Mansion Tax: The Effect of Transfer Taxes on the Residential Real Estate Market [1]

Wojciech Kopczuk and David J. Munroe
Columbia University

January 2, 2014

[1]Wojciech Kopczuk: Department of Economics and School of International and Public Affairs, Columbia University, NBER and CEPR; email address: wojciech.kopczuk@columbia.edu. David Munroe: Department of Economics, Columbia University, email address: djm2166@columbia.edu. We are grateful to Chris Mayer, the Paul Milstein Center for Real Estate and the Real Estate Board of New York for the data. Gilles Duranton, Henrik Kleven, Randy Reback, Maya Rossin-Slater, Anindya Sen, and seminar participants at Columbia, Syracuse, Federal Reserve Board of Governors, Wharton, New York Fed, Toronto, Northwestern and the CEA provided many helpful suggestions.

Abstract

Houses and apartments sold in New York and New Jersey at prices above $1 million are subject to the so-called 1% "mansion tax" imposed on the full value of the transaction. This policy generates a discontinuity (a "notch") in the overall tax liability. We rely on this and other discontinuities to analyze implications of transfer taxes in the real estate market. Using administrative records of property sales, we find robust evidence of substantial bunching and show that the incidence of this tax for transactions local to the discontinuity falls on sellers, may exceed the value of the tax, and is not explained by tax evasion (although supply-side quality adjustments may play a role). Above the notch, the volume of missing transactions exceeds those bunching below the notch. Interpreting our results in the context of an equilibrium bargaining model, we find that the market unravels in the neighborhood of the notch: its presence provides strong incentive for buyers and sellers near the threshold not to transact. Finally, we show that the presence of the tax affects how the market operates *away* from the threshold—taxation increases price reductions during the search process and in the bargaining stage and weakens the relationship between listing and sale prices. We interpret these results as demonstrating that taxation affects the ultimate allocation in this search market.

Short Abstract (100 words)

Using discontinuities in housing transaction taxes in New York and New Jersey we find robust price bunching. Incidence for transactions local to the notch falls on sellers, with no evidence of evasion. The volume of missing transactions above the notch exceeds those bunching (beyond what would be expected from the usual extensive-margin response), indicating strong incentives for buyers and sellers not to transact (market unravels). Away from the threshold, we find increased discounts and weaker relationship between listing and sale prices. Equilibrium bargaining framework highlights that taxation affects the ultimate allocation in this search market.

1 Introduction

Purchasing real estate is a time-consuming and complicated process with large financial stakes and potentially important frictions. Beyond the price, a typical transaction involves many associated costs, including broker's fees, inspection costs, legal fees, title insurance, mortgage application and insurance fees, and moving costs. In this paper, we rely on a particular type of cost—transfer taxes that are imposed on the value of real estate transactions—to understand how frictions affect functioning of this market.

Our objective is fourfold. Real estate transfer taxation is very common and given the importance of this market it is of interest to understand the empirical implications of such taxes. Second, we take advantage of variation in tax incentives and data on both transactions and listings in order to gain better understanding of the importance of search and matching frictions in this market. Third, we use this context to develop a framework for understanding tax incidence and efficiency costs of transaction taxes in search and matching markets more generally. Other contexts where similar issues arise are labor markets and financial transaction taxes. Fourth, our theory and empirics allows for studying the impact of discontinuous incentives on the *existence* of the market itself.

Our empirical approach relies on variation generated by the discontinuous nature of the taxes imposed in New York and New Jersey. In a nutshell, these taxes are levied as a function of the appropriately defined purchase price. A prominent example is the so-called "mansion tax" in New York state (since 1989) and New Jersey (since 2004) that applies to residential transactions of $1 million or more. The tax rate is 1% and is imposed on the *full* value of the transaction so that a $1 million sale is subject to a $10,000 tax liability, while a $999,999 transaction is not subject to the tax at all. In New York City, all real estate transactions are also subject to the real property transfer tax (RPTT) and in New Jersey they are subject to the Realty Transfer Fee (RTF)—both of these schedules happen to have (smaller) discontinuities as well, as we discuss in Section 2. Hence, all of these taxes create tax notches (see Slemrod, 2010), while the introduction of the tax in New Jersey also creates a time discontinuity.[1] Furthermore, the statutory incidence is different for the mansion tax (which is the responsibility of the buyer) than for the New Jersey RTF and New York City RPTT (which are the responsibility of sellers, with the exception of sales of new constructions in NYC). Interestingly, such discontinuities are not uncommon — for example, they are also present in the UK (Besley et al., 2013; Best and Kleven, 2013) and D.C. (Slemrod et al., 2012).

Our results allow us to reach three sets of conclusions. First, and perhaps least surprisingly, we find that the tax distorts the price distribution resulting in significant bunching just below $1 million. This bunching is evident in the distribution of sales in New York displayed in Figures 1 and 2.[2] A similar pattern appears in New Jersey after the introduction of the tax. Figure 3 and Figure 4

[1] There are also *geographic* discontinuities that we do not exploit: the RPTT changes at the New York City border, and both RTF and the pre-2004 mansion tax change at the New Jersey-New York border.

[2] Figure 1 corresponds to the whole state, while Figure 2 is for New York city itself. We chose to present the two figures with different binning and overlaying the fit on just one of them in order to present both the visual evidence of the effects and illustrate salient features of the data (round number bunching) that we control for in the empirical analysis.

demonstrates that the onset of this effect is immediate. The bunching we observe is substantial: our estimates robustly indicate that about \$20,000 worth of transactions shift to the threshold in response to the \$10,000 tax. The strength of this effect does not vary with our proxies for tax evasion and we show, using listings data, that a distortion of similar magnitude is already present when properties are first advertised by sellers, which we interpret as inconsistent with tax evasion. We find some evidence that the effect is weaker (but still very strong) for newly built properties that sell when already finished, which suggests that real adjustments to the characteristics of a property may be part of the effect. Still, we conclude that real responses do not fully account for the extent of bunching, and the tax near the threshold imposes a substantial burden on sellers. Results from smaller discontinuities that shift statutory incidence are consistent with this conclusion.

Second, we build a theoretical framework to illustrate and test for unraveling of the market in the neighborhood of the threshold—the possibility that the tax locally destroys trades of matches with remaining positive surplus. Figures 1, 2 and 3 show that the distribution of prices features a large gap above the threshold and our test relies on the comparison of its size with the extent of bunching below the notch. If all transactions with positive surplus continue to transact, the gap is expected to reflect observations shifting to the threshold, in particular it is expected *not to be bigger* than the extent of bunching. This is a testable prediction. If it fails, it implies that some of the observations are not occurring over and beyond the standard extensive margin response—the phenomon that we refer to as "unraveling". Transactions that could otherwise occur near the notch may not take place at all, because sellers, who face a large burden from these sales, may instead opt out or continue waiting, or because buyers may prefer to continue searching in order to benefit from locally depressed prices.

The implementation of our test for unraveling is straightforward and illustrated in Figure 5 (that we will explain in more details later in the paper). Conceptually, we estimate bunching at the threshold by constructing the counterfactual distribution based on the data to the left of the threshold and comparing this to the observed bunching at the notch. We estimate the gap by constructing the counterfactual distribution based on the data well to the right of the threshold (i.e., affected by the tax and, thus, accounting for the standard extensive-margin response) and comparing this to the observed gap in sales above the notch.[3] We show that the difference between the size of the gap and the extent of bunching is a lower bound for the number of "missing" observations and we use this lower bound to test for the presence of unraveling. Intuitively, the after-tax counterfactual used to estimate the size of the gap is, by construction, already net of the standard extensive margin response (matches that have surplus lower than the tax) and thus the gap only reflects sales that have shifted to the threshold or matches with continued positive surplus in the presence of the tax that do not sell. We argue that these missing transactions would have sold if they were far from the threshold, but are discouraged by the incentives presented by the notch to continue searching.

[3] In practice, we usually simply allow for the shift in the distribution at the threshold to parsimoniously capture the two different counterfactual distributions. We also report results that rely on separate estimation on both side.

We indeed find that more transactions are "missing" than we can observe bunching at the threshold, indicating that the market unravels locally. This effect is large: our baseline estimates indicate that over $40,000 worth of transactions that would still yield positive surplus even with the tax do not take place. In our baseline specification, this corresponds to 2,800 missing transactions in New York city, out of 380,000 that occurred over the whole period. Hence, by our estimates, this one percent tax, applying at a relatively large threshold, managed to eliminate 0.7% of transactions due to the unraveling effect. To reiterate, this response is conceptually different from the standard demand response that is due to higher taxes discouraging transactions with low surplus: it corresponds to eliminating transactions with positive surplus in the presence of the tax. This additional extensive-margin response indicates that the substantial friction introduced by the transaction tax hampers functioning of the market in the neighborhood of the threshold. Given our sources of variation, estimating the standard response would require making strong assumptions about comparability of distributions with and without taxes, and we do not pursue it in this paper.

Recognition and empirical identification of the type of extensive margin response that we focus on is novel to the literature. It is present because of the search frictions and may apply in any market where search is present or population affected by distortionary incentives is not fixed. Much of the literature using "kinks" and "notches" focuses on income tax in the contexts where only intensive margin is of interest and hence this point has not been recognized before.[4] In particular, our results indicate that in situations where taxable units are endogenous (as in housing, but arguably not under the income tax far from filing threshold), exits around kinks/notches cannot be assumed to be the standard extensive margin and hence such responses cannot be easily generalized to reflect the effect on sales elsewhere.

Our third set of conclusions finds evidence indicating that the impact of the tax is not limited to the neighborhood of the threshold, but extends much further. Both price reductions while properties are listed and discounts (the difference between final advertised and sale price) increase permanently above the threshold, indicating that the search and matching process is affected everywhere by the tax. Furthermore, we find that in the presence of the tax listing prices are a weaker signal of the final sale price of the property. Relying on our theoretical arguments, we interpret this greater dispersion of sale price conditional on asking price as corresponding to increased deviation from the efficiency-maximizing matching equilibrium and conclude that a general transaction tax increases inefficiency in the search process.

A small literature focuses on the effect of transfer taxes on the functioning of the real estate market (Benjamin et al., 1993; Van Ommeren and Van Leuvensteijn, 2005; Dachis et al., 2012). Contemporaneously, three other papers (Slemrod et al., 2012; Best and Kleven, 2013; Besley et al., 2013) look at similar distortions to the distribution of final prices (but not listings) in the United Kingdom and Washington, D.C.. These studies focus on the standard extensive margin response (the general effect on sales) to policy changes, rather than incidence, listings and search

[4]More speculatively, we briefly comment in Section 6 on figures from our online appendix that indicate possible changes in patterns of behavior during the search process.

frictions as we do. We are also unique in showing that the extensive margin effect of the tax goes beyond eliminating transactions with negative net-of-tax surplus. Beyond offering the first, to our knowledge, evidence of this type of an effect, these results also cast doubt on generalizing from responses around notches and kinks (where market can unravel) to elsewhere (where only standard extensive margin response should be present) in the presence of matching frictions.

Another strand of literature to which we contribute analyzes the search and matching process in the real estate market. Several studies focus on the role of listing prices and bargaining in determining the final sale outcome (c.f. Han and Strange, 2012, 2013; Merlo and Ortalo-Magne, 2004; Haurin et al., 2010), while a few apply more general search models to real estate data (c.f. Carrillo (2012); Genesove and Han (2012)). However, these studies do not explicitly identify the effect of transaction costs, such as transfer taxes, on outcomes. A related line of study focuses on the role of real estate agents, attempting to unbundle the effect of cost from information provision (Levitt and Syverson, 2008; Jia and Pathak, 2010; Bernheim and Meer, 2013). Finally, a number of empirical papers incorporate information available in real estate listings data to the study of seller behavior in the housing market (Genesove and Mayer, 2001; Carrillo and Pope, 2012).

Our paper is also related to work on behavioral responses to taxation. As in the research on responses to income taxation, we are interested in separating real, timing, avoidance, and evasion responses (Slemrod, 1990; Saez et al., 2012). Contrary to that strand of work, our context requires considering both sides of the market. There has been a recent revival of interest in estimating the incidence of specific taxes/transfers (e.g., Doyle Jr. and Samphantharak, 2008; Mishra et al., 2008; Hastings and Washington, 2010; Marion and Muehlegger, 2011). Real estate tax is more complicated due to non-homogeneity of goods traded, and the closest analogue in the literature is work on incidence of income/payroll taxes or credits (Rothstein, 2010; Saez et al., 2011).

The structure of the paper is as follows. In the next two sections, we discuss the institutional and policy context and our data. In Section 4 we present our theoretical framework. We start by introducing a bargaining framework that illustrates the effect of the tax for a particular match, followed by discussion of frictionless equilibrium and predictions regarding the effect of the tax on the price distribution. We then derive simple testable implications of the presence of frictions to matching. In Section 5, we present empirical results about the distribution of prices, both graphical evidence and local incidence estimates for various types of taxes, relying on price and listings data. We also show evidence for various subsamples in order to shed a light on the role of evasion and real adjustments. In Section 6, we focus on distortions to the matching process near the threshold. We interpret econometric estimates of bunching and missing transactions, and analyze the effect on timing and real estate agent use. In Section 7 we demonstrate the global effect of the tax on discounts and informational content of listings. Conclusions are in the final section.

2 Policy

Real estate transfer taxes are common across the U.S. These taxes are applied to the sale price of real property, and range from as low as no tax in Texas to 2% in Delaware. In New York and New Jersey, the tax rates change discontinuously with total consideration, creating corresponding notches in total tax liability. Table 1 contains details of the relevant tax schedules. One notch arising in both states is due to the mansion tax: a 1% tax on the total consideration for homes costing $1,000,000 or more. Under the mansion taxes of both New York and New Jersey buyers' total tax liability jumps by $10,000 when the sales price moves from $999,999 (where the tax does not apply) to $1,000,000 (where the tax comes into effect). In New Jersey, the mansion tax was introduced on August 1st, 2004, and covers all residential real-estate transactions. In New York state, the mansion tax was introduced in 1989 and applies to the sale of individual coop and condo units, and one-, two-, and three-family homes, with few exceptions.[5]

Real estate sales in New York City and New Jersey are subject to additional taxes with discontinuous average rates. In New York City, the Real Property Transfer Tax (RPTT) applies to residential sales (as defined for the New York mansion tax) with a rate of 1% if the total consideration is $500,000 or less, and 1.425% above $500,000. Commercial sales are also subject to the RPTT at a rate of 1.425% below $500,000 and 2.625% above. Unlike the mansion tax the statutory incidence of the RPTT falls on the seller by law, however it is customary for the buyer to pay the tax when purchasing directly from a sponsor (i.e., purchasing a newly developed condo or a newly offered coop). Thus, the RPTT is a unique tax in that there is variation in the statutory incidence.

Residential sales in New Jersey are subject to the Realty Transfer Fee. This transfer fee (or tax) has a non-linear schedule (see Table 1) that shifts when total consideration is greater than $350,000. The marginal tax rate for consideration above $200,000 is 0.78% if the total price is less than or equal to $350,000, while this tax rate jumps to 0.96% when the total price is greater than $350,000. Moving from a price of $350,000 to $350,001 increases the buyer's tax liability by $630.

3 Data

We study administrative records on real estate transactions in New York state and New Jersey as well as historical real estate listings in New York City. Sales records, which cover the universe of recorded real estate transactions in the given geography and time period, come from three sources: the New York City Department of Finance's (NYCDOF) Annualized Rolling Sales files for 2003–2011 (covering all of NYC), real property transfer reports compiled by the New York state Office of Real Property Services (NYSORPS) for 2002–2006 and 2008–2010 (all of NY State, excluding the five counties of NYC), and SR1A property transfer forms collected in the New Jersey Treasury's

[5]Exceptions are as follows. If a residential unit is partially used for commerce, only the residential share of the total consideration is subject to the tax (although the entire consideration is still used to determine if the tax applies). Similarly, multiple parcels sold in the same transaction are taxed as one unit unless the parcels are evidently not used in conjunction with one another. Vacant lots are exempt from the mansion tax, and, finally, any personal effects sold with the home are deducted from the total consideration for tax purposes (but are subject to state sales taxes).

SR1A file for 1996–2011 (all of NJ). These records contain details of each transaction, including the date of sale, the total consideration paid by the buyer, the address of the property, the property type (e.g., one-, two-, or three-family home, residential coop or condo, etc.), the year of construction of the building (in NYC and NJ) or whether the property is newly constructed (in New York state), and whether the sale is arms-length (NYSORPS only).[6] We also use deeds records for 1996–2008 for NYC collected by an anonymous private data provider from the county records, which indicate whether the buyer relied on a mortgage (see Appendix B for more information about data sources).[7]

We identify sales that are subject to each of the transfer taxes. Misclassification in taxable status, if any, will introduce a bias against finding effects, however it is unlikely to be substantial: our information comes from administrative records that contain sufficient information to classify, even though an explicit taxability flag is not provided. In New Jersey, we consider all "residential" sales to be taxable (mansion tax and Transfer Fee). For New York state, we define all single-parcel residential sales of one-, two-, or three-family homes, condos, and seasonal properties as subject to the mansion tax. Finally, in New York City we define all single-unit (non-commercial) sales of one-, two-, or three-family homes, coops, and condos as taxable.[8]

We match the Rolling Sales data for Manhattan to a subset of historical real estate listings in order to get a broader picture of the effect of the tax on the real-estate search process. Our listing data comes from the Real Estate Board of New York's (REBNY) electronic listing service and covers all closed or off-market listings between 2003 and 2010. REBNY is a trade association of about 300 realty firms operating in New York City and represents a substantial share of listings in Manhattan and Brooklyn. Members are required to post all listings and updates to the electronic listing service. We limit attention to the more complete Manhattan listings, which accounts for approximately 45% of Manhattan sales in the Rolling Sales files.

From the REBNY listing service we observe details of all (REBNY-listed) closed or off market listings since 2003. These data include initial asking price and date of each listing, all subsequent price updates, and the date the property sells or is taken off the market. To acquire the final sale price, we match these listings by precise address (including apartment number) and/or tax lot to the NYCDOF data for Manhattan. We obtain approximately a 90% match rate for listings identified as "closed" in the REBNY data.[9]

Table 2 presents descriptive statistics from the three sources of sales data. Overall, we have records for 3,256,597 taxable sales (with non-zero sale price) spanning 1996–2011. The distribution is skewed: mean sale price is higher than the median. Unsurprisingly, prices are highest in New York City. Although median (and mean) sale price is well below the $1,000,000 threshold of the

[6]This definition excludes sales between current or former relatives, between related companies or partners in business, sales where one of the buyers is also a seller, or sales with "other unusual factors affecting sale price." See Appendix B for more details.

[7]We are grateful to Chris Mayer and the Paul Milstein Center for Real Estate for access to this data.

[8]While multi-parcel sales in New York state are typically subject to the mansion tax, such a sale may be split for tax purposes if structures on adjacent parcels are not used in conjunction with or clearly related to one-another. Since we cannot identify such cases in the New York City and New York state data, we err on the side of caution and exclude all multi-parcel and partially-commercial sales sales from taxable status.

[9]The match-rate is continuous across the tax thresholds. See Appendix B for details on the matching process.

mansion tax, there are still several thousand sales per geography within \$50,000 of the mansion-tax cutoff. Table 3 presents the count of taxable residential sales and median prices over time for the three regions. The growth of housing sales and prices throughout the early 2000s is evident here, as is the subsequent drop in total sales and median price at the onset of the recession in 2007/2008.

Table 4 presents statistics for the matched REBNY listings data. The sales covered by REBNY have much higher prices, on average, than the general NYC rolling sales data (\$1.24 million versus \$658,000). This is due to the REBNY data only covering sales in Manhattan, which is considerably more expensive than the outer boroughs. About 64% of the homes in our sample end up closing (rather than being taken off the market). At the same time, 8% of listings do not close but have a corresponding sale in the rolling sales data, which we interpret as corresponding to either direct sales by owner or sales with a non-REBNY agent. Homes that do not close spend more time on the market (200.5 days vs. 146 for sold properties). These statistics suggest that search frictions are non-negligible in the housing market—the process of finding a buyer for a home is lengthy and sellers are often unable to find a match. There also appears to be bargaining between buyers and sellers: for properties that close and are matched to the rolling sales data, the average discount between the initial asking price and the sale price ($\frac{initial - sale}{initial}$) is 5.9%: 3.2% discount from initial asking to final asking price and a 2.8% discount from final asking price to sale price. However, the median listing in our sample has no price updates between the initial and final asking prices.

4 Theoretical Framework

To interpret our empirical findings, we present a simple model. We first discuss implications of taxation in a bargaining framework given a match between a buyer and a seller. We then characterize the equilibrium and its responsiveness to taxation, absent search frictions. The equilibrium in this situation corresponds to assortative matching. We follow with a discussion of how the price distribution might respond when individuals search and need not transact conditional on matching. Finally, we elaborate on how the equilibrium price dispersion (which is present when there are matching frictions) may respond to taxation, and how these effects can be empirically characterized by relying on observable information (including listings data).

4.1 Bargaining

We start with a bargaining model that clarifies the nature of distortions to the price distribution around the notch and relates tax incidence observed at the notch to the bargaining power that determines incidence elsewhere. For now, we abstract from equilibrium considerations and instead characterize pricing behavior given a match. This is a building block of the equilibrium analysis that we come back to below.The Nash bargaining model itself is formally presented in the Appendix A. Here, we introduce the notation and illustrate key results on Figure 6. The figure corresponds to the lump-sum tax case, but the main insights apply as well to the proportional tax that we discuss in more details in the Appendix.

