.18	-0.12	-0.16	-0.18	-0.07	-0.11	-0.01	-0.13
	$(0.03)^a$	$(0.05)^b$	$(0.04)^a$	$(0.05)^a$	$(0.04)^c$	$(0.04)^a$	$(0.003)^a$	$(0.04)^a$
Distance threshold	5km	2km	5km	5km	3km	3km	3km	3km
Distance LTT trends	N	N	Step	Y	N	N	N	N
Geographical unit	Postal code	Postal code	Postal code	Postal code	Census tract	Distribution Area	Postal code	Postal code
Observations	385188	146338	385188	385188	6705	36965	297174	226995
% Effect of LTT	-19	-13	-18	-20	-8	-12	-20	-13

All regressions include a month indicators, month indicators for Toronto between October 2007 and March 2008, and fixed effects for the appropriate geographic unit. Standard errors in parentheses. *a*, *b*, *c*: corresponding coefficient significant at 1, 5, 10%.

specification mirrors that of column 3 in table 2. The estimated effect is slightly larger at -19% instead of -14%. These two coefficients are nonetheless statistically indistinguishable. In column 2, we replicate again the same regression but consider only a two kilometre band on both sides of the Toronto border. The results are very similar to those obtained with a three kilometre band with an LTT effect of -13%.

In column 3, we return to a five kilometre band but add a coefficient for Toronto postal codes located between two and a half and five kilometres away from the Toronto border. This corresponds to the term $\alpha_{x>k}\chi^{x>k}$ for $k = 2.5$ in equation (4). The (unreported) coefficient for the additional effect of the LTT beyond two and a half kilometres is negative but insignificant. Rather than a two step effect, in column 4 we estimate a spatial trend by interacting the implementation of the LTT with distance to the border. The (unreported) coefficient for this spatial trend is again insignificant. Overall, the results of column 1-4 are suggestive of a slightly stronger effect of the LTT as one moves away from the border towards the centre of Toronto but this effect is not statistically significant.[11]

In column 5, we repeat our preferred regression but use census tracts as our unit of observation instead of postal codes. Census tracts are bigger than postal codes by a factor of more than 30. The effect of the LTT is smaller than with postal codes, -8% instead of -14%. In column 6, we use distribution areas which are smaller than census tracts but still much larger than postal codes. The estimated effect of the LTT on the volume of transaction is -12%. Using coarse geographical units neglects important small scale variation.

In column 7, we repeat our preferred estimation but this time use OLS with standard errors clustered by postal codes rather than a Poisson regression. We see that, at -20%, the OLS estimate of the LTT effect is close to our preferred estimate.[12] In column 8, we repeat the same exercise

---

[11]We experimented with various specifications but we were unable to evidence any robust pattern in the spatial effects of the LTT either within or outside Toronto.

[12]We also used this specification to explore whether accounting for the spatial autocorrelation of errors affected our standard errors. The effects are small. For instance, running the same OLS specification but clustering the error term at the level of three-digit postal codes increases the standard error on the LTT coefficient by about 10%. It remains significant at 1%.

Variable	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
LTT	-0.13	-0.13	-0.11	-0.13	-0.08	-0.15	-0.13	-0.13
	$(0.04)^a$	$(0.04)^a$	$(0.04)^a$	$(0.05)^b$	$(0.04)^b$	$(0.04)^a$	$(0.05)^a$	$(0.06)^b$
House charact.	N	Y	Y	Y	Y	Y	Y	Y
First quarter	06Q1	06Q1	06Q1	06Q1	06Q1	07Q1	07Q1	07Q3
Last quarter	08Q3	08Q3	08Q3	08Q3	08Q3	08Q3	08Q3	08Q2
07Q4 and 08Q1 excluded	N	N	N	N	N	N	Y	Y
Indicators 07Q4 and 08Q1 for TO	N	N	Y	Y	N	N	-	-
Time trends TO	N	N	N	Two	N	N	N	N
Time units	Q	Q	Q	Q	B&A	Q	Q	B&A
Observations	78813	78811	78811	78813	15078	6286	6286	2961
% Effect of LTT	-14	-14	-12	-14	-9	-17	-14	-14

All regressions are Poisson regressions estimated for postal codes within three kilometres of the Toronto border. They all include postal code fixed effects. House charact. is a vector of house characteristics as discussed above. First quarter and Last quarter indicate the time span considered. 07Q4 and 08Q1 excluded indicates the deletion of the last quarter of 2007 and the first quarter of 2008. Indicators 07Q4 and 08Q1 for TO are two dummy variables for Toronto postal codes during the last quarter of 2007 and the first quarter of 2008. Time trends TO indicates the presence of time trends for Toronto postal codes. Finally the time units are Q for quarters and B & A when only two periods are considered, before and after the LTT. Standard errors in parentheses. a, b, c: corresponding coefficient significant at 1, 5, 10%.

using a negative binomial rather than a Poisson regression. This yields the same estimate of the LTT coefficient as the corresponding Poisson regression.