Consider a single match (b, s) between a buyer with a reservation value of b and a seller with reservation value s. Given the price that a buyer and a seller negotiate, p, and a lump-sum tax T imposed on the buyer (as in the case of the mansion tax), the parties end up with surpluses of $S^B = b - T - p$ and $S^S = p - s$, respectively. We assume Nash bargaining with seller's weight β so that the price maximizes $\beta \ln(p - s) + (1 - \beta) \ln(b - T - p)$ and is, thus, set to $p(b, s) - \beta T$, where $p(b, s) \equiv \beta b + (1 - \beta)s$ is the price absent taxes. Consequently, the surplus of each side is equal to $S^B = (1 - \beta)(b - T - s)$ and $S^S = \beta(b - T - s)$, which implies that the parameter β determines how the total surplus $b - T - s$ is split between the two parties.

Necessarily, the incidence of the tax is determined by β as well. Although it follows automatically from this framework, it is worth highlighting that the party with *lower* bargaining power bearsmansion_tax_revisions_DMDec30.lyx . a *lower* share of the tax. This party has a lower claim to the surplus in the first place and, symmetrically, experiences a lower reduction in surplus when the tax is imposed. At the extreme, when $\beta = 0$ or $\beta = 1$, one of the parties has no bargaining power and no surplus, and thus is completely inelastic so that it cannot bear any burden of the tax.

In addition to changing the sale price and the relative distribution of the surplus, transaction taxes may discourage sales altogether. Transactions take place when the surplus is non-negative $(b - T - s \geq 0)$. All matches (b, s) that satisfy $\beta b + (1 - \beta)s = p$, for some p, sell at exactly the same price (equal to $p - \beta T$) or do not sell at all if total surplus is negative. By reducing the surplus, the uniform lump-sum levy discourages some sales. However, the lump-sum tax does not lead to re-ranking of transactions. All prices simply adjust by βT, so that transactions that were occurring at the same price absent the tax continue to sell at equal (although different from the original) prices. This lack of re-ranking is not general: it does not survive considering proportional rather than lump-sum taxation, but it simplifies the following discussion and provides a natural benchmark (and, as discussed, in the appendix, the key qualitative results generalize to the proportional tax case).

We illustrate our formal results regarding the price and sales responses to the tax graphically on Figure 6. The figure shows reservation values of buyers and sellers on the two axes, and also allows for tracing prices. The contract line (the relationship between prices of buyers and sellers) in absence of the tax requires that the prices of buyers and sellers have to be the same: $p^B = p^S$; while in the presence of the tax above the notch it is given by $p^B = p^S + T \cdot (p^S \geq H)$, where H is the notch ($H =$ \$1,000,000 for the mansion tax). This pricing/budget constraint is represented by the solid black line. The straight green line and other lines parallel to it show the locus of matches with the constant value of $p(b, s)$. In the presence of the lump-sum tax all matches on any of these lines sell at the same price, given by $p(b, s) - \beta T$, which corresponds to a point where a given constant price line intersects with the contract line. Transactions in the gray shaded area, marked "Z", have positive surplus and would sell without the tax, but they do not sell when the tax is present because net-of-tax surplus turns negative. Matches in this region reflect the standard extensive margin response.

Price adjustments are affected by the presence of the notch and can be broken-down into four cases that depend on the initial price. Case 1 are transactions that originally occur below the price of H and are not affected. Case 2 are transactions that originally sell at the price between H and $H + \beta T$ and would sell below H if the tax was uniform, but in the presence of the notch occur there. These transactions are illustrated in Figure 6 by the yellow area marked "A" that is bounded by the $p(b, s)$ schedules corresponding to $p(b, s) = H$ and $p(b, s) = H + \beta T$. Given the assumption of maximizing the Nash-bargaining objective function, transactions that originally sell at a higher price than $H + \beta T$ may sell at the notch depending on whether total surplus at the notch or with the tax is higher. We show that for some intermediate range of original (absent-tax) prices above the threshold (Case 3), some transactions will bunch at the notch (region "B"), while others will sell at a new price above the notch (region "C"). However, when the original price is high enough (greater than some *finite* pre-tax price corresponding to the green line) no transaction will bunch, constituting our final case (region "E").

In the appendix we establish that the qualitative characterization of the effect of taxation on prices described by Figure 6 is general. The formal model delivers additional results that are not a priori obvious. For any β there exists a pre-tax price above which transactions are not affected by the notch — this boundary is determined by considering matches involving a seller with reservation value of zero. Furthermore, and less intuitively, there exists a single *finite* bound for such maximum pre-tax prices above which transactions do not bunch that applies uniformly for any value of β. In the lump-sum tax case, the tight (applying as $\beta \to 0$) bound is the solution to $\ln(H) - ln(p) - \frac{H-T}{p} + 1 = 0$ and corresponds to $p \approx \$1,144,717$ when $H = \$1,000,00$ and $T = \$10,000$. In the case of the proportional tax, the bound solves $\ln(H) - \ln(p) - \frac{H - p\ln(1+t)}{p} + 1 = 0$ (where t is the marginal tax rate, 1% in the case of mansion tax) corresponding to $p \approx \$1,155,422$. While this precise characterization is the consequence of functional form assumptions in the case of Nash bargaining, it does provide a theoretical justification for the assumption that we make in our empirical analysis that only matches in some finite omitted region might bunch. In our empirical analysis, we will use this theoretical bound for defining the omitted region in our baseline specification (but of course we explore the sensitivity of this choice).

4.2 Equilibrium

The framework that we have introduced so far focuses on price determination given a match. This is a component of the equilibrium description—given matches that lead to sales, we assume Nash bargaining as the approach for determining the price. The missing component of the model is a description of how matches form. Providing a complete search framework is beyond the scope of this paper and, in fact, we are not aware of a framework in the literature that would incorporate two-sided search in the real estate context (Carrillo, 2012, makes a step in this direction by setting up, but not explicitly solving, a model of this kind). We make two arguments that we then investigate empirically. First, we consider what the distribution of prices reveals about the distribution of the underlying matches and efficiency of the equilibrium. Second, we consider which of the matches are likely to be "stable" in the presence of the tax.

The simplest way to introduce equilibrium consideration in this framework is to assume random matching followed by very large frictions (search costs) preventing both parties from further search so that, conditional on a match, the only decision to make is whether to transact. Under this assumption, there is a matching technology $M(b, s)$ that results in some (smooth) distribution of matches over (b, s) and Figure 6 reflects which of those matches correspond to transactions.

An alternative place to start is by considering a situation with no frictions to search. Suppose that we have an equal number of buyers and sellers. Absent taxation, the overall surplus from a match in our Nash bargaining model is strictly supermodular (it is given by $\ln(b - s) + constant$). Hence, the equilibrium and, simultaneously, the *efficient allocation* that maximizes the overall surplus involves positive assortative matching. In the presence of taxation, the surplus for transactions not subject to the tax and not at the threshold remains $\ln(b - s) + constant$, while the surplus for transactions subject to the tax is $\ln(b - s - T) + constant$ (which is, naturally, also supermodular). Hence, *within* each of these groups maximization of the overall surplus involves assortative matching.

The efficient allocation absent taxation corresponds to an increasing profile of matches (b, s). This profile is illustrated on Figure 6 using a wiggly solid line. If these matches were to remain when the tax is introduced, a match corresponding to point X would be the highest priced one that is not subject to taxation, the match marked by Y would be the lowest priced one that does not shift to the notch, and matches between X and Y would move to the notch. Introduction of the tax may affect which matches take place in the equilibrium. The efficient allocation will retain the main features visible on Figure 6, although the actual equilibrium profile and points X and Y need not coincide with the allocation absent tax distortions. As argued above, the equilibrium matches subject to the tax will continue to be assortatively matched—that is, matches above point Y will lie on an increasing profile. Similarly, matches below point X will form an upward-sloping profile. For matches that are priced at the threshold, the price is fixed at H so that any permutation of residences between buyers would deliver exactly the same surplus (a feature that is arguably peculiar to this model), so that the precise identity of matches between X and Y is indeterminate.[10]

The equilibrium allocation in the presence of frictions will not be efficient, although efficiency serves as a natural reference point. Trivially, our theoretical framework implies that under the efficient allocation variance of the price (or, equivalently, buyer's type) conditional on seller's type (and vice versa) is zero, because the efficient allocation corresponds to an upward sloping line.[11] In the presence of frictions, matches would occur not just on the efficient allocation line as in Figure 6, but rather could be spread over the rest of the region corresponding to surplus from transacting. Intuitively, one might expect that an increase in dispersion of prices for a given type

[10]The location of X and Y may change because fewer transactions may take place in the presence of a tax. Figure 6 appears to preclude this possibility by using the efficient matching schedule that does not involve matches that are crowded out by the tax, but it need not be so in general.

[11]While the efficient allocation is not unique in the bunching region A and B on Figure 6, the price is constant and equal to the threshold level in that region.

of home corresponds to an allocation that is further from the efficient one. While we do not prove this result, it has to be so at least when the frictions are small. To examine whether transaction taxes affect the efficiency of the housing market allocation, we test whether there is an increased dispersion of prices in the presence of the tax. Because matches are indeterminate in the bunching region, this exercise is of interest for transactions that are not local to the threshold—i.e., those far enough from the threshold in either direction. We discuss this further in Section 4.4.

4.3 Measuring the Impact on the Price Distribution

We expect that the tax affects the distribution of home sales by inducing both bunching of sales at the notch and by creating a "gap" in the distribution above the notch, and we estimate both. Our objectives are twofold. First, descriptively, these estimates allow us to quantify the magnitude of distortions to the price distribution. Second, we use these estimates to back out values analogous to the mass in regions A and B of Figure 6. Intuitively, extra transactions at the threshold correspond to regions A and B, while transactions that are missing from the distribution above the threshold (relative to the distribution further to the right) reflect region B. In principle then, these values may be used to recover the mass in region A that is tightly linked to the bargaining parameter, β. However, in what follows, we argue that the tax notch provides a strong incentive for neighboring "productive" matches—those close to the boundary between regions B and C, which have positive surplus in the presence of the tax—to break. Consequently, there may be more transactions missing from the distribution above the notch than locate at the notch itself. We show that our estimates can be used to test for and bound this local extensive-margin response: if such a response is not present, the excess number of transactions at the notch need to exceed the gap in the distribution. We view this test as one of the central contributions of our paper, as it corresponds to testing for unraveling of the market due to the presence of the notch. Equivalently, this is a test of whether the extensive margin response is standard (the gray region in Figure 6)—a condition that is necessary to generalize from estimates based on a notch to behavioral responses to a general tax (as done by Best and Kleven (2013) and Slemrod et al. (2012)). Previewing our results, we find that overwhelmingly the answer is that it is not. The rest of this section serves to define quantities that we estimate and to formalize the test for the local unraveling of the market.

Observed and counterfactual distributions. We first consider how the distribution of sales is distorted by the tax notch. The discussion is graphically illustrated on Figure 5, which is a distribution analogue of Figure 6 (with region labels and coloring corresponding to it). We denote by $F(p)$ the "true" (population) price distribution in the presence of the actual (notched) tax from which our observations are drawn. We denote by $F_T(\cdot)$ the observed cumulative population price distribution—a draw from $F(p)$. In order to characterize and interpret distortions to the distribution, we presume that there is a set of potential matches (from some matching technology) that may result in transactions. We leave the origin of the set of matches unspecified, and simply assume that there is some matching technology that we take as given. Matches are indexed by

i and have associated with them three prices $(p_i, \check{p}_i, \tilde{p}_i)$: the actual price p_i in the presence of a "notched" tax, the shadow price $\check{p}_i$ that would prevail if the tax did not apply anywhere, and the shadow price $\tilde{p}_i$ that would prevail for the same transaction if the tax was proportional everywhere (i.e., involved no threshold). In the context of our model as illustrated on Figure 6, $p_i = \check{p}_i$ for transactions taking place in region D, $p_i = \tilde{p}_i$ in regions C and E, and $\check{p}_i \geq H > p_i$ in regions A and B. We assume that $\check{p}_i > \tilde{p}_i$, which excludes the polar case of incidence fully borne by buyers, but allows for simplifying notation. We also allow for either of the prices to be infinite, corresponding to the transaction not taking place in a given regime. We do not rule out in general that the notched tax affects the equilibrium distribution everywhere (even below the threshold): $F(p)$ corresponds to the actual equilibrium outcome; however, $\check{p}_i$ and $\tilde{p}_i$ are prices specific to matches that form in the observed equilibrium given a notched tax, so that their marginal distributions do not reflect any changes regarding which matches would form if the the tax was removed or replaced by one that is proportional everywhere.

At some abuse of the notation, we use the region descriptors from Figure 6 (e.g., A) to denote the set of matches and the mass (probability) of *transactions* in the corresponding region. Transactions that are distorted by the threshold have $\check{p}_i \geq H > p_i$ and come from a number of different sources: $\tilde{p}_i < H$ (region A except for A'); $\tilde{p}_i = \infty$ (region A'); and $\tilde{p}_i \geq H$ (region B). We note that $A + B = P(\check{p}_i \geq H > p_i)$ and $B = P(p_i < H \leq \check{p}_i \wedge \tilde{p}_i \geq H) = P(p_i < H \leq \tilde{p}_i)$.[12]

Crowd out of productive matches. We allow for the possibility that not every productive match corresponds to a transaction: i.e., that there may be transactions for which $\tilde{p}_i < \check{p}_i < \infty = p_i$—those with sufficient surplus to survive the tax, but not occurring in the presence of the notch. To see why, recall that our basic framework assumed that equilibrium matches in the neighborhood of the notch form in a way similar to those away from it, with only the outcome of the bargaining process affected. However, proximity to the threshold provides strong incentive for some buyers and sellers in matches near the notch to continue searching. Firstly, consider buyer-seller pairs who would move to the threshold if a sale occurs (region B of Figure 6). Sellers in this region—who face a substantial reduction in sale price in moving to the notch—may prefer searching for a buyer who is willing to buy above the notch. Secondly, buyers in the buyer-seller pairs that would transact in the gap region above the notch (region C) may have an incentive to return to the market to find a seller with slightly lower reservation value. Whether this type of local extensive-margin response is present is of interest in its own right, corresponding to both an efficiency loss due to notched tax in markets with search frictions (productive matches that do not transact) and the importance of search in the housing market. We assume that such exits do not occur for transactions below the threshold ($\check{p}_i < H$, region D) and for transactions that have sufficiently high prices (region E, $\tilde{p}_i > \bar{P}$ for some sufficiently large price $\bar{P}$). We denote the mass of such exits that comes from matches that could otherwise sell above the threshold in the presence of proportional taxation

[12] Transactions to the left of the threshold (region D) have $p_i = mansion_tax_revisions_DMDec30.lyx.\check{p}_i$, and those with $p_i = \tilde{p}_i$ correspond to regions C and E. The gray region—transactions with surplus low enough that the tax crowds them out—have $\check{p}_i < \infty = p_i = \tilde{p}_i$.

(regions B and C) by $M \equiv P(H \leq \tilde{p}_i \leq \bar{P} < \infty = p_i)$.

To estimate the mass in these regions, we rely on two counterfactual distributions: one that corresponds to the non-taxable regime and another corresponding to the taxable one. $F_0(p) = P(\check{p}_i < p)$ is the counterfactual distribution corresponding to the non-taxable regime (to the left of the threshold); equivalently, below the taxable threshold, H, F_0 is the true distribution net of transactions that are affected by the presence of the tax ($\check{p}_i \geq H > p_i$). Correspondingly, we define

$$I = F(H) - F_0(H) = P(\check{p}_i \geq H > p_i) = A + B$$

as the number of observations that shift below the threshold due to the tax. Given the observed distribution F_T and an empirical estimate of the counterfactual distribution $\hat{F}_0$, we can construct an empirical estimate of the volume of responsive sales making up I as $\hat{I} = F_T(H) - \hat{F}_0(H)$. In practice, we construct the counterfactual price distribution of sales $\hat{F}_0$ by relying on the actual distribution F_T to the left of the notch, but omitting sales near it—the specifics are in Section 4.5.

Local incidence. Given an estimate of the counterfactual distribution F_0 to the right of the threshold, we can also define a dollar measure $\hat{h}$ as

$$F_0(H + \hat{h}) - F_0(H) = \hat{I} \tag{1}$$

to represent the magnitude of the shift to the threshold. We refer to $\hat{h}$ as "local" or "reduced-form" incidence of the tax. Our preferred interpretation is that it represents the average amount of money that is lost by sellers participating in the marginal transactions affected by the presence of the threshold. However, the value of $\hat{h}$ does not, on its own, inform us about the underlying bargaining power of the two sides of the market and hence does not reflect incidence of the tax away from the threshold. The estimate of bunching, $\hat{I}$, is represented by the blue area in the first panel of Figure 5. The local incidence is obtained by finding the dollar value $\hat{h}$ such that the integral under the counterfactual to the right of the notch is equal to the excess mass (represented as the green area of the first panel of Figure 5). Similar to the "kinked" budget-set methodology outlined by Saez (2010), this local incidence represents the average reduction in price sustained by the marginal buyer-seller pairs that bunch relative to their non-taxed sale price.

We construct and estimate $\hat{h}$ throughout, but our interpretation of $\hat{h}$ as local incidence depends on assumptions about the nature of the counterfactual distribution. With the exception of our data for New Jersey prior to implementation of the tax, we do not observe $\check{p}_i$ for values greater than H at all. When we use data below the notch to project F_0 above the notch, the interpretation of $\hat{h}$ requires additional assumptions. If F_0 below the notch coincides with the distribution absent taxation, then the projected counterfactual above the notch corresponds to the distribution absent taxation as well, so that $\hat{h}$ can be thought of as a reduced form dollar estimate of local incidence (this interpretation also applies when we build our counterfactual using the non-distorted distribution in New Jersey before the tax was introduced). When the untaxed part of the distribution (below the notch) is affected by the tax via general equilibrium effects, the counterfactual estimate above the notch does not have a clear interpretation. In that case, the dollar value $\hat{h}$ that relies on projection of F_0 above the threshold is not directly interpretable (although it remains a convenient way of standardizing the results). However, our estimate of excess mass $\hat{I}$ (and unraveling estimates that depend on it), that relies on F_0 below the threshold only remains valid.

Gap. We also construct a measure of the gap to the right of the notch. We define a counterfactual distribution $F_1(p) = P(\tilde{p}_i < p)$ corresponding to the region subject to the tax: it is the distribution under a proportional tax (but, of course, without allowing for equilibrium adjustment to the set of matches). We presume that there is a known value of $\bar{P}$ such that F_1 and F coincide for prices greater than $H + \bar{P}$ (by Theorems A.3 and A.4, this has to be so in our framework; also recall that we assumed away exit of productive matches for high enough prices) and define the gap in the distribution as

$$G = [F_1(\bar{P}) - F_1(H)] - [F(\bar{P}) - F(H)]$$

i.e., as the difference between the number of transactions taking place in the presence of taxation with and without the notched implementation of the tax. For the estimation, we replace F by F_T, and F_1 by its empirical estimate $\hat{F}_1$. The estimate of $\hat{F}_1(\bar{P}) - \hat{F}_1(H)$ reflects the expected number of observations in regions B and C, while $F_T(\bar{P}) - F_T(H)$ is the actual number of observations in region C. Using our definitions, we can show that

$$G = P(\bar{P} \geq \tilde{p}_i \geq H \wedge (p_i < H \vee p_i = \infty)) = P(p_i < H \leq \tilde{p}_i) + P(H \leq \tilde{p}_i \leq \bar{P} < \infty = p_i) = B + M$$

Recall that M represents the mass of transactions that would have taken place under a proportional tax at prices exceeding the threshold, but do not take place in the presence of the notch. Thus, the gap reflects two effects: exit from the market of productive matches and shift to the threshold.

To reiterate, the gap reflects transactions that are missing from the distribution to the right of the threshold relative to the counterfactual with taxes. In particular, it does not include the standard extensive margin response—matches with surplus small enough that they are no longer economically viable in the presence of the tax (the gray region Z in Figures 5 and 6).

Testing for market unraveling. The gap, G, and behavioral response, I, are related: both partially reflect transactions in region B—those that would sell at prices higher than the threshold in the presence of a continuous tax, but sell at the threshold when it is discontinuous. Clearly,

$$I - G = A - M$$

We report an estimate of $\hat{G} - \hat{I}$ converted to a dollar figure: $\hat{Z} = \hat{h} \cdot \left(\frac{\hat{G}}{\hat{I}} - 1\right)$.[13]

[13] Alternative definitions would be to define $F_1(H + \hat{Z}) - F_1(H) = \hat{G} - \hat{I}$ or $F_0(H + \hat{Z}) - F_0(H) = \hat{G} - \hat{I}$. The choice we implement has two advantages. First, it is in terms of the distribution F_0 so that it is directly comparable to $\hat{h}$. Second, knowing $\hat{h}$ (which we report as well) allows for directly recovering an alternative metric of the gap and

If the market does not unravel in the neighborhood of the notch ($M = 0$), then $I - G = A \geq 0$. Hence, given estimates of I and G, we can then test whether the tax destroys productive matches.

Remark 1. A testable hypothesis $\hat{Z} \leq 0$ corresponds to the lack of market unraveling ($M = 0$).

If the hypothesis of $M = 0$ could not be rejected, one could construct a straightforward estimate of β. With no missing sales, $I - G = A$, so that $\hat{\beta}$ would solve $F_0(H + \hat{\beta} \cdot T) - F_0(H) = \hat{I} - \hat{G}$.[14]

Previewing our results, however, we find that $\hat{Z} \leq 0$ is rejected or, put differently, we find that the size of the gap is larger than the number of transactions that bunch. We conclude that there are transactions that do not take place because of the proximity to the threshold so that the market (partially) unravels in its neighborhood.[15]

One can bound local exit from the market (M) by considering how much missing mass is required to explain our estimates assuming different values of β. In particular, consider the two extreme cases of $\beta = 0$ (buyer captures all surplus) and $\beta = 1$ (seller captures all surplus). In the first case, $A = 0$, while in the second case A corresponds to the mass in the interval of prices $(H, H + T)$. Noting that $\hat{Z}$ is expressed in dollar terms, the dollar-valued mass in the second case is, thus, T. Hence, the implied missing mass when $\hat{Z} > 0$ is between $\hat{Z}$ (when $\beta = 0$) and $\hat{Z} + T$ (when $\beta = 1$). In our discussion of the results, we will refer to the lower bound$\hat{Z}$.

Finally, note that this discussion provides three qualifications of interest more generally when relying on notches and kinks in tax schedules. First, for a clean interpretation of our incidence parameter $\hat{h}$ (and, analogously, for estimating elasticities or other measures of behavioral response based on bunching), the counterfactual distribution F_0 needs to correspond to the situation absent the discontinuity—this is a strong assumption that is violated if there are spillover effects from the notch/kink to the non-taxable region. However, our estimates of bunching $\hat{I}$ and missing mass $\hat{Z}$ do not require such an assumption. Second, the presence of a notch may provide incentives to exit, corresponding to local unraveling of the market. In this case, gap estimates partially reflect such an exit and can be used to test for its presence when combined with the magnitude of the shift to the notch. Third, as the consequence, when there is such market unraveling, extensive responses estimated in the neighborhood of a notch/kink do not generalize to extensive response elsewhere.

4.4 Implications for Efficiency of the Equilibrium Allocation

In order to shed a light on how taxation interacts with search frictions away from the threshold in this market, we proceed as follows. Recall that the equilibrium price for a given match (b, s) is equal to $\beta b + (1 - \beta)s - \beta T$. Conditional on the seller's type, $\text{var}[p|s] = \beta^2 \text{var}[b|s]$. If we could directly observe s, the comparison of the variance of prices conditional on s with and without the

behavioral response, $\frac{\hat{G}}{\hat{I}}$.

[14] We treat T as the lump-sum tax here for simplicity of exposition; the effect of adjusting for the marginal tax of 1% is negligible for the purpose of this exercise.

[15] Naturally, unraveling occurs here because the tax reduces incentives to transact, but in other contexts the incentives may go the other way. Studying a time-notch affecting marriages in Sweden, Persson (2013) compares bunching at the notch and the gap above the notch and finds that in that context discontinuous incentives may encourage transactions at the (time) notch.

tax would constitute a test of the hypothesis that taxation affects price dispersion. Evidence of this kind would suggest that the tax increases deviation from efficiency.

In practice, we are unable to observe the seller's type and instead rely on a set of indicators, X, that proxy for it. In that case,

$$\text{var}[p|X] = \beta^2\text{var}[b|X] + (1-\beta)^2\text{var}[s|X] + 2\beta(1-\beta)\text{cov}[b, s|X]$$

When X contains s, the second and third term are zero. In order to shed a light on how $\beta^2\text{var}[b|s]$ varies with and without taxation, we consider expanding the set of indicators X—as they become more informative about s, the influence of the last two terms declines and the first term should tend towards $\text{var}[b|s]$. We test whether there is a difference between $\text{var}[p|X]$ with and without taxes for a large set of indicators X correlated with seller's type.

One of the primary indicators of seller's type that we consider is the seller's asking price. It is natural to think that this price is correlated with the seller's type, but it may also be endogenous to taxation. In that case, the alternative interpretation of the effect of taxation on $\text{var}[p|X]$ is as a test of whether taxation changes informativeness of this important signal available to buyers.

4.5 Econometric Implementation

In our primary approach, we estimate the price distribution of sales by maximum likelihood as follows. We specify a parametric distribution of prices absent the tax:

$$\ln f_0(p) = g(p) + \alpha D(p) \tag{2}$$

and the distribution in the presence of the tax:

$$\ln f(p) = \ln f_0(p) + \gamma \cdot I(p > H) \tag{3}$$

where the left-hand side is the log of the probability distribution function at price p, $g()$ is a parametric function (a polynomial—third degree in our baseline specification) and D is a set of controls for round numbers. We allow for discontinuity of the density at the threshold (γ) to account for the shift in sale price and global extensive-margin response (gray region Z of Figures 5 and 6) among transactions subject to the tax. We estimate this model using data that excludes some region around the threshold $(H - \underline{\text{P}}, H + \bar{P})$ to ensure that our estimates are not biased by the distortions to the distribution near the notch.[16] Given the observed distribution of prices outside of the omitted region, we estimate the distribution given by equation (3) by maximum likelihood.[17] This procedure yields $\hat{f}_0(p)$, our estimate of the counterfactual distribution function of prices absent

[16] Our theoretical framework actually establishes that there is a value of $\bar{P}$ above which the threshold (though not the tax) is irrelevant for the distribution.

[17] Formula 3 is already the log-likelihood and implementation only requires imposing conditions guaranteeing that $f(p)$ is a probability distribution function, i.e., that it integrates to one over the considered interval.

the tax and $\hat{f}_1 = e^{\hat{\gamma}}\hat{f}_0$, the counterfactual in the presence of the tax.[18] Using these counterfactual distributions, we estimate bunching and gap, local incidence, and bounds on attrition near the threshold as outlined above. Specifically, we estimate the excess mass as the difference between the observed mass in the region $(H - \underline{P}, H)$, and the predicted mass $(\int_{H-\bar{P}}^{H} \hat{f}_0(p)dp)$ and the gap as the difference between the predicted mass in the region $(H, H + P)$ allowing for a discontinuity at H $(\int_{H}^{H+\bar{P}} \hat{f}_1(p)dp)$ and the observed mass in $(H, H + \bar{P})$.

While it is common in existing literature exploiting kinks and notches for identification to include round-number dummies to control for bunching at these points, our implementation of the round-number effects, $D(p)$, is more involved. Our baseline approach is to rely on the maximum likelihood estimation and hence specify the density at any point. In order to parsimoniously capture various forms of bunching, we introduce "bunching" regions for each round number R that extend from $R-b$ to R. Within the bunching regions, the distribution is specified as $g(p)+D_R+D_R \cdot p$ where D_R are the relevant round-number dummies, while it is $g(p)$ otherwise. In practice, we take $b = \$1000$—this allows the bunching region to extend from, for example, \$899,000 to \$900,000 allowing both for bunching at the \$900K level and just below it (which is of relevance when we turn to listings data). Except for a gain in information by allowing for continuous prices, this approach is very close to binning the data in \$1000 bins (and we show the more restrictive "binned" specification as one of our robustness checks). Beyond that, using maximum likelihood instead of fitting a polynomial to binned data replaces the arbitrary zero-mean restriction for the error terms by a natural restriction that the estimated specification represents a distribution function—arguably, a much more natural assumption than OLS.

We allow for round-number bunching at multiples of \$25,000 where rounding is most pronounced in practice. Because the extent of bunching at round numbers may vary depending on the price (perhaps \$1.2 million is not equally as salient as \$600,000), we interact our round-number effects with the price. Furthermore, we allow for separate effects for multiples of \$25,000 and \$50,000. Since our objective is to estimate the counterfactual in the omitted region (in particular, at \$1 million), this approach amounts to assuming that bunching at other round prices is a valid counterfactual for the magnitude of bunching at the tax notch—this is not a directly testable assumption but, as discussed before, data in New Jersey before the introduction of the mansion tax (Figure 3) provides support for this assumption.

All reported standard errors are obtained by bootstrapping the whole procedure 999 times. Note that the estimates of incidence and gap may fall into one of the round-number bunching regions. When this is the case, the estimates are not very sensitive—small changes in parameters correspond to staying in the bunching region. In such cases reported standard errors for estimates of gap and incidence are small, even though standard errors for parameters of the parametric density are not.

[18] We show as a robustness check an estimate of f_0 that only relies on the data to the left of the threshold. Estimates of f_1 using only data on the right involve projecting far out of sample and are sensitive to the specification choice.

5 Distortion to the Price Distribution

We begin by demonstrating that the tax has a causal effect on the distribution of prices and timing of transactions. Response to the tax notch is evident in Figure 1, which shows the empirical distribution of taxable sales in New York with sales grouped into $5,000 bins. There is clear bunching in the sale price just below $1 million and a drop in the volume of sales just above $1 million. Figure A.1 in the Appendix shows analogous patterns at the smaller (0.425%) RPTT notch at $500,000, which also demonstrates some evidence of a response.[19]

We can also verify explicitly that the tax induces bunching by comparing sales in New Jersey before and after the introduction of the tax. Figure 3 presents plots of the (log-scaled) histogram of sales in NJ before and after the tax is introduced with $25,000 bins, with pre-tax distribution adjusted to account for sales growth and inflation, as discussed in the Figure note (and corresponding to our empirical implementation).[20] We see pronounced bunching after the tax is introduced in 2004 and no indication of bunching prior to 2004. Data prior to 2004 also shows that $1 million is no more salient than other multiples of $50,000 before the tax arrives. Moreover, this figure displays clear visual evidence of the gap above the notch in the post-tax distribution, a feature that is not shared with the pre-tax distribution.

As Figure 4 illustrates, the number of sales in NJ just below the threshold clearly increases precisely at the time of the introduction of the tax. Correspondingly, the number of sales above the threshold falls. Focusing on the region within $10,000 of the threshold, it is evident that the increase in the mass below the threshold is larger than the shift from just above $1 million. This difference provides the first clear indication that the local effect of the tax may extend beyond 100% of its value ($10,000). Figure 4 in the Appendix demonstrates that retiming is strong for sales well above the threshold, but that it is unlikely to have lasted for an extended period of time—the excess sales just before the introduction of the tax do not correspond to more than a couple months worth of sales. On the other hand, the pattern of sales in the combined $900,000–1,100,000 range suggests that despite pricing effects post-introduction of the tax, overall retiming of transactions in the neighborhood of the threshold is not particularly strong.[21]

[19] There is also significant bunching at other round price levels (at every $50,000 and, to a lesser extent, remaining $25,000 multiples), which may confound our bunching analysis. Unlike the bunching at $1 million, this round-number bunching occurs in the bin above rather than the one below the round number. A priori it is possible that, although this observed bunching below $1 million is consistent with theoretical predictions, it may simply reflect adjustments to the tax by very small amounts. However, aggregating the data to larger bins in Figure A.2 mostly eliminates such round-number bunching, while continuing to indicate that the response covers more than just the immediate neighborhood of the threshold.

[20] We choose this larger bin size to smooth out the bunching at other multiples of $50,000. The conclusions are the same using $5,000 bins—in the presence of the tax there is excess bunching just below $1,000,000 and a gap just above that is not present when there is no tax. However, the bunching at multiples of $50,000 makes the figure difficult to read when both distributions (pre and post tax) are overlaid.

[21] Appendix Figure A.4, which shows the monthly distribution of sales in New Jersey between $900,000 to $1,100,000 further demonstrates that that responses extend well beyond the $10,000 value of the tax. There is no evidence of the distribution being distorted below $975,000, and clear evidence of a shift of the mass of sales to just under $1 million from as far above as the $1,025,000 to $1,050,000 range. Additionally, Figures 4 and A.4 show patterns that may be consistent with anticipation effects—there is a spike in sales with prices over $1 million just before the introduction of the tax. This is not surprising: the tax had been announced prior to coming into effect, and the lengthy process

Table 5 contains our baseline estimates of how the transfer tax distorts the distribution of housing sales. The first row reports estimates for New York City and corresponds to the specification shown on Figure 2. For our baseline, we use a 3rd order polynomial, while omitting data in the \$990,000-\$1,155,422 region (the upper bound is the theoretical limit discussed in Section 4.1), and allowing for an additional constant shift above the tax threshold. Our estimate of the local incidence parameter in the baseline specification is \$21,542.098: bunching at the threshold is equivalent to all transactions over the following \$21,000 shifting to the threshold. These estimates are consistent with impressions from the graphical evidence presented above. Taken literally as an incidence estimate this corresponds to over 200% incidence of the tax on sellers for the marginal transaction. The fourth column presents the corresponding estimates of $\hat{Z}$: the positive value indicates that there has to be substantial extensive margin response due to the presence of the threshold or, in other words, there is no β that can rationalize behavior if no such response is present.

In Appendix C we discuss robustness of these results to reasonable modifications of our specification, including choice of polynomials, changes in the omitted region and estimating incidence by OLS relying on binned data. We find that our baseline local incidence results are very robust. We also consider "placebo" treatments at other round numbers that, as expected, show no effect.

Our estimate of about 200% reduced-form incidence local to the threshold is consistent across geographies and data sets. In the top panel of Table 5, we report results for New York State (excluding New York City) and New Jersey. Estimates for these regions are remarkably similar—within \$2,000—to those for New York City. In contrast, estimates for New Jersey before the introduction of the tax show no evidence of bunching.

As an alternative, we estimate incidence in New Jersey using the pre-tax period as a counterfactual for the post-tax period and find similar results. We implement this pre/post comparison as follows. We omit transactions within 90 days of the policy change (to avoid the retiming response) and focus on the following year (Oct. 30, 2004–Oct. 29, 2005). We rescale the period before the tax (May 3, 2003 to May 2, 2004) to account for sales growth over time. Specifically, we construct a counterfactual growth factor by taking the ratio of the count of sales within \$2500 of each price from May 3, 2002 to May 2, 2003 to the count of sales from Nov. 5, 2000 through Nov. 4, 2001 (omitting sales between Nov. 2001 and May 2002 to mimic the 180 day gap around the introduction of the tax in August, 2004). Figure 3 shows the corresponding distributions. We find excess sales at the mansion tax threshold as the difference between total post-period sales and adjusted pre-period sales in the region \$990,000–\$999,999, and estimate the incidence as the price, p^*, at which the number of sales in the pre-period between the threshold and p^* is equal to the excess. We estimate the missing mass in the gap in the same way by taking the difference between total sales in the range \$1,000,000–\$1,155,422 in the pre and post periods. We find standard errors for these estimates by bootstrapping this procedure (including the growth factor for the pre-period) 999 times. Our incidence estimate using the pre/post comparison is slightly larger than the baseline

of closing a real estate transaction may allow for the possibility to speed up the timing of final sale. This effect is not long lasting and of no relevance for New York where the tax was introduced long before our data starts.

estimate ($25,000 vs. $21,542).[22]

Note the similarity of our estimate of $\hat{h}$ using cross-sectional data (our baseline specification) and the pre/post comparison in New Jersey. As we discussed above, these two sources of identification make different assumptions about the counterfactual. In the cross-sectional case, the counterfactual distribution is potentially distorted due to general equilibrium effects. In the pre/post comparison, the counterfactual is the distribution absent taxation. Similarity of the estimates of $\hat{h}$ suggests that general equilibrium spillover to the part of the distribution below the notch are not very important.

We find some heterogeneity by property vintage (years since construction) in estimates shown in Table 5, suggesting that some of the local response to the tax may be due to supply-side quality adjustments. We expect that negotiating a purchase of property before construction is finished allows for significant response in terms of the level of finish, appliances and other amenities, allowing for price reductions driven by adjustments in property characteristics. Similarly, older properties may require renovation and hence allow for quality to more readily respond to the tax. In contrast, original sales of apartments or houses after they have already been constructed and finished may have less flexibility. Our data for New York City and New Jersey contain information about year of construction of the property. In particular, in New York City, which is dominated by large apartment buildings, there is a non-trivial number of sales that occur before construction is finished. For New York (but not New Jersey), we find that bunching is very large for sales before construction is complete and for sales that occur three or more years after construction. In contrast, sales that occur soon after construction—presumably original sales of already fully constructed and equipped properties—show smaller, but still significant (exceeding $10,000 incidence estimate), bunching. Recall though that the introduction of the tax in New Jersey induced bunching immediately (Figures 4 and A.4), so that investment-related responses are unlikely to explain the bulk of the response. We interpret these results as evidence that supply-side response along the quality/finish dimension is important: quality adjustments may perhaps explain around half of the price shift, still leaving local incidence of over 100%. While we do not pursue full welfare analysis, note also that such a tax-motivated response corresponds to welfare loss by the same logic that implies that taxable income response reflects the efficiency cost of income tax (Saez et al., 2012).

In the first and second panels of Table 6, we report estimates for the other taxes. Our estimates of response to the NYC RPTT and NJ RTF are consistent with the mansion tax estimates. We find no evidence of response to the small ($600) New Jersey RTF threshold. For the $2125 RPTT notch that applies only in New York City, we find a response for new sales, but not for old sales (and we find no effect in the rest of New York State where the tax does not apply). The variability of these results coincides with shifts in statutory incidence. Like the mansion tax, the RPTT on new sales is

[22] We find little change in local incidence over time. Figure A.5 in the appendix displays monthly incidence estimates for NJ. Prior to the introduction of the mansion tax in August 2004 estimates show no response to the threshold. Once the tax is introduced, prices quickly respond—incidence estimates reach the $20,000 level within four months. We find no evidence that the response to the tax is changing with the housing boom and bust. We also see no obvious relationship between incidence and the real estate market in New York. Table A.3 presents incidence estimates for all three geographies over time. While the estimates vary somewhat from year to year and region to region, we see no clear pattern over time and they all hover around our baseline estimates of $20,000.

the responsibility of the buyer and we find evidence of a response that is consistent with the mansion tax estimates, albeit somewhat smaller: \$1758.225 represents an 82.7% local incidence on sellers of the \$2125 increase in tax liability. The RPTT on old sales and the NJ RTF schedule are the responsibility of the seller, and in none of these cases do we find any evidence of response. Both of these results are consistent with our reduced-form incidence estimates based on the mansion tax: a 100% burden of the tax on sellers should correspond to no price change for the sellers when they are the party with statutory incidence. Alternatively, switching incidence can correspond to changes in the salience of the tax—perhaps a tax imposed on sellers is less salient than the one imposed on the buyers.[23] One concern is that the response to the tax may be driven by tax evasion, which we investigate by examining proxies for the availability of evasion opportunities. In the third panel of Table 6, we split the NYC sample by coop vs. condo status. Coop transactions have to be approved by coop boards that have the power to veto sales. In particular, there is anecdotal evidence that coop boards disapprove transactions that occur below the expected (or, perhaps, desired) market price. One might expect that if underreporting of the price is the important margin of response to the tax, the extent of bunching in coop apartments should be smaller than otherwise. This is indeed what we find, although the margin is small: \$16,292.354 for coops vs. \$23,292.602 for non-coops.

We also investigate a more direct proxy for evasion—the nature of the transaction—presented at the bottom of Table 6. All-cash transactions involve fewer parties (in particular, no financing) and more liquidity, which may increase the likelihood of side payments. We find the opposite: incidence of \$16,018 for cash transactions versus \$20,676.666 when the sale has an associated mortgage. In the context of real estate transactions, tax evasion is certainly possible, but one might expect that it is not completely straightforward: both parties have to agree and money has to change hands at some point during the long closing process. Evasion in this context likely requires an aspect of trust between the two parties. Our New York State data contains a dummy for whether a transaction is "arms-length" (i.e.., between related parties). We find no evidence that arms-length transactions involve more bunching (\$23,169 for arms-length sales, \$23,786 for non-arms-length).

Finally, examining listings data for Manhattan, we find comparable bunching in seller listing prices, suggesting that evasion is not a driving force of the observed sale-price response. Figure 7 shows the smoothed distribution of listing prices around the mansion tax threshold for properties sold and matched to the tax data.[24] There are three prices shown for sold listings: the initial

[23]For both the RTF and the RPTT there is a small dominated region where sellers would be better off accepting a lower sale price below the threshold, which is offset by a lower tax bill. Thus, there should be a small amount of bunching and a small gap at the RTF and RPTT thresholds when sellers remit. That we find no evidence of this is potential support of these taxes' limited salience. The dollar value of these discontinuities is small, however, so we refrain from drawing sharp conclusions on this point.

[24]Since bunching is more prominant in listings data, we adjust the distributions on the graphs to remove the common round number bunching. Specifically, we regress the log of the per-\$25,000-bin count on a cubic in price and dummy variables for multiples of \$50k and \$100k interacted with the price. We then subtract the predicted bunching effect from the actual counts. The remaining peaks in the data are the result of noise and do not necessarily correspond with salient round numbers. Figure A.6 and Figure A.7 in the appendix show the unadjusted distribution and the distribution of listing prices for all listed properties in Manhattan, respectively.

asking price, the final price in the listing data, and the sale price. Among properties that sell, bunching appears most prominent for the final asking price, followed by the sale price, and the initial asking price. These visual perceptions are confirmed by our estimates in Table 7 that find substantial bunching for both initial and final listing prices that actually exceed the response at the sale stage.[25] The response of the listing prices indicates that sellers internalize the presence of the tax (which is the responsibility of the buyer) even before meeting the buyer. Since these listings responses occur before the seller identifies a buyer who would be willing to engage in tax evasion, we find it unlikely that the ultimate sale price response is driven by such cheating.

The evidence so far shows clearly that transaction taxes distort the distribution of sale prices. We find some evidence of supply-side response in quality adjustments, as well as differences in estimates based on the side of the market responsible for the tax. Our tests of tax evasion are weak, but do not suggest that this is the main force. These results reflect local reduced-form incidence estimates—the adjustment of prices in response to the threshold. By themselves, they do not reveal the strength of bargaining power and are not informative about the incidence of the tax away from the threshold. As discussed in our theoretical section, understanding the bargaining power and, relatedly, the possibility of unraveling in the market requires investigating the size of the gap in the distribution as well (measured by $\hat{Z}$).

6 Unraveling: Market Distortions Local to the Threshold

Our estimates of $\hat{Z}$ imply that there is significant unraveling of the market local to the tax threshold. In general, we find the number of sales bunching at the threshold to be smaller than the number of sales missing from the gap, translating into a positive value of $\hat{Z}$, as reported in Table 5. As discussed in Section 4.3, a positive sign on $\hat{Z}$ cannot be reconciled with positive values of β and instead indicates unraveling of the market in the neighborhood of the threshold. Hence, our results show that the threshold design of the tax discourages transactions that would have taken place if the rate was the same but discontinuity was not present—even after controlling for the usual extensive margin response (sales with positive surplus in absence of the tax, but negative surplus when taxed) we find that there are sales that do not occur. Moreover, the presence of unravelling suggests that the matching process is an important part of real estate sales, and this process may be disrupted by the tax.