Following the argument made by Bertrand, Duflo, and Mullainathan (2004), serially correlated transactions at the postal code level could bias downwards the standard errors for the LTT coefficient in tables 2 and 3. To explore this issue, table 4 reports a series of results using transactions data aggregated over longer periods of time, quarters or more instead of months. Although coarser time controls may have a small effect on the point estimates, they will reduce any problem caused by serial correlation.

Column 1 of table 4 reproduces the specification of column 1 of table 2 but uses a quarterly aggregation instead of a monthly aggregation. While the point estimates differ only a little (-0.146 vs. -0.131), the standard errors are virtually identical (0.0358 vs. 0.0356). Column 2 adds house characteristics to the regression and thus mirrors column 2 of table 2 with quarterly data. The results are the same: a slightly lower point estimate and identical standard errors. Column 3 adds indicator variables for Toronto during the last quarter of 2007 (to condition out the pre-LTT spike of transactions) and first quarter of 2008 (to condition out the post-LTT reflection of the previous spike). This corresponds to column 4 of table 2 with quarterly data and leads to yet again the same result. Column 4 enriches the specification further by adding two time trends, before and after the LTT. This mirrors column 6 of table 2. Similar standard errors are again indicative of a lack of serial correlation.

In column 5 of table 4, time aggregation is taken to the extreme and we consider only two time periods, before and after the LTT. Because it includes house characteristics, this specification

mirrors column 2 with two time periods. The standard errors remain identical. In column 6, we return to the specification of column 2 but ignore 2006 to consider only transactions from 2007 and 2008. This alternative strategy also reduces the number of quarters before the implementation of the LTT and should allow us to reduce any problem caused by serial correlation. Relative to column 2, the standard errors increase by less than 10% while the number of observation is about halved. Column 7 duplicates the same regression again but excludes observations from the last quarter of 2007 and the first quarter of 2008. Despite another reduction in the number of observations, standard errors increase only slightly. In column 8, we estimate an even more extreme specification that only considers transactions from the third quarter of 2007 and the second quarter of 2008. The point estimates are still the same and the loss of precision remains small. Despite using less than 3,000 observation instead of close to 227,000 in column 2 of table 2, the standard error for the LTT coefficient in column 8 of table 4 is less than double.

From tables 2 and 4, we conclude that the effect of the LTT on transaction volumes is between -10 and -20%, with our preferred estimate being -14%. This conclusion is robust to our choice of geography, controls for house characteristics, time aggregation, and choice of estimation method. Controlling for the peak of transactions in Toronto prior to the implementation of the LTT makes a small difference to our estimates. Controlling correctly for spatially differentiated trends does not affect our results. Trying to minimise the effect of serial correlation in transactions causes virtually no change to the standard errors. Finally, one might worry that this decline in transactions is only temporary and caused by the reluctance of sellers to incur a nominal loss after the implementation of the LTT (Genesove and Mayer, 2009). This concern is misplaced because prices were still increasing during the first half of 2008 according to the Terranet National Bank House Price Index.[13] Besides, as we show next, the price reduction caused by the LTT is small and thus unlikely to have such effects.

7. Results for prices

For our sample, the average amount of LTT paid is about 1.1 percent of the sale price. Consistent with this, inspection of figure 2 suggests that the effect of the LTT on prices is small. In this section, we refine the estimate of the price effect of the LTT suggested by figure 2, with the caveat that measuring so small an effect accurately is difficult, even with our high-quality data.

Our strategy for estimating the effect of the LTT on the price of single-family homes is much the same as the one we employ to estimate the effect of the LTT on the number of sales. For sales within a three kilometre band around the Toronto border between January 2006 and August 2008 we estimate,

$$\log p_{jt} = \delta_{i(j)} + \sum_t \alpha_t \chi^t + \alpha_z z_{jt} + \lambda \chi_{jt}^{\text{TO}} \chi_{jt}^{\text{LTT}} + \sum_{\tau \in T} \alpha_\tau \chi_{jt}^{\tau} + \mu_{jt}. \tag{5}$$

Although this equation resembles our estimating equation (4) for transaction volumes, several differences are worth highlighting. First, we use the log of the sale price for individual transactions as the dependent variable. The unit of observation is a particular transaction, j, sold in a given

___
[13]See http://www.housepriceindex.ca/Default.aspx.

Table 5. Effect of the LTT on prices

Variable	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
LTT (% change)	-1.31	-0.76	-0.88	-1.43	-0.81	-1.19	-1.48	-0.81
	$(0.62)^b$	$(0.42)^c$	$(0.49)^c$	$(0.79)^c$	$(0.42)^c$	$(0.47)^b$	$(0.60)^b$	(0.64)
House characteristics	N	Y	Y	Y	Y	Y	Y	Y
Indicators TO +/-3 m.	N	N	Y	Y	Y	Y	Y	Y
Time trends TO	N	N	N	Y	N	N	N	N
Distance threshold	3km	3km	3km	3km	5km	5km	5km	2km
Distance LTT trends	N	N	N	N	N	Step	Y	N
Observations	19785	19783	19783	19783	32994	32994	32994	12773

All regressions are OLS regressions estimated with postal codes fixed effects and indicator variables for each month from January 2006 to August 2008. House characteristics is a vector of house characteristics as discussed above. Indicators TO +/- 3 m are six dummy variables for Toronto postal codes during the last three months of 2007 and the first three of 2008. Time trends TO indicates the presence of one or two time trends for Toronto postal codes. Distance threshold is the maximum distance to the Toronto border for a postal code to be included. Distance LTT trend denotes the inclusion of an interaction term between LTT and distance to the Toronto border for Toronto postal codes. Standard errors clustered by postal code in parentheses. *a*, *b*, *c*: corresponding coefficient significant at 1, 5, 10%.

month, t. Since the dependent variable, log transaction price, is continuous, we estimate equation (5) with OLS (rather than with Poisson, as we used for equation 4). Price is a function of its location captured by its postal code, $\delta_{i(j)}$.[14] As in our earlier volume regressions, equation (4), we include monthly fixed effects, $\sum_t \alpha_t \chi^t$, in equation (5). House characteristics, $\alpha_z z_{jt}$ are also included and play a more important role. We expect the Toronto market during the months immediately before and after the implementation of the LTT to experience some turmoil which might affect prices. Therefore, we control for month effects in Toronto from October 2007 to March 2008, as we do in equation (4), with the term $\sum_{\tau \in T} \alpha_\tau \chi^\tau_{jt}$. Finally, as in equation (4), the effect of the LTT is captured by the term in $\lambda \chi^{TO}_{jt} \chi^{LTT}_{jt}$.

Column 1 presents our baseline specification using only month indicators and postal code fixed effects as controls. This estimation mirrors the first regression of table 2 and estimates the following simplified version of equation (5): $\log p_{jt} = \delta_{i(j)} + \sum_t \alpha_t \chi^t + \lambda \chi^{TO}_{jt} \chi^{LTT}_{jt} + \mu_{jt}$. In this regression the coefficient of the LTT indicator is negative and significantly different from zero at the 5 percent level. As expected the effect of the LTT on property prices is modest at -1.3%. This effect is slightly above the average LTT rate but statistically indistinguishable from it.

In column 2, we add house characteristics, the term in $\alpha_z z_{it}$ in equation (5), to mirror the second column of table 2. We use the same broad set of property characteristics as previously. The coefficient of the LTT indicator is now slightly smaller at -0.8%. However it does not differ statistically from the tax rate nor from the estimate in column 1. We do not report here the coefficients for house characteristics. Most of them are highly significant and their signs and magnitudes are unsurprising. Because of their importance in explaining houses prices we retain house characteristics in subsequent price regressions.

---

[14]Because we cannot track properties over time (only their postal code), we cannot use a repeated-sale index methodology (Case and Shiller, 1987).

In column 3, as in column 3 of table 2, we add 6 monthly indicators variables for October 2007 to March 2008. These correspond to the term in $\sum_{\tau \in T} \alpha_\tau \chi_{it}^\tau$ in equation (5) and lead to a marginally larger effect of the LTT at -0.9%. Interestingly, we find a small significant jump in prices of around 3% for Toronto in December 2007.[15] This suggests that the heating up of the market in the run-up to the LTT led some buyers to pay more than the amount of the LTT in their rush to avoid it. This regression is our preferred estimation. Although, they do not make any difference to our results, we retain these six monthly indicators to maintain symmetry with our transactions volume regressions.

In column 4, we add a time trend for all Toronto transactions. For reasons similar to those in our previous set of regressions, adding a time trend increases the LTT coefficient, albeit not as dramatically as for transaction volumes. We experimented with time effects more extensively and, as with our volumes regressions, we conclude that our results are somewhat sensitive to the inclusion of time trends and the exact specification for them.