To be specific, our baseline estimate of $\hat{Z} = \$43,861.766$ implies that the tax eliminates transactions that correspond to a range of original prices with the length of between $\$43,861.766$ and $\$43,861.766 + \text{tax} \approx \$54,000$. This corresponds to about 2,800 sales over the 2003-2011 period. Altogether, we find that the tax induces a \$20,000 price-range worth of transactions from above the threshold to bunch at the threshold, and discourages another $\$50,000$ or so of sales from occurring at prices just above the threshold. The finding that $\hat{Z}$ is substantial and positive is robust. Firstly,

[25] We do not find conclusive evidence of a similar response of listings prices to the NYC RPTT. This is consistent with the results that we discussed before: we find a response to the RPTT only for new sales where the tax applies to buyers. However, the number of new sales in the listing data is very small and we run into power issues.

as can be seen in Table 5, $\hat{Z}$ ranges from 37,409.87 and 43,861.766 across the three geographic regions that we consider (NYC, NYS, and NJ). Secondly, $\hat{Z}$ is small and economically insignificant in NJ before the tax is introduced. Thirdly, as discussed in Appendix C, the estimates of $\hat{Z}$ are robust to the specification choices. Moreover, using data for NJ and constructing our counterfactual distribution using data from before the tax is introduced gives an estimate of $\hat{Z}$ somewhat smaller than the baseline maximum likelihood estimate and with large standard errors but with the point estimates that are still economically significant.

This extensive-margin response highlights an important margin of efficiency loss due to the transaction tax notch. As discussed above, a positive estimate of $\hat{Z}$ suggests a very specific extensive-margin response, which we refer to as unraveling: some buyer-seller pairs who have a positive joint surplus under the tax (regions B and C of Figure 6) are exiting the market. This does not imply that these parties do not trade at all—buyers and sellers can continue to search and some of them may ultimately transact at different prices away from the bunching/gap region—but it provides evidence that the market in the region just above the threshold is unraveling. Note that this is different from the usual extensive margin response in which buyer-seller pairs who would transact in absence of the tax find that the tax reduces their joint surplus below zero and so the sale does not occur. Our estimation procedure explicitly controls for this traditional extensive-margin response by allowing for a level shift (discontinuity) in the distribution above the notch.

Examining real-estate listings data for New York City, we find suggestive evidence that the further tax disrupts the search process in the neighborhood of the notch. We interpret the presence of substantial bunching in the listings price (discussed above) as evidence that the tax influences seller search behavior. We also find that those who list just above the notch (between \$1M and \$1.075M) are still very likely to sell below one million (see the relationship between listing and sale price in Figure A.8 in the Appendix), and are more likely to sell than those who list below the notch or much higher above the notch (Figure A.9)—despite spending more time on the market (Figure A.10). Interestingly, those who list just above the notch are more likely to leave their original REBNY realtor and sell with another realtor or on their own (Figure A.11), suggesting an additional margin of adjustment to the tax—saving on realtor fees to compensate for a lower sale price.

The evidence thus far points to extensive distortions to price, unraveling of the market, and some disruption of the matching process local to the transfer tax notch. However, many of these responses occur because of how the tax is implemented—prices can adjust below the notch to avoid the tax and the notch creates specific local incentives for buyers and sellers to break matches. In what follows, we examine how the tax affects home sales more generally, away from the notch.

7 Global Market Distortions

In this section, we show evidence indicating that the transfer tax may distort the matching process everywhere above the notch. Conditional on initial listing price, sellers are taking larger discounts in the presence of the tax. While this higher discount could be explained by a shift in bargaining power or by endogenous listing prices, we find evidence that the efficiency of matches themselves is distorted by the tax. In particular, the variance of sale price conditional on property characteristics increases with the tax. While these results are largely descriptive—we must rely on observations below the threshold to form a counterfactual above—they do show sharp effects and, as discussed in Section 4.4, we interpret this higher variance in selling price as a decrease in the efficiency of the matching process.

Price discounts, which we define as the percent drop from listing to sale price, increase under the transfer tax. Figure 8 shows that the discount from the initial price to the final advertised price (i.e., before a buyer is identified) and to the final sale price increase as the initial listing price moves above \$1,000,000. We present the median and 75th percentiles of the distribution of discounts (many listings are not revised) in the figures. The effect is not immediate at the \$1,000,000 threshold, because the tax applies to the sale price and not the initial price, and it is the latter that constitutes the running variable here: close to \$1 million a small discount is sufficient to bring the sale price below the notch. However, the increase is persistent well above the threshold. We find analogous evidence for the discount from final to sale price—see Appendix Figure A.12—suggesting that the price response is slowly revealed and reinforced throughout the search process by distorting the initial prices, subsequent revisions and, finally, during the bargaining stage.[26]

We investigate the relationship between transfer taxes and price discounts more formally and find that the increase above the notch is significant and persistent. We regress the discount from initial asking price to sale price on a linear spline in initial asking price, with nodes at every multiple of \$100,000 (restricting the sample to listings with initial prices between \$500,000 and \$1,500,000). We follow the same procedure for the discount from initial asking price to final asking price. We plot the difference between the predicted discount (first price and final price) at each node and the predicted discount at \$1,000,000 in Figure 9. These estimates show a significant jump in the price discounts at the notch that shows no signs of reversing before reaching \$1.5 million.[27]

There are several explanations for the increase in discounts. Firstly, it could be that sale prices above the notch are not changing, but the response is driven entirely by a change in asking prices (and we know from Section 5 respond to the tax). In this case, since the final outcomes would be unchanged there could be little efficiency loss due to the tax. Secondly, it could be that the tax increases buyers' bargaining power. This may not entail any welfare loss due to the tax, but rather a redistribution of surplus from sellers to buyers. These two explanations still reveal that the tax affects behavior, although not the final outcomes. Alternatively, as we argue next, it could be that the tax disrupts the search process and reduces match quality with associated efficiency losses.

We find evidence consistent with transfer taxes disrupting the buyer-seller match process: conditional on seller characteristics, the tax increases the variance of sale prices. We investigate the

[26] This can also be seen in Figure A.13, in which we focus on the mean discounts from the initial price that allow us to decompose the response (but blur the response that is predominantly present at high quantiles). Roughly half of the response is due to price revisions and half due to discounts at the bargaining stage.

[27] The corresponding point estimates (and related slope estimates) are in Appendix Table A.4.

relationship between asking prices and the variance of sale-prices with a two-stage spline estimation procedure. We first regress sale price on a linear spline in asking price (nodes at every $100,000 between $500,000 and $1,500,000). We then estimate by median regression the relationship between the squared residuals from this first stage and a linear spline in asking price. We estimate standard errors using a clustered bootstrap procedure.[28]

We find a significant increase in the variance of sale price as (initial or final) asking price crosses the $1,000,000 notch. In Figure 10 we plot the predicted dispersion at the given node (from the two-stage spline procedure) relative to that at the notch.[29] This increased dispersion is pronounced and persistent well above the notch. The estimates show that, in general, the predicted variance below the notch is very close to the predicted variance at $1,000,000, while the variance of sale price conditional on asking price is significantly higher above the notch than below the notch. Even as asking price rises to $1,500,000, the variance of sale price does not return to pre-notch levels. This is inconsistent with asking prices and discounts simply scaling up without changes to other aspects of the matching process. As with price discount, these estimates suggest that the real estate market is affected by the transfer tax, even far above the notch—in the presence of the tax, asking price is a noisier signal of final sale price.[30]

In Section 4.4 we argued that an increase in the dispersion in sale price conditional on seller characteristics implies a movement further away from the optimal allocation of assortative matching. An increase in the variance of sale price conditional on asking price is suggestive of such an efficiency loss, but it is not conclusive: the ideal measure is the variance of price conditional on seller type. Asking price is an imperfect proxy for seller type—if asking price is endogenous to the tax, the increase in variance may be driven by a change in the composition of sellers who list at each price.

We repeat the two-step estimates of variance by asking price including controls for property characteristics, in order to better approximate the variance of sale price conditional on seller type. In both stages we control for year-of-sale fixed effects, zip code fixed effects, building type (single-family home, multi-family home, apartment in walkup, apartment in elevator building, etc.), whether the sale is of a new unit, and the log of years since construction (plus an indicator for missing years since construction). We also plot the results from this procedure in Figure 10. Including these controls somewhat reduces the effect of the tax on the variance of sale price, suggesting that endogeneity of listing prices may have influenced the previous estimates, but the difference is not large and not even statistically significant except near the threshold. In particular, there is still a significant increase in the dispersion of sale price above the notch. We interpret this increase as evidence that

[28] We use median regression in the second stage, because squaring the residuals makes these specifications sensitive to outliers. The bootstrap is as follows: for each observation, we resample first-stage residuals from the 50 nearest observations (by listing price). We use these residuals and the first-stage predicted values to construct a bootstrap sample and re-estimate the two-step process. We iterate this process 999 to acquire the distribution of estimates.

[29] We present the numerical estimates of the difference between the predicted value at each node and the predicted value at $1,000,000 in Table A.5. The estimates are fairly consistent for both asking prices and are insensitive to the choice of quantile or mean regression in the first stage—see also Figure 10.

[30] This increase in dispersion is also confirmed in the raw data. In Figures A.14 and A.15 we plot the variance of sale price conditional on initial asking price and final asking price, respectively. In both cases, the variance of sale price increases above the notch.

the transfer tax reduces efficiency of the housing market by disrupting the matching process.

8 Conclusions

Our analysis of real-estate transaction tax notches reveals substantial price response local to tax thresholds, the importance of search in the housing market, and that transaction taxes may increase search-related inefficiencies. The price responses that we identify suggest that sellers local to the threshold take large price cuts—greater than the cost of the tax—although this may partially be driven by quality adjustments. The local response also highlights the importance of search in the housing market—not only do we find that sellers alter their search behavior around the notch, but also that productive matches near the notch fail to transact. Finally, we find that in the presence of transaction taxes the relationship between seller asking price and final sale price weakens, and that this persists when we control for property characteristics. This effect extends far above the threshold. We interpret this increase in price dispersion as a movement away from the efficient allocation of positive assortative matching in the housing market and as revealing the distortion due to taxation globally.

Our evidence of exit from markets near tax thresholds raises important issues for implementation and interpretation of studies relying on bunching at notches for identification. To date, most of this literature has assumed one-sided markets and abstracted from extensive-margin responses. Our framework highlights that there are two different types of exit from markets. One is standard—transactions with low surplus do not take place. In the income tax context, this is akin to the labor force participation decision.[31] The second type of exit, which has not been previously recognized in this literature, is unraveling of the market near the notch that corresponds to destroying productive matches. Both of these responses imply that bunching does not fully characterize the consequences of a notch. While in some contexts (e.g., an income tax notch at a high value) it is reasonable to ignore the first type of extensive-margin response, the second type of response is likely to be intrinsic to any matching context in which parties have an option to continue search (e.g., firms may continue searching for a worker willing to accept a wage below the notch; employees may make different occupational choices). The possibility of this second type of response undermines the assumption that the excess mass bunching at the threshold is identical to the missing mass in the gap (as in Kleven and Wassem, 2013). Our empirical framework relaxes this assumption and allows us to explicitly test for exit local to the notch.

More generally, our results suggest that taxes may introduce inefficiencies into other search markets. Labor markets are, perhaps, the most obvious example of a market with non-trivial search costs where matches are subject to taxation. Many labor-market regulations follow notched designs, such as notched income and wage taxation or requirements to provide health insurance or comply with Value Added Tax if the number of employees crosses a given threshold. That firms might face such discrete costs to increasing scale may lead not simply to supply- or demand-

[31] See Marx, 2013 for explicit modeling of this decision to exit the market.

side adjustments, but perhaps to the destruction of equilibrium opportunities that require costly search—in the same way that productive real-estate matches are discouraged near the transaction tax notch. Moreover, our finding that transaction taxes increase search-related inefficiencies well above the threshold suggests that even non-notched policies may lead to less efficient worker-firm matches. Of course, the housing market differs from the labor market in ways that may make the mechanisms we study particularly pronounced. Firstly, the availability of alternative margins of adjustment affects the ease of moving below the tax threshold. Secondly, the housing market is a "spot" market where the current price determines tax treatment; this is not necessarily the case in labor markets, where contracts may be long lasting and as earnings adjust over time they move further away from the notch. Nonetheless, our results underscore the importance of considering how taxation affects search, especially in the context of policies that follow a threshold design.

References

Benjamin, John D., N. Edward Coulson, and Shiawee X. Yang, "Real estate transfer taxes and property values: The Philadelphia story," *The Journal of Real Estate Finance and Economics*, 1993, *7* (2), 151–157.

Bernheim, Douglas B. and Jonathan Meer, "Do Real Estate Brokers Add Value When Listing Services Are Unbundled?," *Economic Inquiry*, 2013. Forthcoming.

Besley, Timothy, Neil Meads, and Paolo Surico, "The Incidence of a Transactions Tax: Evidence from a Stamp Duty Holiday," 2013. London School of Economics, mimeo.

Best, Michael and Henrik Kleven, "Property Transaction Taxes and the Housing Market: Evidence from Notches and Housing Stimulus in the UK," April 2013. London School of Economics, mimeo.

Carrillo, Paul E., "An Empirical Stationary Equilibrium Search Model of the Housing Market," *International Economic Review*, 2012, *53* (1), 203–234.

_ **and Jaren C. Pope**, "Are Homes Hot or Cold Potatoes? The Distribution of Marketing Time in the Housing Market," *Regional Science and Urban Economics*, 2012, *42* (1), 189–197.

Dachis, Ben, Gilles Duranton, and Matthew A. Turner, "The Effects of Land Transfer Taxes on Real Estate Markets: Evidence from a Natural Experiment in Toronto," *Journal of Economic Geography*, 2012, *12* (2), 327–354.

Doyle Jr., Joseph J. and Krislert Samphantharak, "$2.00 Gas! Studying the Effects of a Gas Tax Moratorium," *Journal of Public Economics*, 2008, *92* (3-4), 869–884.

Genesove, David and Christopher Mayer, "Loss Aversion and Seller behavior: Evidence from the Housing Market," *Quarterly Journal of Economics*, 2001, *116* (4), 1233–1260.

_ **and Lu Han**, "Search and Matching in the Market for Existing Homes," *Journal of Urban Economics*, 2012, *72*, 31–45.

Han, Lu and William Strange, "What is the Role of the Asking Price for a House?," 2012. University of Toronto, mimeo.

_ **and** _ , "Bidding Wars for Houses," *Real Estate Economics*, 2013. Forthcoming.

Hastings, Justine and Ebonya Washington, "The First of the Month Effect: Consumer Behavior and Store Responses," *American Economic Journal: Economic Policy*, 2010, *2* (2), 142–162.

Haurin, Donald R., Jessica L. Haurin, Taylor Nadauld, and Anthony Sanders, "List Prices, Sale Prices and Marketing Time: An Application to US Housing Markets," *Real Estate Economics*, 2010, *38*

(4), 659–685.

Jia, Panle and Parag A. Pathak, "The Impact of Commissions on Home Sales in Greater Boston," *American Economic Review Papers and Proceedings*, 2010.

Kleven, Henrik J. and Mazhar Wassem, "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan," *Quarterly Journal of Economics*, May 2013, *128* (2), 669–723.

Levitt, Steven D. and Chad Syverson, "Market distortions when agents are better informed: The value of information in real estate transactions," *Review of Economics and Statistics*, 2008, *90* (4), 599–611.

Marion, Justin and Erich Muehlegger, "Fuel Tax Incidence and Supply Conditions," *Journal of Public Economics*, October 2011, *95* (9-10), 1202–12.

Marx, Benjamin M., "Regulatory Hurdles and Growth of Charitable Organizations: Evidence from a Dynamic Bunching Design," 2013. Columbia University, mimeo.

Merlo, Antonio and Francois Ortalo-Magne, "Bargaining Over Residential Real Estate: Evidence from England," *Journal of Urban Economics*, 2004, *56*, 192–216.

Mishra, Prachi, Arvind Subramanian, and Petia Topalova, "Policies, Enforcement, and Customs Evasion: Evidence from India," *Journal of Public Economics*, October 2008, *92* (10-11), 1907–1925.

Persson, Petra, "Social Insurance and the Marriage Market," January 2013. Columbia University, mimeo.

Rothstein, Jesse, "Is the EITC as Good as an NIT? Conditional Cash Transfers and Tax Incidence," *American Economic Journal: Economic Policy*, 2010, *2* (1), 177–208.

Saez, Emmanuel, "Do Taxpayers Bunch at Kink Points?," *American Economic Journal: Economic Policy*, 2010, *2* (3), 180–212.

_ **, Joel B. Slemrod, and Seth H. Giertz**, "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review," *Journal of Economic Literature*, March 2012, *50* (1), 3–50.

_ **, Manos Matsaganis, and Panos Tsakloglou**, "Earnings Determination and Taxes: Evidence from a Cohort Based Payroll Tax Reform in Greece," *Quarterly Journal of Economics*, February 2011, *127* (1), 493–533.

Slemrod, Joel B., "Optimal Taxation and Optimal Tax Systems," *Journal of Economic Perspectives*, Winter 1990, *4* (1), 157–78.

_ , "Buenas Notches: Lines and Notches in Tax System Design," September 2010. University of Michigan, mimeo.

_ **, Caroline Weber, and Hui Shan**, "The Lock-in Effect of Housing Transfer Taxes: Evidence from a Notched Change in D.C. Policy," 2012. University of Michigan, mimeo.

Van Ommeren, Jos and Michiel Van Leuvensteijn, "New Evidence of the Effect of Transaction Costs on Residential Mobility," *Journal of Regional Science*, 2005, *45* (4), 681–702.

Table 1: Real Estate Transfer Tax Schedules

Tax	Threshold ($)	Rate Below	Rate Above	Jump	Statutory Incidence
Mansion Tax (NY & NJ)	$1,000,000	0%	1%	$10,000	Buyer
RPTT (NYC, Residential)	$500,000	1%	1.425%	$2,125	Seller*
RPTT (NYC, Commercial)	$500,000	1.425%	2.625%	$6,000	Seller
RTF (NJ)	$350,000	0.78%**	0.96%	$630	Seller

Notes: *Buyer remits tax if the sale is of newly developed property (otherwise Seller remits). **NJ RTF schedule features non-linear tax schedule below $350,000, all of which changes when the sale price crosses the notch; 0.78% and 0.96% are simply the marginal rates faced above and below the notch.

Table 2: Sample Statistics for Taxable Sales

	NYC (2003 – 2011)	NYS (2002 – 2010)	NJ (1996 – 2011)
Number of Sales	380629	1172708	1703260
Sales ∈ ($950K, $1.05M)	7932	6242	7556
Median Price	405600	159900	200000
Mean Price	660719	258363	262122

Notes: NYC data is from the Department of Finance Rolling Sales file for 2003–2011 (taxable sales defined as single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos). Data for NYS from Office of Real Property Service deeds records for 2002–2006 and 2008–2010 (taxable defined as all single-parcel residential sales of one-, two-, or three-family homes). Data for NJ from the State Treasury SR1A file for 1996–2011 (taxable defined as any residential sale).

Table 3: Median Price of Taxable Sales Over Time

	NYC		NYS		NJ	
Year	n	Price	n	Price	n	Price
1996		.	.	.	111759	127000
1997	.	.	.	.	115470	130000
1998	.	.	.	.	131485	137500
1999	.	.	.	.	139167	143000
2000	.	.	.	.	136891	151000
2001	.	.	.	.	136733	169000
2002	.	.	163491	132000	145718	197000
2003	47679	293000	167709	149500	148906	235000
2004	53342	340000	175766	165000	159220	270155
2005	52310	395460	175873	184640	155340	315000
2006	47973	445000	152220	170000	127630	327000
2007	48552	480000	5934	287000	108790	321050
2008	40354	475000	117000	162500	86151	288000
2009	31368	420000	110408	155500	83407	257500
2010	27132	463000	104307	160666	80944	250000
2011	31919	456000	.	.	28766	250000

Notes: NYC data is from the Department of Finance Rolling Sales file for 2003–2011 (taxable sales defined as single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos). Data for NYS from the Office of Real Property Services deeds records for 2002–2006 and 2008–2010 (taxable defined as all single-parcel residential sales of one-, two-, or three-family homes). NYS observations in 2007 are from sales made in 2007, but recorded in 2008–2011 and omits sales recorded in 2007. Data for NJ from the State Treasury SR1A file for 1996–2011 (taxable defined as any residential sale).

Table 4: REBNY Listings Sample Statistics

All Listings

	Sold	Matched, but not Closed	Days on Market	Initial Asking Price	Final Asking Price	Discount (First to Final)
Mean	0.671	0.103	197.859	1604547	1602670	0.019
Median	1.000	0.000	110.000	899000	875000	0.000
n	71875	71875	67550	71875	71875	71875

Closed and Matched Listings

	Days on Market	Initial Asking Price	Final Asking Price	Sale Price	Discount (First to Final)	Discount (First to Sale)	Discount (Final to Sale)
Mean	146.107	1384028	1435747	1241209	0.032	0.059	0.028
Median	80.000	825000	799000	784052	0.000	0.043	0.023
n	40680	44320	44320	44193	43506	43309	43309

Notes: Data from the Real Estate Board of New York's listing service; represents all REBNY listings in Manhattan between 2003 and 2010 that are closed or off market. Sold is an indicator equal to one if the final status of the listing is "Closed." Days on the market is calculated as the number of days between the initial active listing and the final status of "in contract" (if the property sells with REBNY) or "permanently off market" (otherwise). Initial asking price is the asking price on the listing when first active; final asking price is the price listed immediately prior to the listing being "in contract" or being taken off the market (if unsold). Sale price is the price reported in the NYC DOF data and is available only for REBNY listings that have a match in the DOF data (sale price of 0 is considered missing). Discount is defined as $1 - \frac{final\ price}{first\ price}$ and is windsorized at the 1st and 99th percentiles. 5,940 listings have invalid listing and off-market dates (missing or obviously misreported), and are omitted from days on market calculations. "Matched, but not closed" is an indicator that a listing has a match in the NYC DOF data, but is never reported as "Closed" by the REBNY agent.