In column 5, we replicate our preferred regression of column 3 but we extend our sample of observations to consider all properties sold within five kilometres of the Toronto border instead of 3. This specification mirrors the first column of table 3 for transaction volumes. At -0.8% our estimated coefficient is slightly smaller but statistically indistinguishable from our preferred estimate of -0.9%. In column 6, we explore the possibility of a spatially differentiated effect of the LTT by adding a coefficient for Toronto properties between two and a half and five kilometres away from the Toronto border. This (unreported) coefficient is positive but small and insignificant. In column 7, we consider instead a spatial trend by interacting the implementation of the LTT with distance to the border. The (unreported) coefficient for this spatial trend is again tiny and insignificant. Finally, in column 8, we replicate our preferred specification of column 3 but restrict our sample to all properties sold within two kilometres of the Toronto border instead of 3. The coefficient is very close to that in our preferred specification (and that using a five kilometre threshold in column 5). With fewer transactions, the standard errors are now slightly larger, making the coefficient statistically insignificant.

In all of our price regressions we examine the effect on price of a dichotomous variable indicating whether or not the transaction was subject to the LTT. We also experiment with other measures of the LTT, in particular with an interaction between the LTT indicator and the tax rate to which the transaction is was subject. This experiment did not lead to statistically significant estimates of the effect of the LTT on price. Partitioning our sample into more and less expensive houses and repeating our estimations on the resulting sub-samples provides weak evidence for the hypothesis the price effect is stronger for properties that pay a higher rate.

Overall, the results of our price regressions suggest that the LTT led to a decrease in Toronto property prices of about the same magnitude as the tax.

---

[15]The coefficients for the other months are not significant.

8. A simple model of residential location with property taxes

We now consider a simple model of residential location. Our goals are to understand the implications of the LTT for location decisions, to understand the welfare implications the LTT, and to provide the foundations for our welfare calculation.

Consider an economy with two possible locations and two agents. The locations are 'Toronto', denoted T, and a 'suburb', denoted S.[16] The two agents are an incumbent, I, who owns a house in Toronto but who may move to the suburb, and an entrant, E, who will move to Toronto or the suburb. Since there are only two possible locations, one of the two agents must locate in Toronto, and the other in the suburb. Let $j \in \{E, I\}$ index players.

Initially both entrant and incumbent draw an employment or educational opportunity in Toronto. This opportunity has monetary value u_j^T if agent j locates in Toronto, and zero otherwise. We suppose that each player draws u_j^T from a uniform distribution with support $[\underline{u}, \bar{u}]$, and that these draws are common knowledge.

Next, the incumbent chooses the price p^T at which he offers to sell his house to the entrant. The entrant then chooses a location $l_E \in \{S, T\}$. If the entrant rejects the incumbent's offer, he chooses $l_E = S$ and moves to the suburbs, while the incumbent remains in Toronto. If the entrant accepts, he chooses $l_E = T$, he pays p^T to the incumbent and occupies the incumbent's house, while the incumbent moves to the suburbs.

All suburban housing provides both entrant and incumbent with the same monetary payoff, $u^S < \underline{u}$, and can be purchased for price p^S. To ease exposition suppose that the incumbent does not pay for housing if he remains in Toronto. Each agent acts to maximise the difference between the value associated with his location, and the price paid for that location.

Suppose that the market price at each location is such that the seller captures all of the surplus, as in 'take it or leave it' bargaining, and that the entrant moves to Toronto when he is indifferent between locations. Since the seller captures all surplus, we must have $u^S = p^S$. That is, whether the entrant or incumbent ends up in the suburbs, the price of the suburban house is equal to the value of occupying it. Our assumption that the incumbent capture all surplus requires that the entrant be indifferent between the suburbs and Toronto whenever he moves to Toronto. Therefore, if the incumbent sells to the entrant then the sale price must satisfy $u_E^T - p^T = u^S - p^S$. This requires that $p^T = u_E^T$.

The incumbent will sell to the entrant if and only if his payoff from doing so exceeds his payoff from staying in Toronto. That is, if and only if $p^T + u^S - p^S > u_I^T$. It follows that the incumbent moves to the suburbs (and the entrant to Toronto) if and only if $u_E^T > u_I^T$ and that if this transaction occurs then $p^T = u_E^T$.

We now suppose that a land transfer tax is imposed in Toronto. A land transfer tax is only paid when an individual purchases a house. It is not paid if the house does not change hands. As a stylised description of this tax, suppose that the amount τ is collected from the entrant if he purchases the incumbent's house and moves to Toronto, but not otherwise.

---

[16]This model could easily be extended to consider moves within Toronto as well. However, since our data do not track buyer and seller locations, we do not pursue this investigation.

Let τ denote a land transfer tax. Just as without the LTT, the incumbent captures all of the surplus. Thus the equilibrium price must solve $u_E^T - p^T - \tau = u^S - p^S$ so that $p^T = u_E^T - \tau$. The incumbent chooses to sell if and only if doing so makes him better off. That is, if and only if $p^T + u^S - p^S > u_I^T$. This condition holds if and only if $u_E^T - \tau > u_I^T$. Thus we see that the land transfer tax prevents beneficial trades that would occur in its absence.

It is also of interest to consider equilibrium behaviour in the presence of an ordinary property tax. Suppose that property in Toronto is subject to an ordinary property tax, $\tau_o \in R_+$, that must be paid by the incumbent regardless of whether the property changes hands.[17]

In this case, if the incumbent captures all surplus from any real estate transaction, we have $u_E^T - p^T = u^S - p^S$ so that $p^T = u_E^T$. The incumbent is willing to sell at this price if and only if $p^T - \tau + u^S - p^S > u_I^T - \tau$. This condition holds if and only if $u_E^T > u_I^T$. Therefore, unlike the LTT, exactly the same trades occur under an ordinary property tax as occur without any property tax at all.

We now turn to welfare analysis. Define welfare, W, to be the sum of the utility levels of the entrant and incumbent. Thus, $W = u_E^T + u^S$ if the entrant moves to Toronto and incumbent to the suburbs, and $W = u_I^T + u^S$ otherwise. It follows that welfare is maximised if and only if the agent who draws the best opportunity lives in Toronto. In an untaxed equilibrium and when there is an ordinary property tax the agent who draws the best opportunity always locates in Toronto. It follows that such equilibria are socially optimal. In the presence of a land transfer tax the market allocation is not socially optimal. When property in Toronto is subject to a land transfer tax, an entrant who draws an opportunity $u_E^T \in (u_I^T, u_I^T + \tau)$ does not move to Toronto. The resulting welfare loss is $u_E^T - u_I^T \in [0, \tau]$.

To use this intuition to estimate the loss of welfare caused by Toronto's LTT, we must deal with two complications. First, when the Toronto LTT was imposed an almost identical provincial LTT was already in place, while the welfare loss calculated above results from increasing the LTT from zero to τ. To correctly reflect this situation, let τ_O denote the old or pre-existing provincial LTT. A marginal transaction is one for which u_E^T lies in the interval $[u_I^T + \tau_O,\ u_I^T + \tau_O + \tau]$. Such a transaction occurs before the new Toronto LTT, but not after, and entails a welfare loss of $u_E^T - u_I^T \in [\tau_O, \tau_O + \tau]$. That is, the magnitude of the welfare loss associated with each marginal foregone transaction is at least equal to the old tax, and is no greater than the sum of the two taxes.

The second complication we must overcome is that, while we are able to estimate the frequency of foregone transactions resulting from Toronto's LTT, we do not observe utility levels associated with foregone transactions. To overcome this problem, we use our assumption that u_E^T and u_I^T are independent draws from a uniform distribution on $[\underline{u}, \overline{u}]$. In this case, the welfare loss of each foregone transaction is the expected value of $u_E^T - u_I^T$, conditional on the transaction not occurring only because of the new LTT.

Figure 6 illustrates the space of possible draws of u_E^T and u_I^T. The shaded triangle describes pairs (u_E^T, u_I^T) for which transactions occur before and after the imposition of the Toronto LTT. The unshaded region describes pairs (u_E^T, u_I^T) for which transactions do not occur before or after the

---

[17]Our conclusions are unchanged if we require the incumbent to pay the property tax if he stays in Toronto and the Entrant to pay if he locates in Toronto.