Table 5: Response to Mansion Tax, by Region and Years Since Construction

Sample		Incidence	Std. Error	$\hat{Z}$	Std. Error	n
NYC		21542.098	1150.878	43861.766	4142.953	102493
NYS (excl. NYC)		23227.515	1084.482	41610.588	4334.170	108462
NJ Post Tax		21477.388	1474.300	37409.873	4310.896	111936
NJ Pre Tax		-784.065	38.892	2958.261	285.872	57836
NJ Pre/Post Comparison†		25000.000	8515.132	14223.330	11628.94	2020
NYC (Yrs. Since Contr.)	<0	37329.701	16009.929	-13709.572	42837.004	559
	0	13759.258	7671.441	-2451.076	27109.852	1048
	1	11309.339	3457.124	45053.550	15619.853	3422
	2	14118.294	3145.743	47311.896	15259.170	4388
	3	24467.069	5927.603	36654.826	26001.895	2253
	4–6	25586.634	6045.138	82880.894	29444.354	2433
	7+	22780.508	1275.592	48877.574	5193.370	72128
NYS (excl. NYC)	Old	22677.619	1283.443	40945.774	4264.702	104576
	New	34254.287	8350.604	62173.984	25965.831	3886
NJ Post Tax (Yrs. Since Constr.)	0	24949.365	10542.512	19598.806	33121.546	988
	1	24684.039	7485.882	55501.171	26732.423	1896
	2	23730.361	7183.484	54344.928	25413.421	1773
	3	24950.530	8043.393	25353.913	25917.361	1882
	4–6	19718.802	4392.694	41103.750	17066.774	6148
	7+	19967.636	1633.977	40047.784	4989.520	85551
NJ Pre Tax (Yrs. Since Constr.)	0	-351.083	4149.632	1619.879	26316.982	723
	1	-142.337	6362.861	4103.325	71219.647	852
	2	-798.644	1743.317	3214.845	13783.929	1082
	3	-2121.267	2519.545	8725.490	10998.073	1249
	4–6	-661.898	537.723	9250.812	18873.573	3335
	7+	-778.278	46.481	3226.236	380.842	36530

Notes: Estimates from baseline procedure (3rd-order polynomial, omit $990k–$1,155,422). NYC data is from the Department of Finance Rolling Sales file for 2003–2011 (taxable sales defined as single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos). Data for NYS from the Office of Real Property Services deeds records for 2002–2006 and 2008–2010 (taxable defined as all single-parcel residential sales of one-, two-, or three-family homes). Data for NJ from the State Treasury SR1A file for 1996–2011 (taxable defined as any residential sale). NJ sample restricted to sales recorded between 1996 and 2003 for pre-tax estimates and after August, 2004 for estimates in presence of the tax. Years since construction defined as the difference between the year the property was built and the year of sale. New sales in New York state are defined as any being flagged as new construction.

†Estimates for NJ Pre/Post comparison using NJ taxable sales omitting 90 days around the implementation of the policy: from Oct. 30, 2004 to Oct. 29, 2005 (post-period) and May 3, 2003 to May 2, 2004 (pre-period). Incidence estimate is the price at which the number of sales in the pre-period to the right of the threshold equal the difference between the number of sales in the bunching region ($990,000–$999,999) in the post- and adjusted pre-periods—pre-period distribution is adjusted as described in text. $\hat{Z}$ calculated as in the text, where the excess number of sales bunching at the gap is the difference between the post- and adjusted pre-period distributions in the bunching region, while the gap is calculated as the difference between the distributions above the notch ($1M–$1,155,422). Number of observations listed for the pre/post comparison is the total count of taxable sales between $990,000 and $1,155,422 for these dates.

Table 6: Heterogeneity in Response by Notch and Sub-Sample

Geography	Sample	Incidence	Std. Error	$\hat{Z}$	Std. Error	n
			RPTT			
NYC	New	1758.225	751.923	3071.659	1975.290	21683
	Old	-390.461	28.339	488.460	69.802	259840
NYS	New	-679.032	510.129	7.829	332.599	19147
	Old	-889.542	18.945	1105.190	56.531	687807
			RTF			
NJ	Post Aug. 2004	-699.552	15.333	-268.101	44.131	546882
	Pre Aug. 2004	-591.400	20.099	-300.348	68.344	836832
			Mansion Tax			
NYC	All Coops	16292.354	2184.555	58245.269	8590.430	26950
	All Non-Coops	23292.602	1027.253	38817.535	4612.629	75543
	Old Non-Coops	24196.113	535.699	37340.124	5067.660	63971
NYC Deeds	Cash Only	16018.361	2356.852	74976.653	8814.378	28339
	Mortgage	20676.666	2223.237	119838.215	11273.628	49421
NYS	Arms-Length	23168.557	1163.640	43274.495	4606.535	97936
	Non-Arms-Length	23786.484	3149.902	24496.898	12693.600	10526

Notes: RPTT and RTF estimates based on 5th-order polynomial, omitting sales between $490,000 and $550,000 (for NYC and NYS) and between $340,000 and $400,000 (for NJ), Mansion tax estimates as in the baseline specification. NYC data is from the Department of Finance Rolling Sales file for 2003–2011. Data for NYS from the Office of Real Property Services deeds records for 2002–2006 and 2008–2010 (excluding NYC) restricted to all single-parcel residential sales of one-, two-, or three-family homes. Sales in NYC are defined as single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos. New sales are defined as any sale occurring within three years of unit's construction (in NYC) or any sale flagged as new construction (in NYS, excluding NYC). Data for NJ from the State Treasury SR1A file for 1996–2011 (taxable defined as any residential sale). Coops are identified in the rolling sales data as sales with associated building codes equal to "Coops - Walkup Apartments" or "Coops - Elevator Apartments."NYC Deeds Records data from deeds records collected by private data provider (taxable defined as any residential sale). Non-arms-length sales in NYS defined by the Office of Real Property Services as a sale of real property between relatives or former relatives, related companies or partners in business, where one of the buyers is also a seller, or "other unusual factors affecting sale price" (ex. divorce or bankruptcy).

Table 7: Mansion Tax: Listings

Sample	Price	Incidence	Std. Error	$\hat{Z}$	Std. Error	n
All	First	24443.666	166.150	101702.509	10589.033	36232
	Final	34992.894	3380.599	76606.311	9122.157	35714
Sold	First	24363.369	353.901	95113.750	12794.903	25112
	Final	38445.445	4289.510	68828.242	10640.681	24755
	Sale	19148.096	2089.630	53521.404	8077.065	24474
Unsold	First	24700.339	2638.516	113962.040	21381.710	7612
	Final	32926.816	6670.229	94839.590	19962.747	7539

Notes: Data from the Real Estate Board of New York's listing service; represents all REBNY listings between 2003 and 2010 that are closed or off market. Unsold sample defined as all listings with final status not equal to "closed." Sold sample defined as all listings that match to a NYC Department of Finance sale record with final status equal to "closed." First price is the initial price posted on the listing. Final price is the last price posted while the listing is active (prior to status being changed to "in contract" or "off market"). Sale price is the recorded price from the NYC Department of Finance.

Figure 1: Distribution of Taxable Sales in New York State

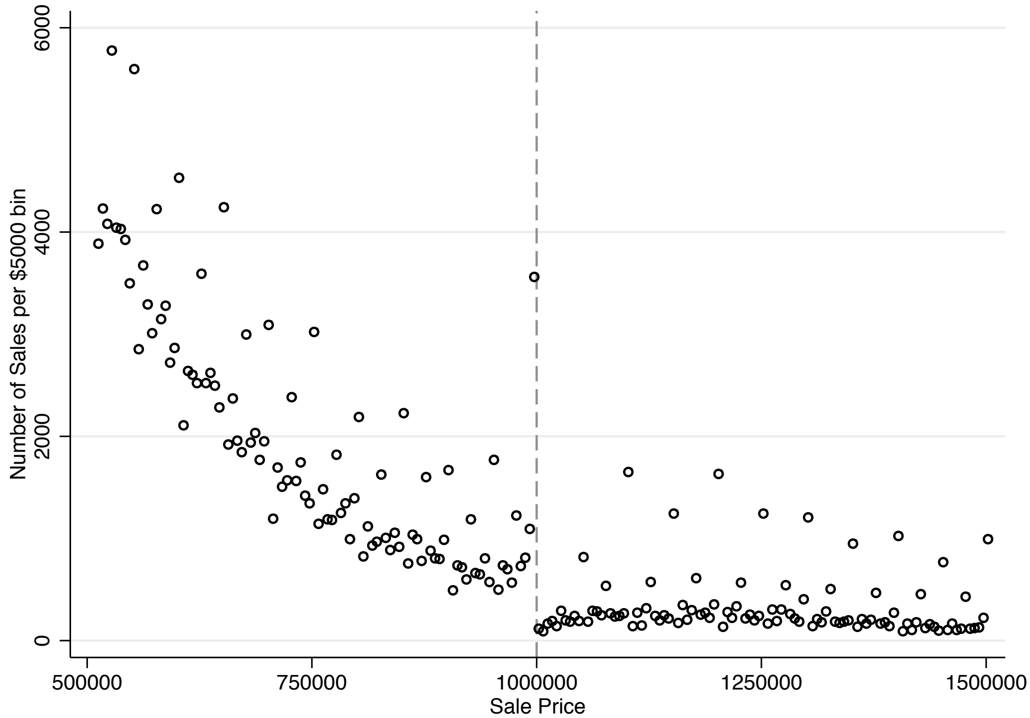

Notes: Plot of the number of mansion-tax eligible sales in each $5,000 price bin between $510,000 and $1,500,000. Data from the NYC Rolling Sales file for 2003–2011 (taxable sales defined as single-unit non-commercial sales of one-, two-, or three-family properties) and from the N.Y. State Office of Real Property Service deeds records for 2002–2006 and 2008–2010 (taxable defined as single-parcel residential sales of one-, two-, or three-family homes).

Figure 2: Distribution of Taxable Sales in New York City

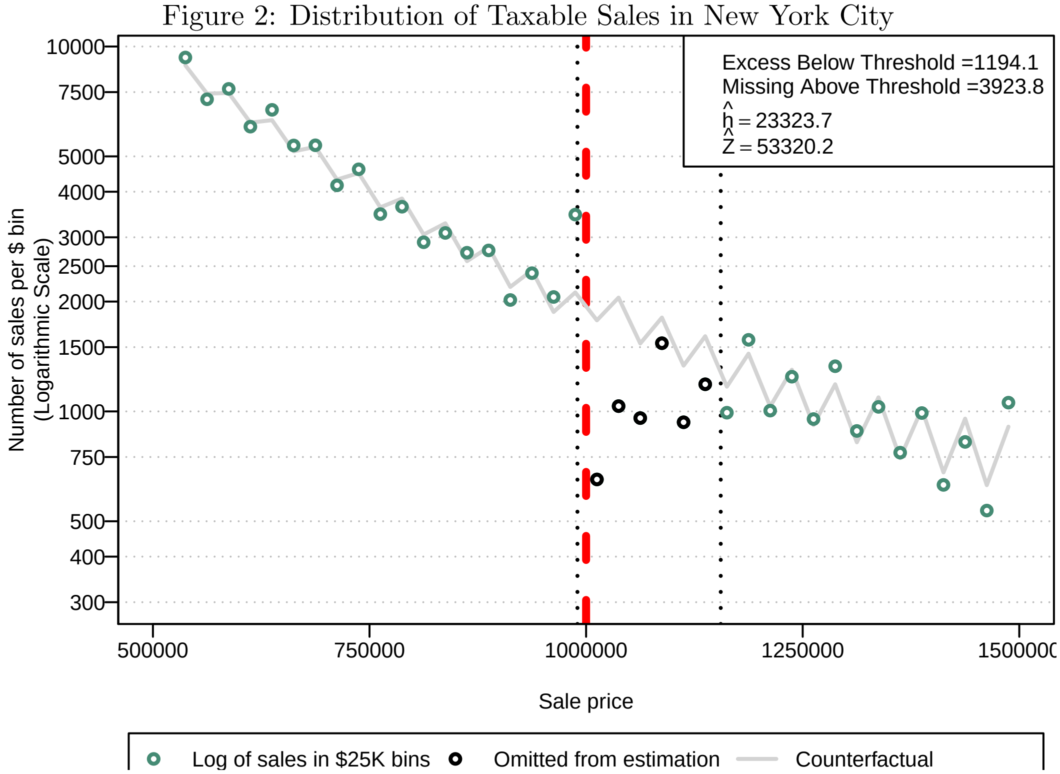

Notes: Plot of the number of mansion-tax eligible sales in each $25,000 price bin between $510,000 and $1,500,000. Data from the NYC Rolling Sales file for 2003–2011. Fit corresponds to the baseline specification in Table 5.

Figure 3: Distribution of NJ Sales Pre- and Post-Mansion Tax

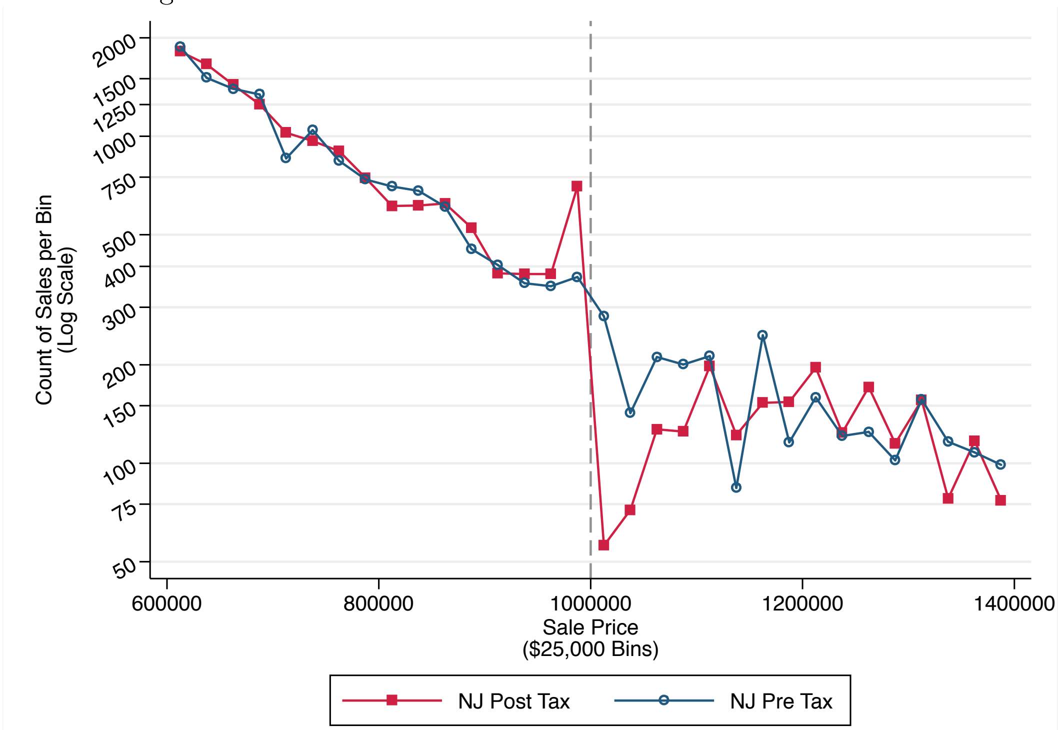

Notes: Plot of the number of mansion-tax eligible sales in each $25,000 price bin between $510,000 and $1,500,000 before and after the introduction of the tax. Data from NJ Treasury SR1A file for 1996–20011 (taxable defined as any residential sale). We implement this pre/post comparison as follows. We omit transactions within 90 days of the policy change (to avoid the retiming response) and focus on the following year (Oct. 30, 2004–Oct. 29, 2005). We rescale the period before the tax (May 3, 2003 to May 2, 2004) to account for sales growth over time. Specifically, we construct a counterfactual growth factor by taking the ratio of the count of sales within $2500 of each price from May 3, 2002 to May 2, 2003 to the count of sales from Nov. 5, 2000 through Nov. 4, 2001 (omitting sales between Nov. 2001 and May 2002 to mimic the 180 day gap around the introduction of the tax in August, 2004).

Figure 4: New Jersey Monthly Sales Above $990,000

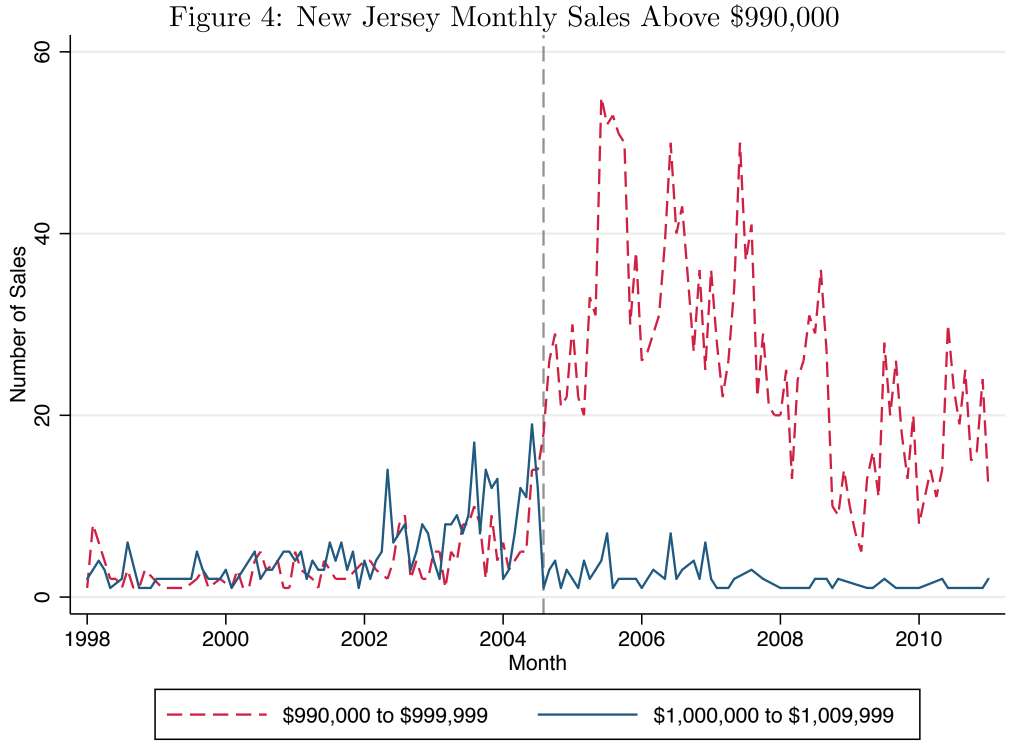

Notes: Total taxable NJ sales in given price range by month. Data from NJ Treasury SR1A file for 1998–2011 (taxable defined as any residential sale). Mansion tax introduced in August, 2004 (denoted by gray dashed line).