Figure 6. Equilibrium location decisions before and after an increase in the LTT from τ_O to $\tau_O + \tau$.
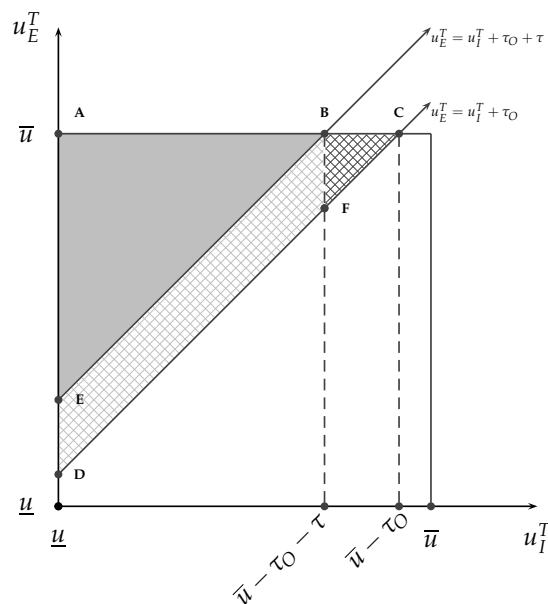
new tax. The two hatched areas describe transactions foregone because of the Toronto LTT. This is the region of the utility space where transactions occur before the new tax, but not after.

Since we have assumed that pairs (u_E^T, u_I^T) are uniformly distributed, the probability of drawing a pair (u_E^T, u_I^T) that results in a transaction before the land transfer tax is

$$\frac{\text{Area of } \triangle ACD}{[\overline{u} - \underline{u}]^2}.$$

Similarly, the probability of drawing a pair (u_E^T, u_I^T) that results in a transaction after the land transfer tax is

$$\frac{\text{Area of } \triangle ABE}{[\overline{u} - \underline{u}]^2}.$$

We estimate that the LTT causes about a 14% decline in the number of transactions. In the context of figure 6, this is the conditional probability of a transaction not occurring with the LTT given that it would have occurred without. Under our assumption of uniformity, this conditional probability is the ratio of the area of $\lozenge BCDE$ to $\triangle ACD$. Thus we have

$$\frac{\text{Area of } \lozenge BCDE}{\text{Area of } \triangle ACD} = 0.14$$

Since the two triangles $\triangle ABE$ and $\triangle ACD$ are right-angled and isosceles, and since we know that the distance $|\overline{BC}| = \tau$, it follows that $\overline{u} - \tau_O \approx 13.77\tau$.[18]

---

[18] This ratio may be written

$$\frac{\frac{1}{2}(\overline{u} - \tau)^2 - \frac{1}{2}(\overline{u} - \tau - \tau_0)^2}{\frac{1}{2}(\overline{u} - \tau)^2} = 0.14.$$

A little algebra gives the desired result.

It is now straightforward to calculate the two conditional probabilities,

$$\text{Prob}\left[(u_E^T, u_I^T) \in \Diamond BFDE | (u_E^T, u_I^T) \in \Diamond BCDE\right] \approx 0.962$$

and,

$$\text{Prob}\left[(u_E^T, u_I^T) \in \triangle BCF | (u_E^T, u_I^T) \in \Diamond BCDE\right] \approx 0.038.$$

These probabilities are, respectively, the probability that the utility levels for a foregone transaction lie in the region $\Diamond BFDE$ and $\triangle ACD$.

It remains to calculate the expected value of $u_E^T - u_I^T$ conditional on the associated transaction not occurring because of the Toronto LTT. First note that, conditional on such a foregone transaction falling in $\Diamond BCDE$, $E(u_E^T - u_I^T) = \tau_O + \tau/2$. [19] Similarly, conditional on a foregone transaction falling in $\triangle BCF$, it is not difficult to show that $E(u_E^T - u_I^T) = \tau_O + \tau/4$.

It follows immediately that the expected welfare loss of a transaction foregone because of the Toronto LTT is

$$\begin{aligned} \Delta W &= \tau_O + 0.038 \times \tau/4 + 0.962 \times \tau/2 \\ &= \tau_O + 0.491\tau. \end{aligned}$$

In our sample the average Toronto LTT exaction is about 1.1%. This is almost identical to the provincial LTT exaction. The value of an average house in Toronto is about $ 400,000. Thus, we estimate that $\tau_O = \tau = \$4,400$. Substituting into the expression above, we estimate that the welfare loss from each transaction foregone because of the Toronto LTT at $ 6,559.

The average annual number of transactions recorded in our data for 2004-2006 is about 21,200 for Toronto. Multiplying 21,200 by the estimated 14 percent effect of the LTT gives us a decrease of about 3,000 in the number of single-family homes sold in an average year in Toronto. Multiplying this number by the welfare loss associated with each transaction gives a total welfare loss of 19.5 million dollars.

Three comments are in order. First, the MLS reflects many, but not all single family home sales. While no definitive estimates of the share of transactions covered by the MLS are available, our best guess is that these data record about 85 percent of total single-family-home sales. Our estimate of the cost of the LTT is low because it does not reflect all single family home sales, and because it does not reflect condominium sales.

Second, our estimate is made under the seemingly strong assumption that utility levels are drawn from a uniform distribution. In fact, while this assumption is dramatically easier to work with than others we considered, the final estimate of the welfare loss does not appear to be especially sensitive to this assumption. Using different distributional assumptions principally

---

[19]From figure 6, we have

$$\begin{aligned} E\left(u_E^T - u_I^T | (u_E^T, u_I^T) \in \Diamond EBFD\right) &= \frac{1}{\tau(\overline{u} - \tau - \tau_0 - \underline{u})} \int_{\underline{u}}^{\overline{u} - \tau_0 - \tau} \int_{x + \tau_0}^{x + \tau_0 + \tau} (y - x) dy dx \\ &= \tau_0 + \frac{\tau_2}{2}. \end{aligned}$$

serves to affect the expectation of $u_E^T - u_I^T$ conditional on a foregone transaction. Under current assumptions, this expectation evaluates to $\tau_O + 0.491\tau$. Under other distributional assumptions, this expectation is bounded by τ_O below and $\tau_O + \tau$ above. The lower bound is obtained if every foregone transaction is exactly as costly as the worst transaction that could occur without the Toronto LTT, and the upper bound is obtained if every foregone transaction is exactly as costly as the best transaction that is foregone because of the LTT.

Third, it is important to note that an LTT is a tax on property just as is an ordinary property tax. Therefore, to the extent that an ordinary property tax distorts a consumer's choice of land and housing capital, so does an LTT. Thus, the LTT truly appears to be dominated by an ordinary property tax. All of the distortions caused by an ordinary property tax are also caused by an LTT, but the disincentive for moving is unique to the LTT.

9. Conclusions

Using data that describe 139,266 single family home sales, we exploit a natural experiment to estimate the effects of a new LTT on Toronto's real estate market. Consistent with our intuition that an LTT operates as a tax on moving, we find that Toronto's tax decreases the volume of real estate transaction by about 14%. That is, Toronto's LTT leads to a 14% decrease in residential mobility. We also find that the LTT is approximately fully capitalised into land prices.

The theoretical framework we develop makes our intuition about the effects of the LTT precise and offers a way to calculate the welfare loss associated with Toronto's LTT. This welfare loss, effectively the cost of foregone mobility, is substantial, about 1$ for every 8$ in tax revenue raised or about 19 million dollars per year for Toronto. In contrast, an ordinary property tax has no impact on mobility, and therefore avoids the excess burden associated with the LTT.

Land transfer taxes are pervasive. In particular , they are used by national or sub-national governments in Australia, Canada, France, the United Kingdom, and the United States. This analysis has immediate policy implications for these jurisdictions. In particular, that welfare would be improved by shifting the burden of public finance from land transfer taxes to ordinary property taxes.

References

38th Legislature of Ontario (2006) *Bill 53, An Act to revise the City of Toronto Act, 1997*. Toronto, Ontario: Legislative Assembly of Ontario.

Angrist, J.D., Lavy, V. (1999) Using Maimonides' rule to estimate the effect of class size on scholastic achievement. *Quarterly Journal of Economics*, 114: 533–575.

Bayer, P., Ferreira, F., McMillan, R. (2007) A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115: 588–637.

Benjamin, J.D., Coulson, N.E., Yang, S.X. (1993) Real estate transfer taxes and property values: The Philadelphia story. *Journal of Real Estate Finance and Economics*, 7: 151–157.

Bertrand, M., Duflo, E., Mullainathan, S. (2004) How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119: 249–275.

Black, S.E. (1999) Do better schools matter? Parental valuation of elementary education. *Quarterly Journal of Economics*, 114: 577–599.

Case, K.E., Shiller R.J. (1987) Prices of single family homes since 1970: New indexes for four cities. *New England Economic Review*, 1987: 45–56.

City of Toronto. 2008. *Statistical Information, 2007 Financial Report*. Toronto, Ontario: City of Toronto.

Duranton, G., Gobillon, L., Overman H.G. (2011) Assessing the effects of local taxation using microgeographic data. *Economic Journal*, forthcoming.

Federation of Tax Administrators (2006) State real estate transfer taxes. *FTA Bulletin B-03/06* .

Ferreira, F., Gyourko, J. (2009) Do political parties matter? Evidence from US cities. *Quarterly Journal of Economics*, 124: 399–422.

Genesove, D., Mayer, C. (2009) Loss aversion and seller behavior: Evidence from the housing market. *Quarterly Journal of Economics*, 116: 1233–1260.

Gibbons, S., Machin, S. (2003) Valuing english primary schools. *Journal of Urban Economics*, 53: 197–219.