Figure 5: Incidence and Gap Concepts

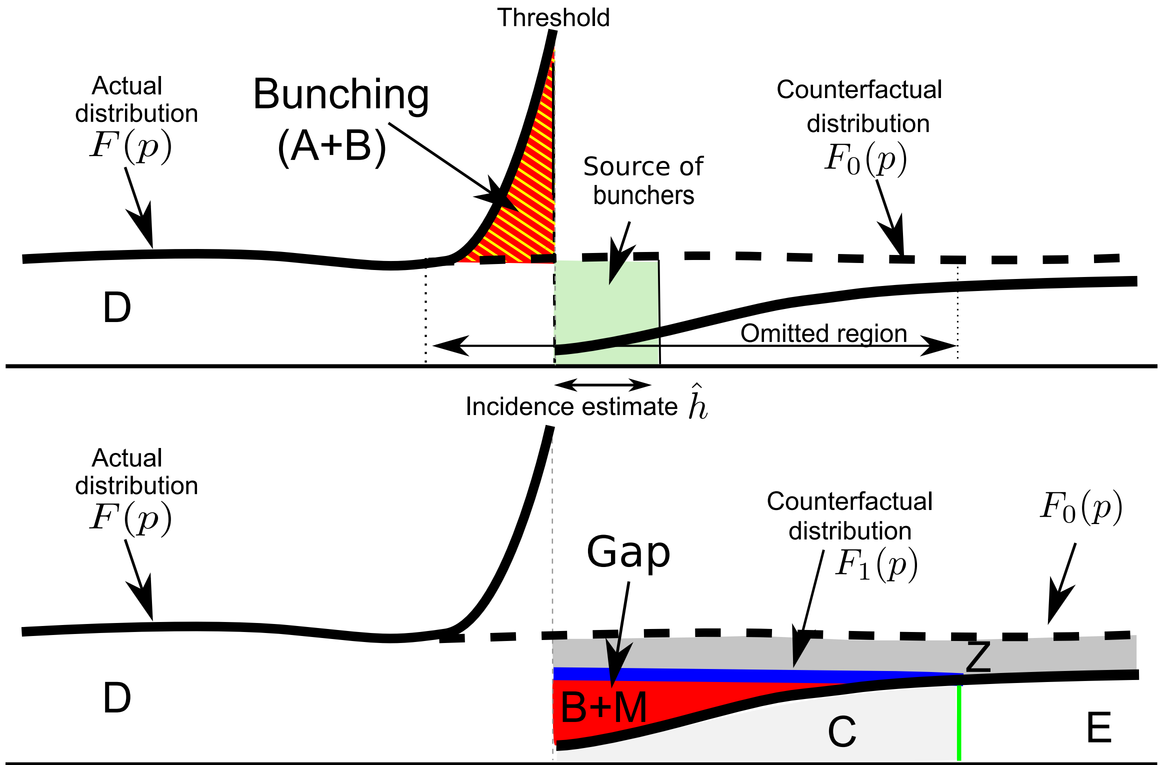

Figure 6: Bunching at the notch and efficient allocation

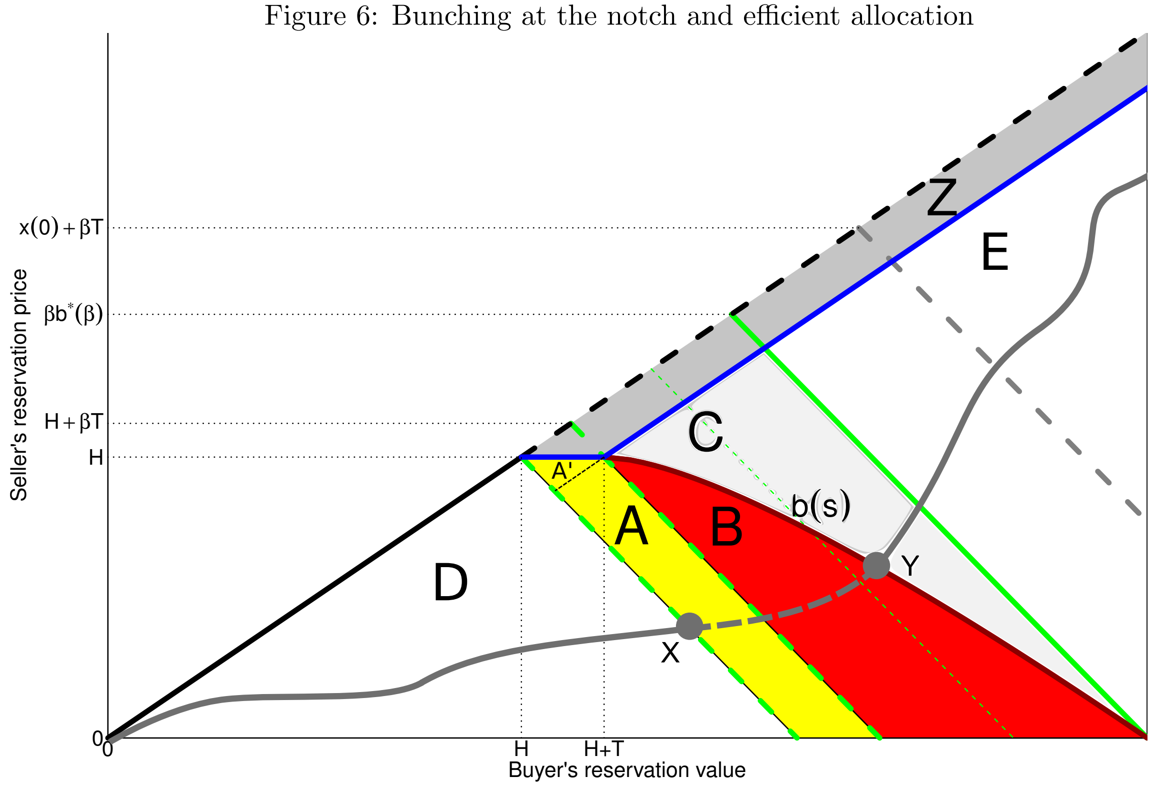

Figure 7: Distribution of Real-Estate Listing Prices in NYC (Sold Properties Only)

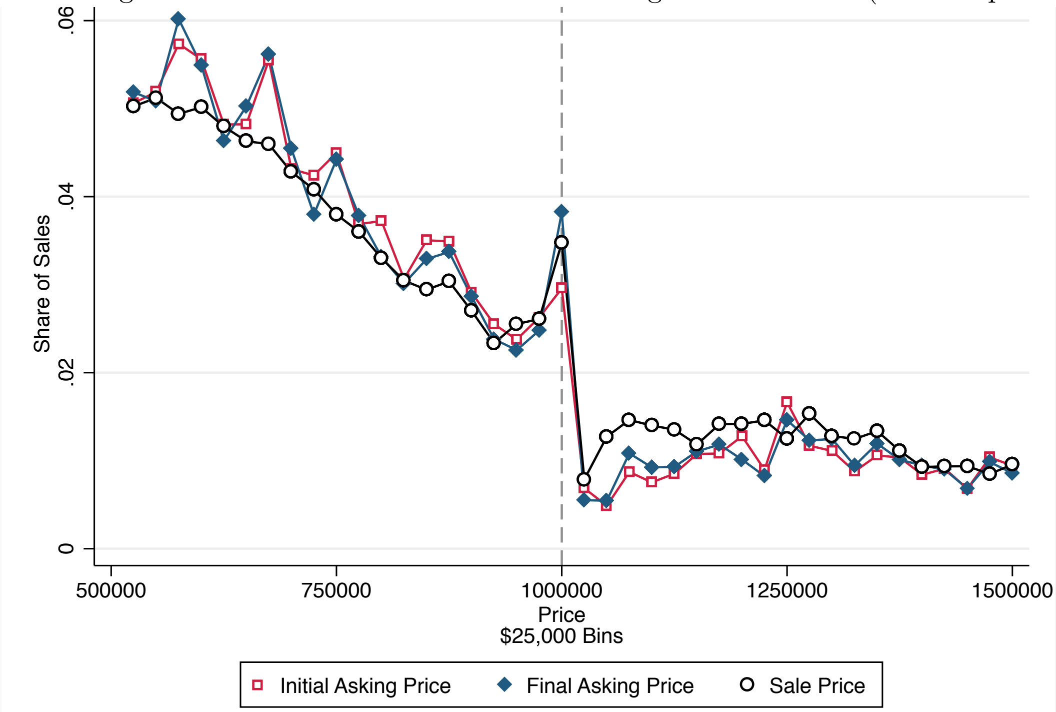

Notes: Data from REBNY listings matched to NYC Department of Finance sales records. Sample restricted to "sold" listings: last listing status is "closed" and property can be matched to NYC sales data. Smoothed plot of the distribution that accounts for round-number bunching. The log of the per-$25,000-bin counts are regressed on a cubic in price and dummy variables for multiples of $50,000 and $100,000 interacted with the price. Predicted bunching for round-number bins are then subtracted from the corresponding counts.

Figure 8: Median & 75th Percentile Price Discounts by Initial Asking Price

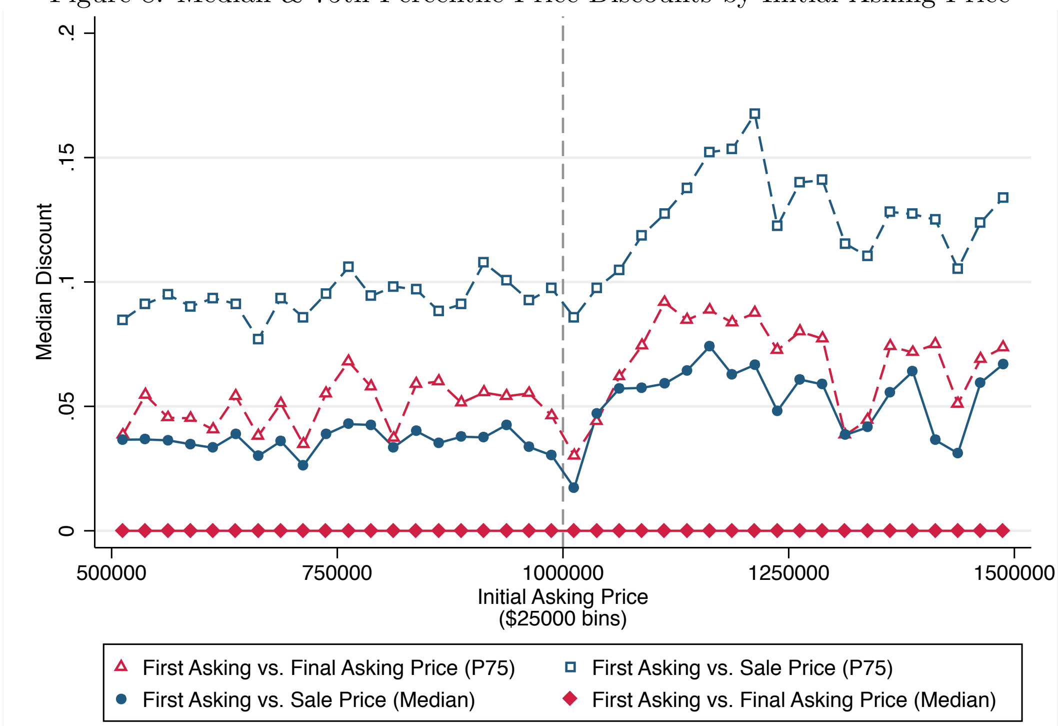

Note: Plot of the median and 75th percentile discount from initial asking to sale price (= 1 - final/initial) and initial asking to final asking price (= 1 - sale/initial) per $25,000 initial-asking-price bin. Data from REBNY listings—sample includes all closed REBNY-listed properties in the range $500,000–1,500,000 that match to NYC DOF data.

Figure 9: Predicted Price Discounts by Initial Asking Price (Relative to $1,000,000)

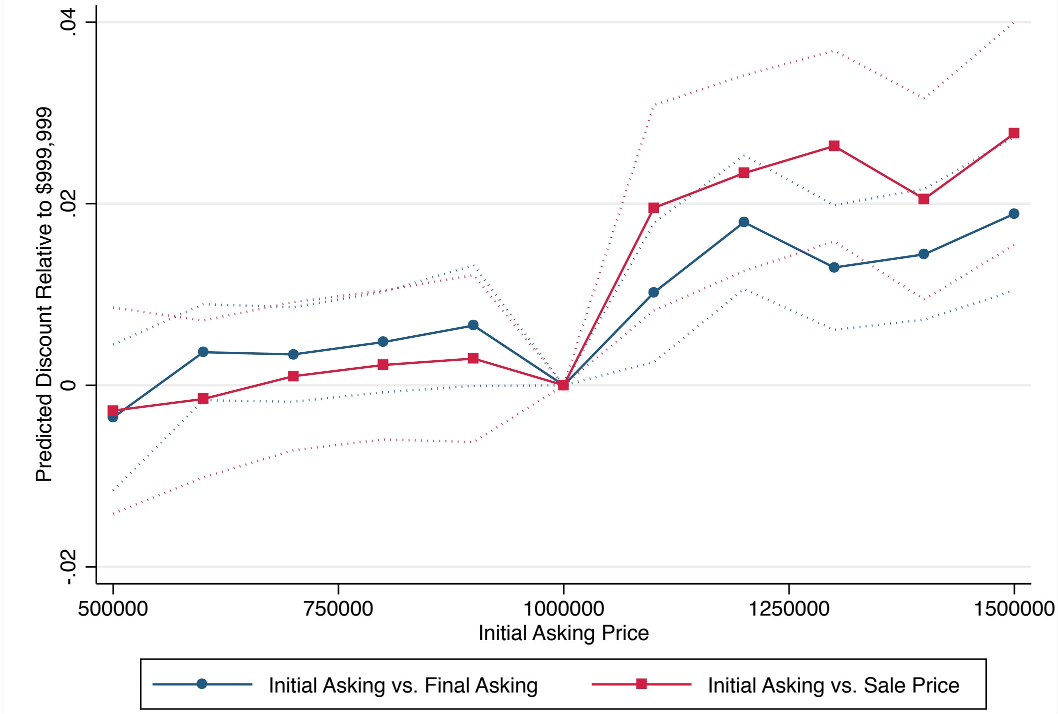

Note: Plot of difference between predicted discounts by initial asking price and prediction at $1,000,000 from regressions of discount (first to final price or first to sale price) on a linear spline in initial asking price with $100K knots between $500,000 and $1,500,000. Dashed lines represent 95% confidence intervals.

Figure 10: Predicted Dispersion of Log of Sale Price by Initial Asking Price (Relative to $1,000,000)

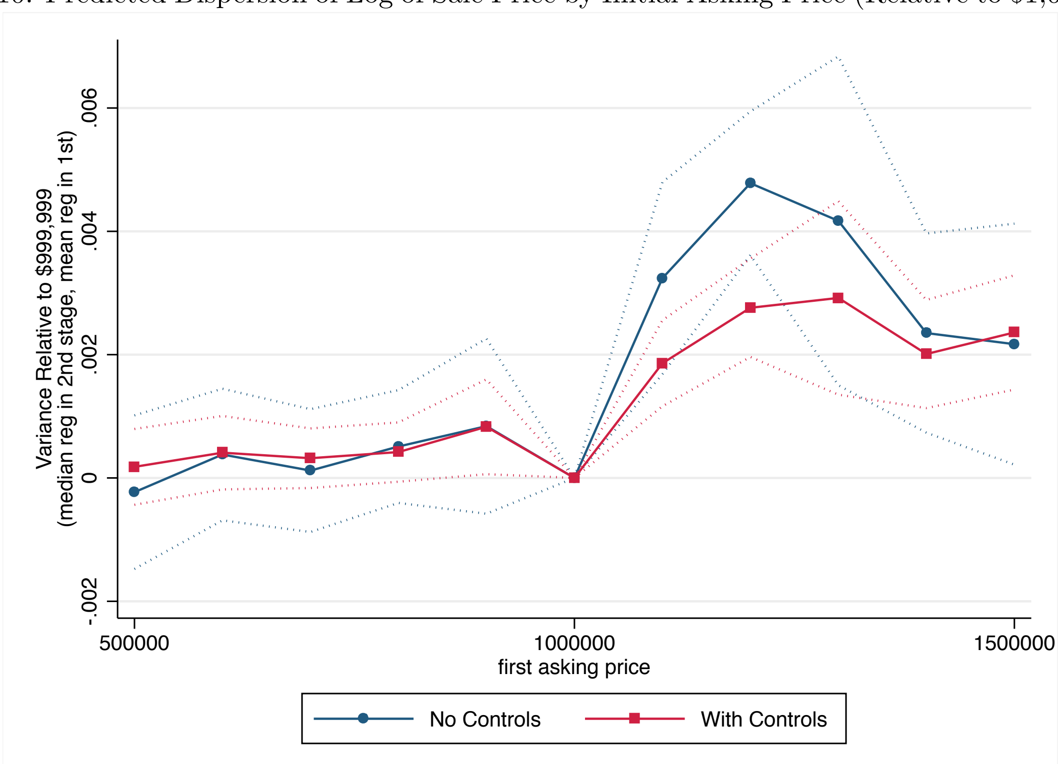

Notes: Plots of the difference between predicted values at given initial asking price and predicted value at $1,000,000 from the following procedure. The log of sale price is regressed on a linear spline in the log of initial asking price with $100,000 knots between $500,000 and $1,500,000. Squared residuals from this first stage are then regressed on a linear spline in log of initial asking price (using median regression; results are sensitive to outliers). Controls in the indicated results include year of sale, zipcode, building type, whether the sale is of a new unit, and the log of years since construction. Dashed lines represent 95% confidence intervals from 999 wild bootstrap replications of the two-stage procedure, resampling residuals in the first stage by asking-price clusters.

A Bargaining model

To identify buyer-seller pairs that move to the notch, we compare maximized surplus above the notch to surplus when price is at the notch. Surplus at the notch is given by $\beta\ln(H-s)+(1-\beta)\ln(b-H)$, while maximized surplus when the tax is due is given by $\beta\ln(\beta(b-s-T))+(1-\beta)\ln((1-\beta)(b-s-T))$. Transactions in this category that sell at the threshold satisfy:

$$f(b,s;\beta)\equiv\beta\ln(H-s)+(1-\beta)\ln(b-H)-\beta\ln(\beta)-(1-\beta)\ln(1-\beta)-\ln(b-s-T)\geq 0$$
$$(1-\beta)s+\beta b\geq H+\beta T \text{ and } b-T-s\geq 0,\ b\geq H+T,\ s\leq H$$

at the notch. That is, the surplus at the threshold has to be higher than under the alternative of selling with the tax, and the non-trivial case corresponds to transactions with positive surplus that could otherwise sell at a price higher than the threshold. We show the following results:

Lemma A.1. *Fix $0<\beta<1$ and consider matches (b,s) that satisfy $p(b,s)\geq H+\beta T$, $b-T-s\geq 0$, $b\geq H+T$, and $0\leq s\leq H$*
(a) For any value of $0\leq s\leq H$ there exists $b(s)>-\frac{(1-\beta)}{\beta}s+\frac{H+\beta T}{\beta}$ such that $f(b(s),s;\beta)=0$
(b) Matches (b,s) that satisfy $b\in\left[-\frac{(1-\beta)}{\beta}s+\frac{H+\beta T}{\beta},b(s)\right]$ locate at the notch and those with $b>b(s)$ sell with the tax at $p(b,s)-\beta T$
(c) Matches (b',s') that would otherwise sell at the same price as $(b(s),s)$ (i.e., $p(b',s')=p(b(s),s)$) sell at the notch when $s'\leq s$ and sell with the tax otherwise
(d) $b(0)$ is finite, and matches (b,s) that absent the tax would sell at prices higher than the corresponding price $p(b(0),0)=\beta b(0)$ will never bunch.

Proof. For part (a) and (b) note that for a transaction that would otherwise sell exactly at $H+\beta T$ the notch is preferred so that $f\left(-\frac{(1-\beta)}{\beta}s+\frac{H+\beta T}{\beta},s;\beta\right)\geq 0$; that $\frac{\partial f}{\partial b}=\frac{1-\beta}{b-H}-\frac{1}{b-s-T}=-\frac{H+\beta T-(\beta b+(1-\beta)s)-T}{(b-H)(b-s-T)}<0$ because $(1-\beta)s+\beta b\geq H+\beta T$; and finally that $\lim_{b\to\infty}f(b,s;\beta)=\beta\ln(H-s)-\beta\ln(\beta)-(1-\beta)\ln(1-\beta)+(1-\beta)\ln\left(\frac{b-H}{b-s-T}\right)-\beta\ln(b-s-T)=-\infty$ because all but last term converge to finite values as b increases. Hence, for each $0\leq s\leq H$, there is $b(s)$ that solves $f(b(s),s;\beta)=0$, and $b(s)$ separates positive from negative values of $f(b,s;\beta)$.
Part (c): to evaluate the effect of a change in s holding $p(b,s)$ constant, substitute $b'=-\frac{1-\beta}{\beta}s'+\frac{p(b(s),s)}{\beta}$ into $f(b,s;\beta)$ and totally differentiate with respect to s to obtain $\frac{df}{ds}=-\frac{\beta}{H-s}-\frac{(1-\beta)^2}{\beta(b-H)}+\frac{1}{\beta}\frac{1}{b-s-T}=\frac{1}{\beta}\left(\frac{1}{b-s-T}-\frac{\beta}{(H-s)/\beta}-\frac{1-\beta}{(b-H)/(1-\beta)}\right)$. Note that convexity of $\frac{1}{x}$ implies that $\frac{\beta}{(H-s)/\beta}+\frac{1-\beta}{(b-H)/(1-\beta)}\leq\frac{1}{\beta(H-s)/\beta+(1-\beta)(b-H)/(1-\beta)}=\frac{1}{b-s}$and since $\frac{1}{b-s}<\frac{1}{b-s-T}$ we have $\frac{df}{ds}>0$.
Part (d): finiteness of $b(0)$ follows from part (a). When $p(b,s)>\beta b(0)$, then there is $b'<b$ such that $p(b',s)=\beta b(0)=p(b(0),0)$. Part (c) implies that (b',s) sells with the tax because $s\geq 0$. Because $b>b'$, $\frac{\partial f}{\partial b}<0$ then implies that (b,s) also does not locate at the notch. □

The lemma establishes the existence and shape of the schedule $b(s)$, which is marked by a solid red line in Figure 6. Given the seller reservation value, matches with buyers to the left of this schedule bunch at the notch, while those above it sell with the tax. Parts (a) and (b) of the lemma establish that $b(s)$ exists and is unique for any s lower than the threshold. Part (c) shows that the slope of $b(s)$ is flatter than that of the constant-price schedules, and Part (d) shows that the schedule $b(s)$ intersects the horizontal axis at some finite value so that for sufficiently high original prices transactions will never bunch. We re-state this last observation in the following corollary:

Corollary A.2. *Transactions at the notch satisfy $H\leq p(b,s)\leq\beta b^*(\beta)<\infty$ where $b^*(\beta)$ is defined as $f(b^*(\beta),0;\beta)=0$ or, explicitly,* it solves $\beta\ln(H)+(1-\beta)\ln(b-H)-\beta\ln(\beta)-(1-\beta)\ln(1-\beta)-\ln(b-T)=0$.

The corollary follows from part (d) of the lemma. As the original price increases, the attractiveness of the notch declines so that only transactions with sufficiently high overall surplus (sufficiently low reservation price of the seller) continue to bunch. For some price, even the seller with zero reservation value will no longer agree to bunch at the notch and hence no matches corresponding to higher $p(b,s)$ will bunch either. The bound in the Corollary depends on β. Interestingly, one can show that there is a uniform and finite bound for all β, so that transactions above some finite price are never induced to bunch, regardless of the value of β.