Glaeser, E.L., Luttmer E.F.P. (1997) The misallocation of housing under rent control. Working Paper 6220, National Bureau of Economic Research.

Glaeser, Glaeser, E.L., Luttmer E.F.P. (2003) The misallocation of housing under rent control. *American Economic Review*, 93: 1027–1046.

Hahn, J., Todd, P., Van der Klaauw W. (2001) Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69: 201–209.

Holmes, T.J. (1998) The effect of state policies on the location of manufacturing: Evidence from state borders. *Journal of Political Economy*, 106: 667–705.

Imbens, G.W., Lemieux, T. (2008) Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, 142: 615–635.

Ioannides, Y.M., Kan, K. (1996) Structural estimation of residential mobility and housing tenure choice. *Journal of Regional Science*, 36: 365–364.

Lemieux, T., Milligan, K. (2008) Incentive effects of social assistance: A regression discontinuity approach. *Journal of Econometrics*, 142: 807–828.

Nowlan, D.M. (2007) *Economic Implications of the Proposed City of Toronto Land Transfer Tax, Attachment 1 to 'New Taxation Measures Supplemental Report - City of Toronto Act 2006, Toronto City Council agenda' (July 16th, 2007)*. Toronto: City of Toronto.

Van Ommeren, J., Van Leuvensteijn, M. (2005) New evidence of the effect of transaction costs on residential mobility. *Journal of Regional Science*, 45: 681–702.

10. Appendix: Construction of figures 1 and 2

Top panel figure 1:

This figure is based on 67,502 transaction that occur during a pre-LTT period from February to August of 2007 and a post-LTT period during the corresponding months of 2008. After consolidating a small number of postal codes with centroids exactly the same distance from the border, these transactions give us 32,818 postal codes for which we can count transactions before and after the LTT. For each postal code we calculate the raw change in the number of transactions (since many postal codes do not record a transaction before the LTT, calculating the percentage change in transaction entails dropping postal codes). This results in a set of 32,818 postal code changes in transaction volumes. We drop postal codes further than ten kilometres from the border, leaving us with 19,466.

The top panel of figure 1 is a lowess plot of the way these changes in transactions vary as the distance to the Toronto border varies. Dotted lines are constructed by bootstrapping this lowess estimation and show 95% and 5% pointwise confidence bounds.

Bottom panel figure 1:

This figure is based on 45,856 transaction that occur within five kilometres of the Toronto border between January 2005 and August 2008. For each month during this time period, we calculate the number of transactions on each side of the border. From this, we calculate

$$\frac{\text{Toronto monthly transactions} - \text{Suburban monthly transactions}}{\text{Toronto monthly transactions}} \times 100,$$

which yields a sequence describing the percentage change in price across the municipal border for 44 months.

The bottom panel of figure 1 plots these 44 points. Since we are able to display the universe of data (rather than an estimated mean), we do not provide confidence bounds. Thus, for December of 2007 we see that about 45% more transactions occurred in the five kilometre band inside of the border than in the five kilometre band outside.

Top panel figure 2:

This figure is based on 67,502 transaction that occur during a pre-LTT period from February to August of 2007 and a post-LTT period during the corresponding months of 2008. After consolidating a small number of postal codes with centroids exactly the same distance from the border, these transactions give us 32,818 postal codes for which can observe transactions before and after the LTT. For each postal code we calculate the percentage change in transaction price,

$$\frac{\text{Post-LTT price - Pre-LTT price}}{\text{Pre-LTT price}} \times 100.$$

Since this quantity is undefined if there are zero transactions in a postal code pre-LTT, we are left with a set of 5,464 postal codes for which we observe percentage change in price that occurs around the inception of the LTT.

The top panel of figure 2 is a lowess plot of the way these changes in prices are distributed across space, as the distance to the Toronto border varies. Dotted lines are constructed by bootstrapping this lowess estimation and show 95% and 5% pointwise confidence bounds.

Bottom panel figure 2:

This figure is based on 45,856 transaction that occur within five kilometres of the Toronto border between January 2005 and August 2008. For each month during this time period, we calculate the monthly average transaction price for each side of the border. From this we calculate

$$\frac{\text{Toronto monthly mean transaction price} - \text{Suburban monthly mean transaction price}}{\text{Toronto monthly mean transaction price}} \times 100,$$

which yields a sequence describing the percentage change in price across the municipal border for 44 months.

The bottom panel of figure 1 plots these 44 points. Since we are able to display the universe of data (rather than an estimated mean), we do not provide confidence bounds. Thus, for December of 2007 we see that the average house sold for about 15% more in Toronto than in the suburbs.