Theorem A.3. *For any $\beta > 0$, transactions at the notch satisfy $H \leq p(b,s) \leq \beta b^*(\beta) < x(0) < \infty$, where $f(b^*(\beta),0;\beta) = 0$ and $x(\beta) \equiv \beta(b^*(\beta) - T)$ for any $\beta \in (0,1)$, $x(0) = \lim_{\beta\to 0} x(\beta)$ and the value of $x(0)$ is the solution to $\ln(H) - ln(x) - \frac{H-T}{x} + 1 = 0$.*

Proof. In what follows we change variables as $x = \beta(b - T)$ (because it turns out that $\lim_{\beta\to 0} b^*(\beta) = \infty$). Note also that part (a) of the Lemma applied to $(b^*(\beta), 0)$ implies that $x(\beta) = \beta(b^*(\beta) - T) > H$ so that we don't need to consider $x \leq H$. Define the net benefit of locating at the notch for a given x as $g(x,\beta) \equiv f(\frac{x}{\beta} + T, 0; \beta)$ or more explicitly

$$\begin{aligned} g(x,\beta) &\equiv \beta \ln(H) + (1-\beta)\ln\left(\frac{x - \beta(H-T)}{\beta}\right) - \ln\left(\frac{x}{\beta}\right) - \beta\ln(\beta) - (1-\beta)\ln(1-\beta) \\ &= \beta\ln(H) - \ln(x) + (1-\beta)\left[\ln(x - \beta(H-T)) - \ln(1-\beta)\right] \end{aligned}$$

We are interested in properties of $x(\beta)$ that solves $g(x(\beta),\beta) = 0$.

Denote by x^* the solution of $\ln(H) - \ln(x) - \frac{H-T}{x} + 1 = 0$. $x^*$is independent of β and finite. Note that $g(x,\beta)$ is continuous in x on $[H, x^*]$ and that $g(H,\beta) = (1-\beta)\ln\left[\frac{(1-\beta)H+\beta T}{(1-\beta)H}\right] > 0$ and $g(x^*,\beta) = \beta\ln(H) - \beta\ln(x^*) + (1-\beta)\ln\left[\frac{x^*-\beta(H-T)}{(1-\beta)x^*}\right] < \beta\ln(H) - \beta\ln(x^*) + (1-\beta)\left[\frac{x^*-\beta(H-T)}{(1-\beta)x^*} - 1\right] = \beta\left[\ln(H) - \ln(x^*) - \frac{H-T}{x^*} + 1\right] = 0$ (using $\ln(x) < x - 1$ and the definition of x^*). Hence, for every β, $g(x,\beta) = 0$ has a solution on (H, x^*) and, in particular, $\lim_{\beta\to 0} x(\beta)$ has to be finite.

For all $\beta \in (0,1)$, $x(\beta)$ solves $g(x(\beta),\beta) = 0$ so that $g_x x'(\beta) + g_\beta = 0$. Note that $g_x = \frac{1-\beta}{x-\beta(H-T)} - \frac{1}{x} = -\frac{\beta(x+H-T)}{x(x-\beta(H-T))}$. Clearly, for any $x > H - T$ we have $\lim_{\beta\to 0} g_x = 0$. Because $\lim_{\beta\to 0} x(\beta) > H - T$ and is finite, we also have $\lim_{\beta\to 0} g_x(x(\beta),\beta) = 0$ and $\lim_{\beta\to 0} \mid x'(\beta) \mid < \infty$. Consequently, $0 = \lim_{\beta\to 0} g_x x'(\beta) + g_\beta = \lim_{\beta\to 0} g_\beta(x(\beta),\beta) = \lim_{\beta\to 0}\left\{\ln(H) - \ln(x(\beta) - \beta(H-T)) + \ln(1-\beta) - \frac{(1-\beta)(H-T)}{x(\beta)-\beta(H-T)} + 1\right\} = \ln(H) - \ln(x(0)) - \frac{H-T}{x(0)} + 1$ as in the statement of the proposition. □

Example. When $H = 1,000,000$, $T = 10,000$, and $x(0) \approx \$1,144,717$, transactions that absent the tax would sell above this value will not bunch regardless of the value of β.

The gray dashed line on Figure 6 illustrates the bound, which corresponds to the price $x(0)$. As β changes, the slope of the corresponding line will change but it will always correspond to the price of $x(0)$. The sharp bound for a given β (the solid green line) always lies to the left of this uniform bound and converges to it as β tends to zero. While this bound is irrelevant given β, it is of natural interest when β is unknown.

A.1 Proportional tax

While considering a lump-sum tax simplifies the analysis, transaction taxes, including the mansion tax, are typically proportional. However, the results are only slightly affected when the tax is proportional. Intuitively, incentives for bunching at the notch are always determined by the level of the loss due to taxation (both due to the tax itself and any distortionary impact it might cause), rather than the rate of the tax (this is standard intensive/extensive margin distinction). The proportional tax induces re-ranking, but retains qualitative features of the solution described above. In the presence of the proportional tax the Nash bargaining outcome is given by $p^s = q(b,s;t) \equiv \beta\frac{b}{1+t} + (1-\beta)s$, where $q(b,s,t)$ denotes the seller's given types and the marginal tax rate, and $p^b = (1+t)p^s = \beta b + (1-\beta)s(1+t)$ so that the overall surplus from the transaction is equal to

$$\beta\ln(\beta) + (1-\beta)\ln(1-\beta) + \ln(b - s(1+t)) - \beta\ln(1+t)$$

As in the case of the lump-sum tax, the price $q(b,s;t)$ is linear in types so that the locus of matches with constant price remains linear (although the slope is affected by the tax rate), as in Figure 6.

It is also straightforward to show an analogous result to Lemma A.1 for the proportional case. Holding s constant, the net benefit to locate at the notch declines with b and becomes negative for sufficiently high b (because $(1-\beta)\ln(b-H)-\ln(b-s(1+t))$ declines in b). Thus, a schedule analogous to $b(s)$ also exists in the proportional case. Similarly to part (c) of the Lemma, the net surplus from switching to the notch declines in s holding $q(b,s;t)$ constant, so that the price corresponding to $b(0)$ has to constitute the upper bound of the region affected by the presence of the tax.

Finally, there is a straightforward relationship between the bounds corresponding to the lump-sum and proportional tax. To see it, note that for a given match (b,s) the value of locating at the notch is the same regardless of whether the tax is proportional or lump-sum because that allocation does not involve any tax. Consider $s=0$, and the value of b, $b^*(\beta;T)$, that as before represents a match that is indifferent between locating at the notch given the value of T. Simple inspection of the surplus for the lump-sum and proportional tax cases shows that the indifference will hold for the proportional tax as well (because the surplus will be the same as under the lump-sum tax) when the marginal tax rate is such that $\ln(b^*(\beta;T))-\beta\ln(1+t)=\ln(b^*(\beta;T)-T)$ so that $(1+t)^\beta=\frac{b^*(\beta;T)}{b^*(\beta;T)-T}$. Thus, given β, the bounds for T map into the bounds for t by this relationship.

Theorem A.3 describes a uniform bound for prices corresponding to transactions that might be affected by the lump-sum tax of T. Because bounds for proportional and lump-sum taxes are related for any β, that Theorem can be adapted to identify the corresponding bound in the proportional tax case.

Theorem A.4. *Given marginal tax rate t, define x^* as the solution to $\ln(H)-\ln(x)-\frac{H-x\ln(1+t)}{x}+1=0$. For any $\beta>0$, transactions at the notch need to satisfy $H\le q(b,s;t)\le x^*/(1+t)<\infty$, and x^* is the lowest such bound.*

Proof. To obtain the analogue of Theorem A.3, recall that given T the Theorem established the existence of the upper bound of undistorted prices below which transactions (might) relocate to the notch. Denote by $x(T)$ the uniform bound for βb identified in Theorem A.3 for a given value of T. For any β and tax rate t that satisfy $\ln(x(T)/\beta)-\beta\ln(1+t)\ge\ln(x(T)/\beta-T)$, $x(T)$ would equal or exceed the undistorted price bound for transactions relocating to the notch under proportional tax. We will find the value of t for which $x(T)$ is the smallest such a bound for any positive β. Rewrite this inequality as $\ln(x(T))-\beta\ln(1+t)\ge\ln(x(T)-\beta T)$. Note that it holds with equality when $\beta=0$. Taking derivatives of both sides with respect to β,we obtain $-\ln(1+t)$ and $-\frac{T}{x(T)-\beta T}$ respectively. In order for the inequality to hold in the neighborhood of $\beta=0$, we need to have $\ln(1+t)\le\frac{T}{x(T)-\beta T}$ for small β and when that's the case the inequality will hold for any value of β because the left-hand side is constant while the right-hand side is increasing in β. The bound will be tight when we do in fact have equality at $\beta=0$ so that $\ln(1+t)=\frac{T}{x(T)}$. Accordingly, when this relationship holds, substituting $x\ln(1+t)$ for T in the equation defining $x(0)$ in Theorem A.3 (which leads to the formula in the statement of the proposition) and solving for x will yield exactly the same solution $x^*(t)$. Finally, this procedure shows that when $s=0$, only matches with $b\le x^*(t)/\beta$ may bunch. Correspondingly, only transactions that satisfy $q(b,s;t)\le q(x^*(t)/\beta,0;t)=x^*(t)/(1+t)$ might bunch. □

Clearly, the value x^* that solves this formula is also the solution to the equation in Theorem A.3 when $T=x^*\ln(1+t)$, and the proof makes it clear that this is the right "conversion" between the proportional and lump-sum tax cases. The theorem provides a bound in terms of prices distorted by the tax $q(b,s,t)$. However, because $q(b,s,t)(1+t)\ge q(b,s,0)$ it also provides a (weaker) bound in terms of prices that are not distorted ($t=0$): transactions that bunch need to satisfy $H\le q(b,s;0)\le x^*<\infty$.

Example. For the New York and New Jersey mansion tax, $H=1{,}000{,}000$, $t=0.01$, $x^*\approx\$1{,}155{,}422$. Hence, regardless of the value of β, transactions that absent the tax would occur at prices above $1,155,422 will never bunch, while those that would occur below might bunch. Transactions that do bunch, would otherwise sell (in the presence of the tax) at no more than $x^*/(1+t)=\$1{,}143{,}982$. This is the same bound as the one corresponding to the lump-sum tax of $1{,}155{,}422\cdot\ln(1+.01)\equiv 11496.83$

B Data Appendix

New York City Department of Finance Annualized Rolling Sales. The New York City Department of Finance (NYCDOF) Annualized Rolling Sales files contain details on real-property transactions for the five boroughs from 2003 to the present (we use the data through 2011). The data are realized by the NYCDOF on a quarterly basis and are derived from the universe of transfer-tax filings (which are mandatory for all residential and commercial sales). Geographic detail for each sale includes the street address (and zip code), the tax lot (borough-block-lot number), and the neighborhood (Chelsea, Tribeca, Upper West Side, etc.). The Rolling Sales files contain limited details about the properties themselves, including square footage, number of units (residential and commercial), tax class (residential, owned by utility co., or all other property), and building class category (a more detailed property code—for example, one-family homes, two-family homes, residential vacant land, walk-up condo, etc.). Transaction details in the data include the sale price and date. A sale price of $0 indicates a transfer of ownership without cash consideration (ex. from parents to children).

New York City properties are subject to the mansion tax if they are single-, double-, or triple-family homes, or individual condo or co-op units. We define taxable sales as those transactions of a single residential unit (and no commercial units) with a building classification of "one family homes," "two family homes," "three family homes," "tax class 1 condos," "coops - walkup apartments," "coops - elevator apartments," "special condo billing lots/condo-rental," "condos - walkup apartments," "condos - elevator apartments," "condos - 2–10 unit residential," "condos - 2–10 unit with commercial use," or "condo coops/condops." We define co-ops as a building code of "coops - walkup apartments," or "coops - elevator apartments." We define a commercial sale to be a transaction with at least one commercial unit (and no residential units) or a tax class of 3 or 4.

New York State Office of Real Property Service SalesWeb. The New York State Office of Real Property Service (NYSORPS) publishes sales records for all real-property transactions (excluding New York City) recorded between 2002–2006 and 2008–2010 available through the "SalesWeb" database. Since deeds are recorded after the sale, this data includes a small number of sales from 2007. The database is compiled by ORPS from filings of the State of New York Property Transfer Report (form RP-5217).

The NYS deeds records indicate several details about each transaction and property. Transaction-specifc details include the sale price and date, the date the deed was recorded (and recording details such as book and page number), the buyer's, seller's, and attorney's name and address (often missing), the number of parcels included in the transaction, and details about the relationship between the buyer and the seller (whether the sale is between relatives, whether the buyer is also a seller, whether one party is a business or the government, etc.). Of particular interest to us is whether the sale is defined by the state as arms-length. The data dictionary defines an arms-length sale as "a sale of real property in the open market, between an informed and willing buyer and seller where neither is under any compulsion to participate in the transaction, unaffected by any unusual conditions indicating a reasonable possibility that the full sales price is not equal to the fair market value of the property assuming fee ownership", which excludes sales between current or former relatives, related companies or partners in business, sales where one of the buyers is also a seller, or sales with "other unusual factors affecting sale price."Property details include the square footage, assessed value (for property-tax purposes), address (including street address, county, zip code, school district), and the property class (one-family home, condo, etc.). We consider as subject to the mansion tax all single-unit sales with property class equal to one-, two-, or three-family residence, residential condo, or a seasonal residence.

New Jersey Treasury SR1A File. We make use of sales records from the New Jersey Treasury's SR1A file for 1996–2011, which contains records of all SR1A forms filed at the time of sale (the form is mandatory in the state for all residential sales). Each record includes the sale price and date the deed was drawn, buyer and seller name and address (often missing), deed recording details (date submitted, date recorded, document number), and whether there are additional lots associated with the sale. Property details include land value, tax lot, square footage, and property class. We define taxable sales as those with a residential property class.

New York City County Register Deeds Records. These data are collected from the county registers for the five counties in New York City: Bronx County, Kings County, New York County, Queen's County, and Richmond County. The records were collected by an anonymous private firm and made available to us by the Paul Millstein

Center for Real Estate at the Columbia Graduate School of Business.

These data include additional detail as compared to the Rolling Sales files, although at the expense of precision. Prices in this data set are rounded to the nearest $100, which leads to misallocation of sales to one side of a tax notch. Transaction details include the sale price and date, an indicator for whether the unit is newly constructed, the number of parcels being sold, whether the purchase was made in cash (i.e. whether a mortgage is associated with the sale), and indicators for private lenders and within-family sales. Property details are limited to address, zip code, and county.

Data Cleaning. We begin by dropping all transactions with a price below $100 (1,658,639 in NY State, 954,241 in NJ, and 274,118). The bulk of these transactions have a zero price, representing transfers of property between parties not associated with a proper sale (e.g., a gift or inheritance). This restriction is relatively innocuous, as our analysis focuses on sales around each tax notch (although this choice does affect the descriptive statistics). More importantly, we attempt to identify and discard all duplicate records. In New York State, we identify duplicates as sales that occur within 90 days of one another at the same street number in the same grid number (a unique tax lot id). Of these 48,073 duplicates, we always keep the later sale (in case duplicates are representative of updates to the records). For New Jersey, since we do not observe tax lots, we identify all duplicate sales that occur at the same standardized address within 90 days of one another and drop all but the final duplicate (343,221). Finally, for NYC we identify duplicates as properties in the same borough at the same standardized address that sell within 90 days of one another (20,420). While these duplicates represent a large number of sales, and there are several ways one could define duplicates, our estimates are insensitive to whether and how we clean duplicates (e.g., cleaning NY state based on address or NYC based on tax lot).

Real Estate Board of New York Listings Service. We have collected residential real-estate listings from the Real Estate Board of New York's (REBNY) electronic listing service. REBNY is a trade association of about 300 realty firms operating in Manhattan and Brooklyn. REBNY accounts for about 50% of all residential real-estate listings in these boroughs. A condition of REBNY membership is that realtors are required to post all listings and updates to the listing service within 24 hours.

Using the REBNY listing service, we have collected all "closed" (i.e. sold) or "permanently off market" residential listings posted between 2003 (when the electronic listings are first available) and 2010. REBNY listings include the typical details available on a real-estate listing: asking price, address, date on the market and a description of the property. Additionally, we observe all updates to each listing (and the dates of each update), which lets us see how asking prices evolve and determine the length of time a property is on the market. Finally, we observe the final outcome of the listing: whether the property is sold or taken off the market.

We create several variables for each REBNY listing. We define the initial asking price as the first posted price on the listing, and the final asking price as the last posted price while the listing is "active." We identify the length of time that a listing spends on the market as the number of days between the initial posting and the date that the listing is updated as "in contract." We define the discount between two prices as the percent drop in price—$\frac{p_0-p_1}{p_0}$, where p_0 and p_1 are prices and p_0 is posted before p_1.

One caveat to the REBNY listings is that the price is often not updated at the time of sale. To overcome this, we match REBNY listings to the NYCDOF data by address and date. Of the 48,220 closed REBNY listings for Manhattan, we achieve a match rate of 92%. Non-matches fall in a number of categories. Sales in some condop buildings are missing from the DOF data due to a clerical error at the NYC DOF. Some transactions contain only street address or a non-standard way of specifying the apartment number (in particular, commercial units and unusual properties such as storage units fall in this category). Occasionally, the same building may have two different street addresses and a unit may be listed differently in the two databases.At the same time, of the 23,655 Manhattan listings that are not reported as closed in the REBNY listings database, we find 7,425 corresponding sales in the NYCDOF data. We treat such matches as an indication that the property was sold without the REBNY realtor (either sold by the owner or using another realtor).

C Robustness of incidence estimates

Table A.1 demonstrates that our estimates are quite robust to variety of estimation approaches. Incidence estimates are very consistent, and gap estimates vary somewhat but remain positive and large in most specification checks that we consider. Intuitively, there are good reasons for why results may vary as one adjusts the order of polynomials and the omitted region. Both incidence and gap estimated using cross-sectional data (the only exception to it our estimates for NJ that rely on pre/post comparison) involve prediction out of sample (into the omitted region). As the size of the omitted region increases, one has to predict far out of sample so that the "forecast" error is bound to increase. Furthermore, very flexible polynomials that can fit data in sample well are not restricted in their behavior in the omitted region and in some cases may generate non-monotonicity or explosive behavior within the omitted region — overfitting is not the right approach for predicting out of sample. On the other hand, the omitted region that is too small generates bias in the estimates of the counterfactual. Nevertheless, our results are robust to reasonable modifications of our baseline specification as discussed below.

While our preferred specification uses a third-order polynomial, our incidence estimates are not too sensitive to this choice. The second through fifth rows of Table A.1 present estimates that we obtain using different orders—the results are similar, although inspection of the fit of the data suggests that very low-order polynomials cannot capture properly the shape of the distribution, while very high-order polynomials (not reported) introduce very unrealistic behavior in the omitted region. As the result, there is a bit of sensitivity to the order of polynomials in the gap estimates, which are positive and significant for all specifications, but shrink somewhat for higher orders of the polynomial.

The results are only somewhat sensitive to selecting a narrow omitted region. The estimates in the sixth through eighth rows of Table A.1 illustrate that a smaller omitted region leads to smaller incidence estimate ($3000 to $5000 less than the baseline). We do not estimate $\hat{Z}$ for the narrowest specifications since it does not make sense to restrict the gap to be so small, especially given the visual evidence of the width of the gap. Relatedly, the estimate of $\hat{Z}$ using the omitted regions through $1.1M are smaller than the baseline — understable, given the visual evidence in Figures 1, 2 and 3 indicating that the gap extends further than that (consistently with the theoretical argument as well).

On the other hand, our results are robust to extending the omitted region beyond the baseline, as is seen in Rows 10 through 12. The estimates of $\hat{Z}$ are consistently large, positive, and significant, although less precise as we use less and less data (and need to predict the counterfactual over a larger range). Reassuringly, the incidence estimates change little as we vary the upper bound of the omitted region. Similarly, none of the estimates are too sensitive to expanding the omitted region below the threshold. The results in rows twelve through fourteen of Table A.1 show that both the incidence and gap estimates grow as the bunching region is expanded below the threshold, however differences in estimates are economically small and not statistically distinguishable from the baseline. Naturally, the standard errors also grow as the omitted region is expanded.

We also estimate our counterfactuals for bunching and gap separately using only data below and above the omitted region (respectively). We present in Row 16 our estimate using a 3rd order polynomial and data below the omittted region for the bunching/incidence counterfactual and a 1st order polynomial using data above the omitted region for the missing mass, and a 2nd order above the omitted region in Row 17. Again, incidence and $\hat{Z}$ are comparable with our baseline. Furthermore, bootstrapped standard errors increase significantly as the order of polynomial increases, underscoring our earlier point that allowing for overfitting by estimating high order polynomials that are then used to project into the omitted region is a questionable approach. This observation (and visual inspection of the fit) justified our choice of the baseline specification that relies on the 3rd order polynomial and only a level shift at the threshold.

Our baseline estimate is also not sensitive to allowing for a discontinuity at the threshold. The baseline specification relies on the data both below and above the omitted region. Since the latter is distorted by the tax we rudimentarily control for it by allowing for a level shift in the distribution. The estimate in row 18 of Table A.1 demonstrates that incidence and gap increase slightly when we do not allow for this discontinuity.

For completeness, we also estimate analogous specification by OLS—this is the standard approach in the recent public finance work on notches and kinks—but we note that any of these methods involves specifying the parametric density function and the maximum likelihood estimation is a natural choice that guarantees that the estimates satisfy the law of probability rather than the hard-to-interpret mean zero residual restriction. Additionally, by requiring the data to be binned, OLS will throw out information. We report the OLS results obtained by binning into $5000 and $10,000 bins in rows 17 and 18 and conclude that they are quantitatively similar to the baseline.

The placebo estimates in Table A.2 show that our estimates for NYC are not spurious. Using the same procedure, we estimate the incidence and gap for all commercial sales (which are not subject to the mansion tax) and for residential sales at other multiples of $100,000. In all cases, we find small negative incidence estimates and relatively small gap estimates.

Figure A.1: Distribution of Sales in New York City around the $500,000 RPTT tax notch

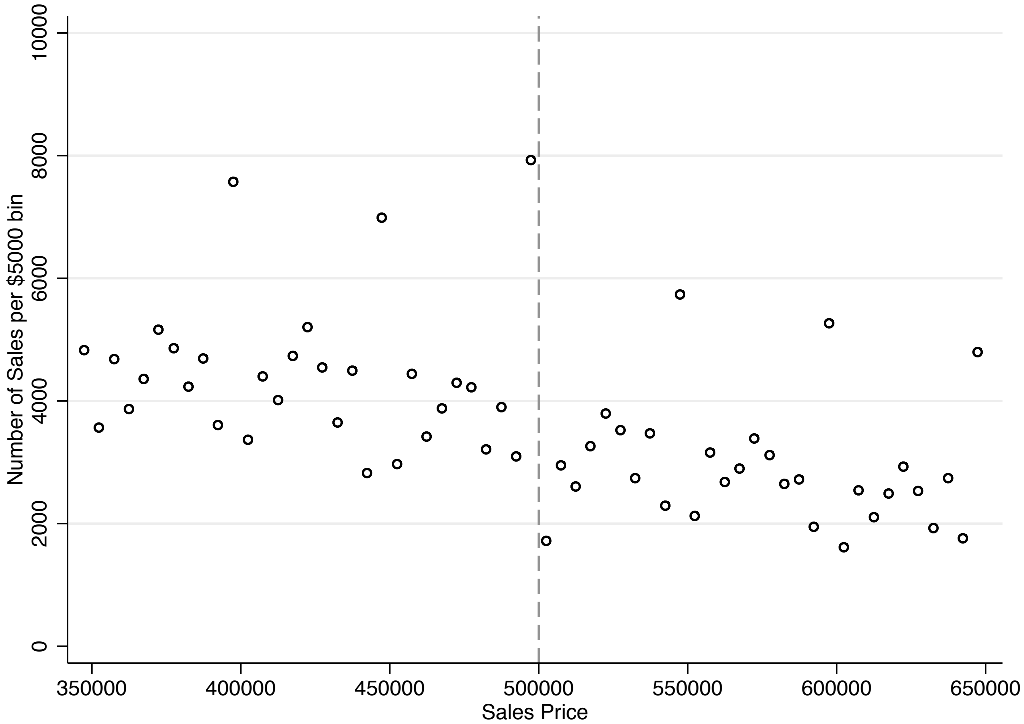

Notes: Plot of the number of sales in each $5,000 price bin between $350,000 and $650,000. Data from the NYC Rolling Sales file for 2003–2011. Both commercial and non-commercial sales are subject to the NYC RPTT.

Figure A.2: Distribution of Taxable Sales in New York State

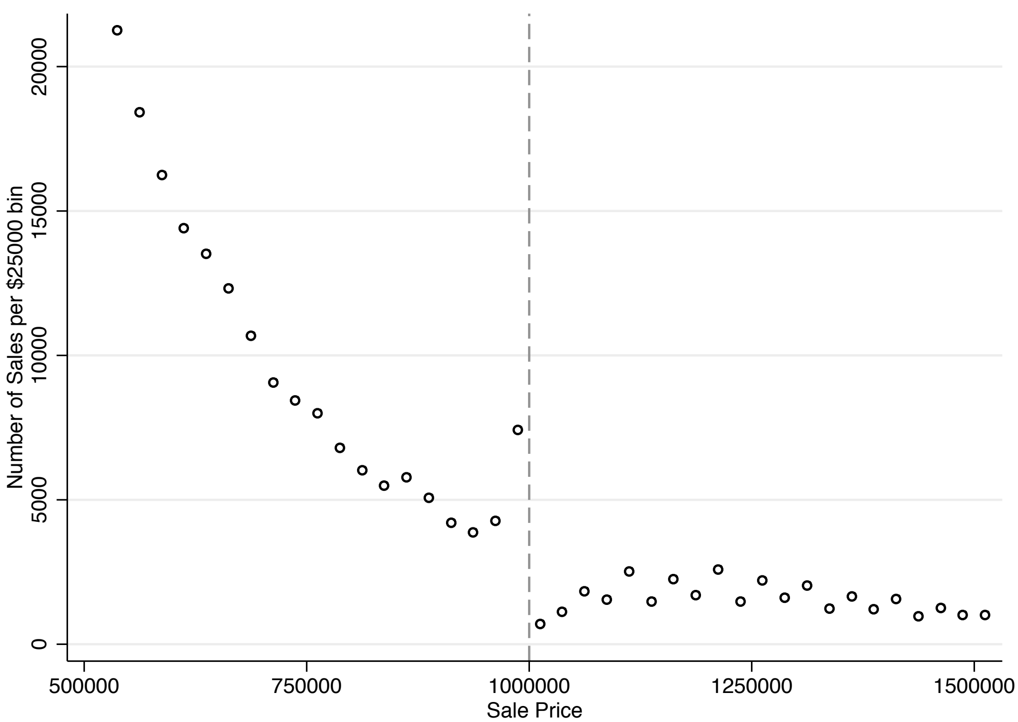

Notes: Plot of the number of mansion-tax eligible sales in each $25,000 price bin between $510,000 and $1,500,000. Data from the NYC Rolling Sales file for 2003–2011 (taxable sales defined as single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos) and from N.Y. State Office of Real Property Service deeds records for 2002–2006 and 2008–2010 (taxable defined as all single-parcel residential sales of one-, two-, or three-family homes).

Figure A.3: New Jersey Monthly Sales Above $990,000

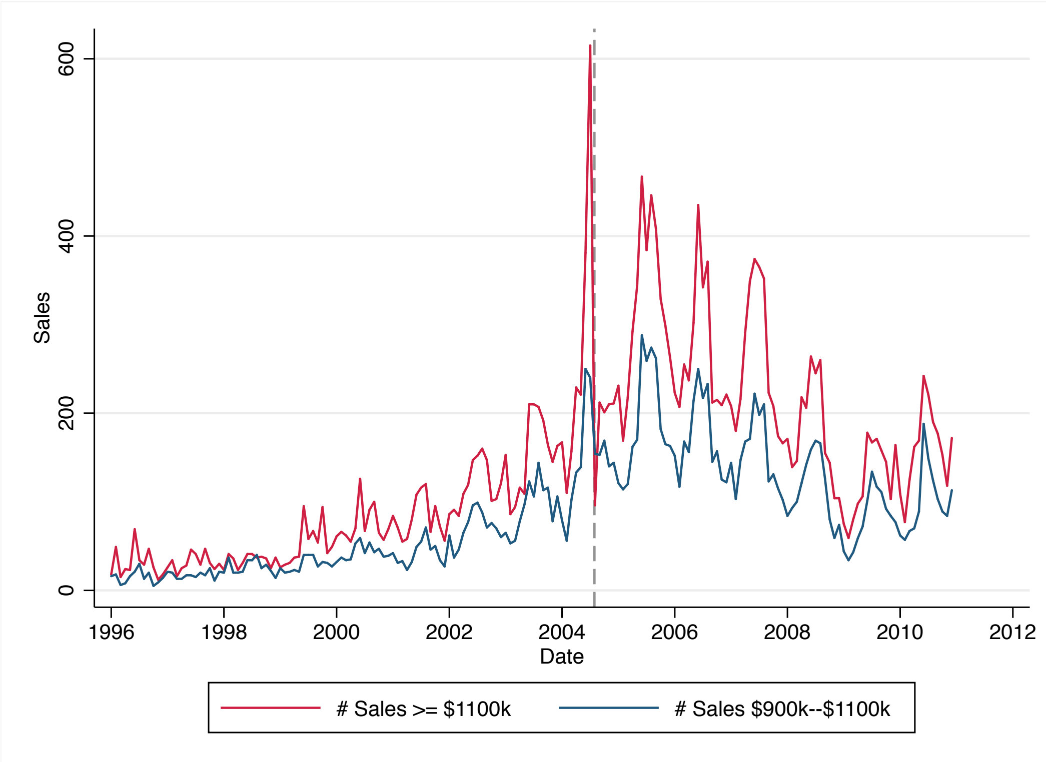

Notes: Total taxable NJ sales in given price range by month. Data from NJ Treasury SR1A file for 1996–2011 (taxable defined as any residential sale). Mansion tax introduced in August, 2004 (indicated by dashed gray line).

Figure A.4: Distribution of Monthly Sales in New Jersey ($900k – $1M)

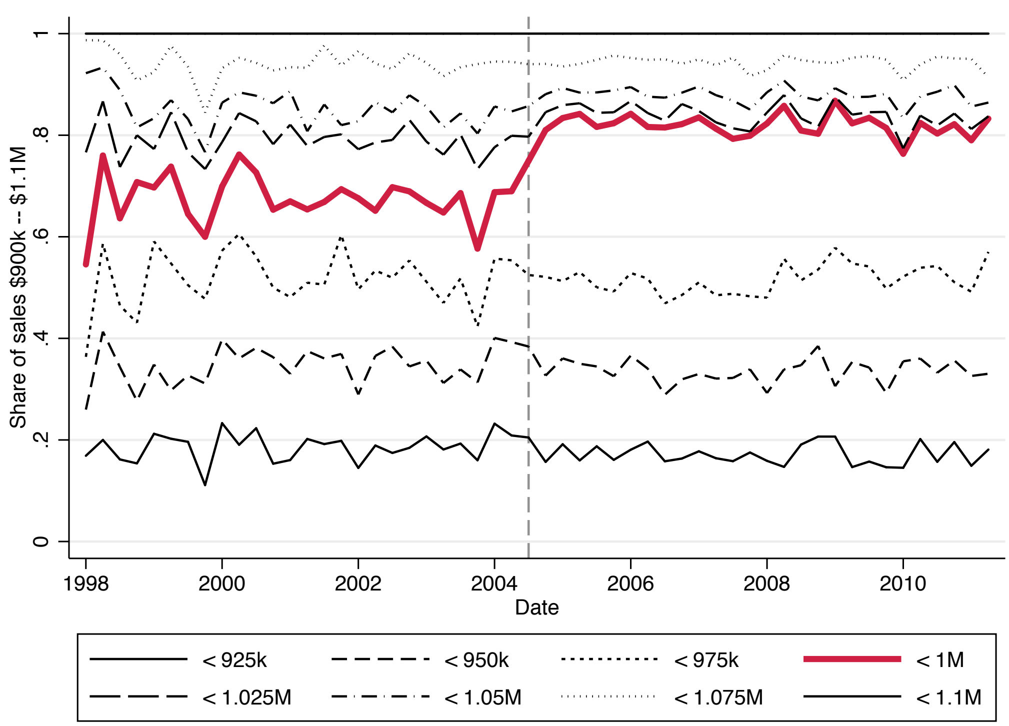

Notes: Number of taxable sales in given range as a share of total sales between $900,000 and $1,100,000 by month. Data from NJ Treasury SR1A file for 1998–2011 (taxable defined as any residential sale). Mansion tax introduced in August, 2004 (denoted by gray dashed line).

Figure A.5: NJ Local Incidence Over Time

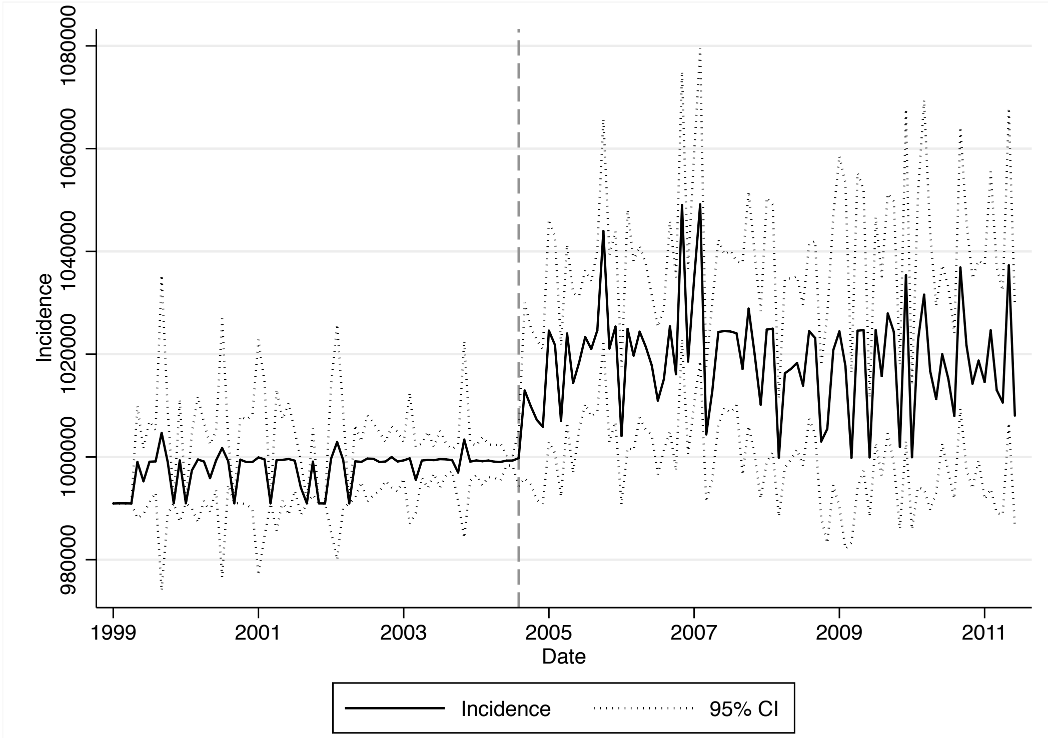

Notes: Monthly baseline local incidence estimates and 95% confidence intervals for NJ. Data from NJ Treasury SR1A file.

Figure A.6: Distribution of Real-Estate Listing Prices in NYC (Sold Properties Only)

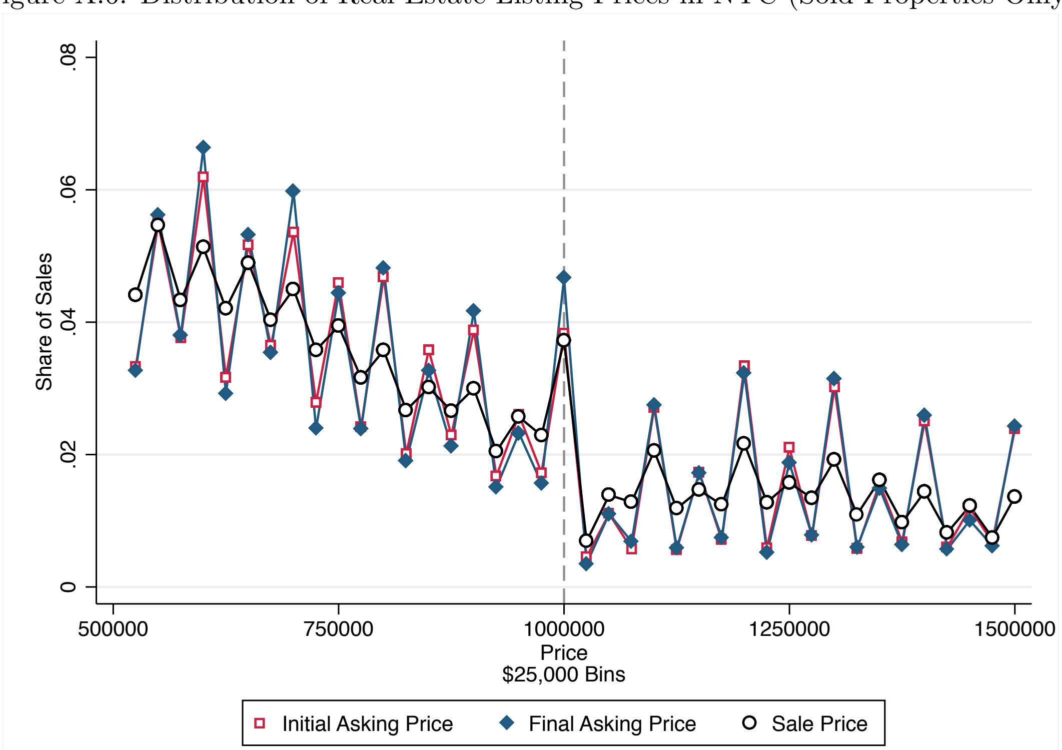

Notes: Data from REBNY listings matched to NYC Department of Finance sales records. Sample restricted to "sold" listings: last listing status is "closed" and property can be matched to NYC sales data. Plot of the number of listings per $25,000 bin as a share of all sales between $500,000 and $1,500,000 (bins centered so that the threshold bin spans $975,001–$1,000,000).

Figure A.7: Distribution of Real-Estate Listing Prices in NYC (All REBNY-Listed Properties)

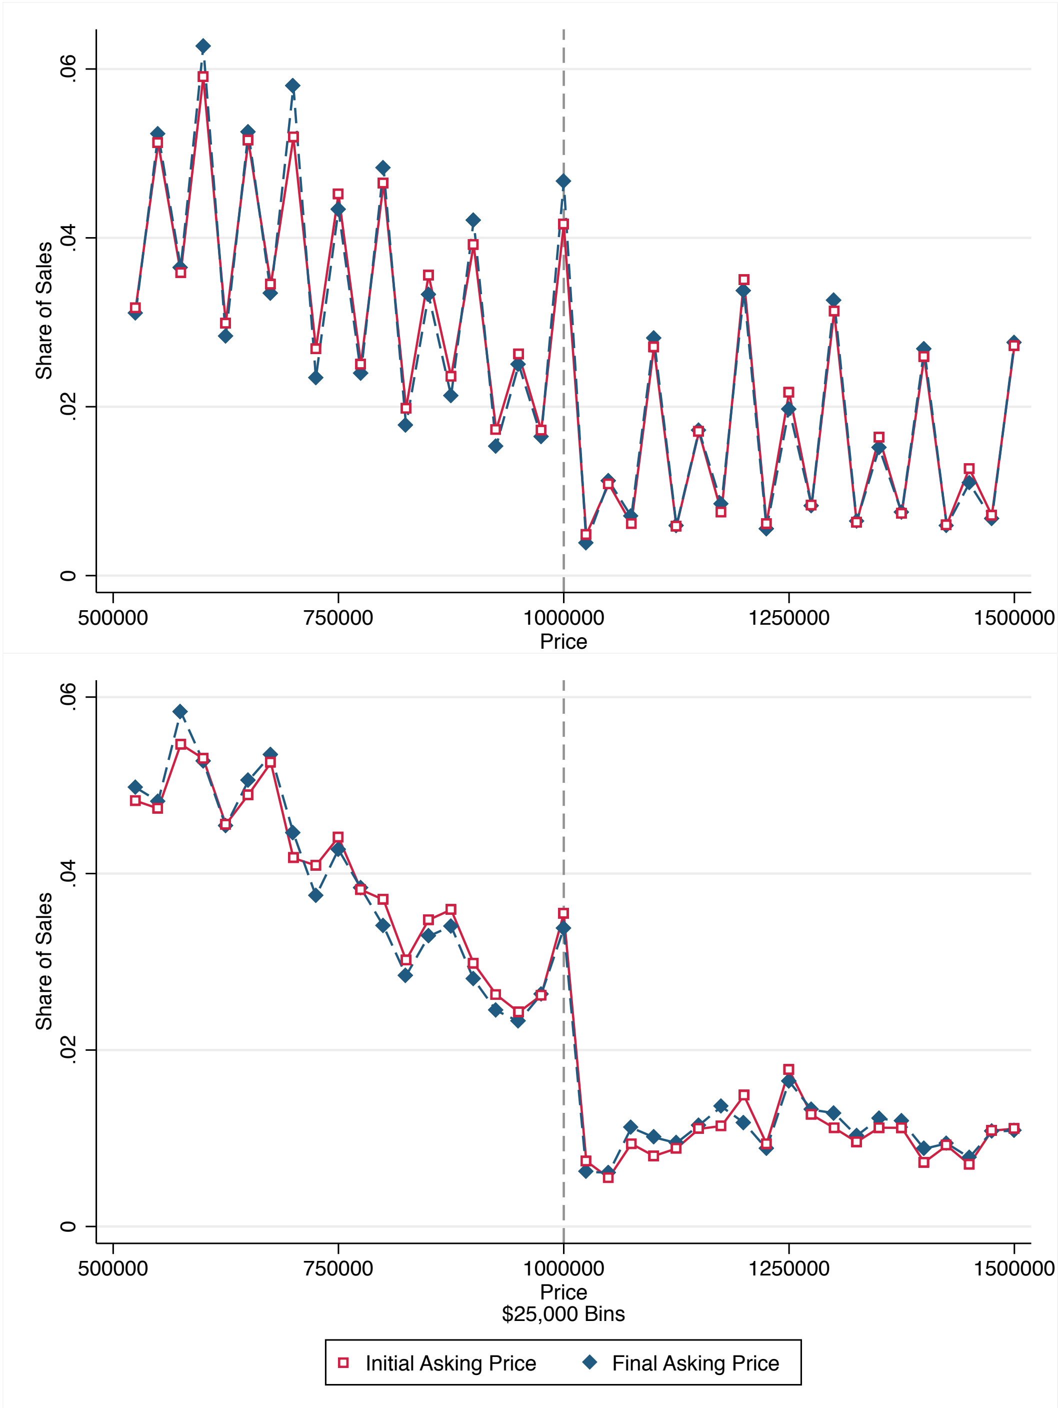

Notes: Data from REBNY listings. Sample includes all REBNY-listed sales in the given range. Panel (a) presents a plot of the number of listings per $25,000 bin as a share of all listings between $500,000 and $1,500,000 (bins centered so that the threshold bin spans $975,001–$1,000,000). Panel (b) presents a smoothed plot of the distribution that accounts for round-number bunching: the log of the per-bin counts from panel (a) are regressed on a cubic in price and dummy variables for multiples of $50,000 and $100,000 interacted with the price. Predicted bunching for round-number bins are then subtracted from the corresponding counts.

Figure A.8: Distribution of Sale Price by Initial Asking Price (with Quantile Regression)

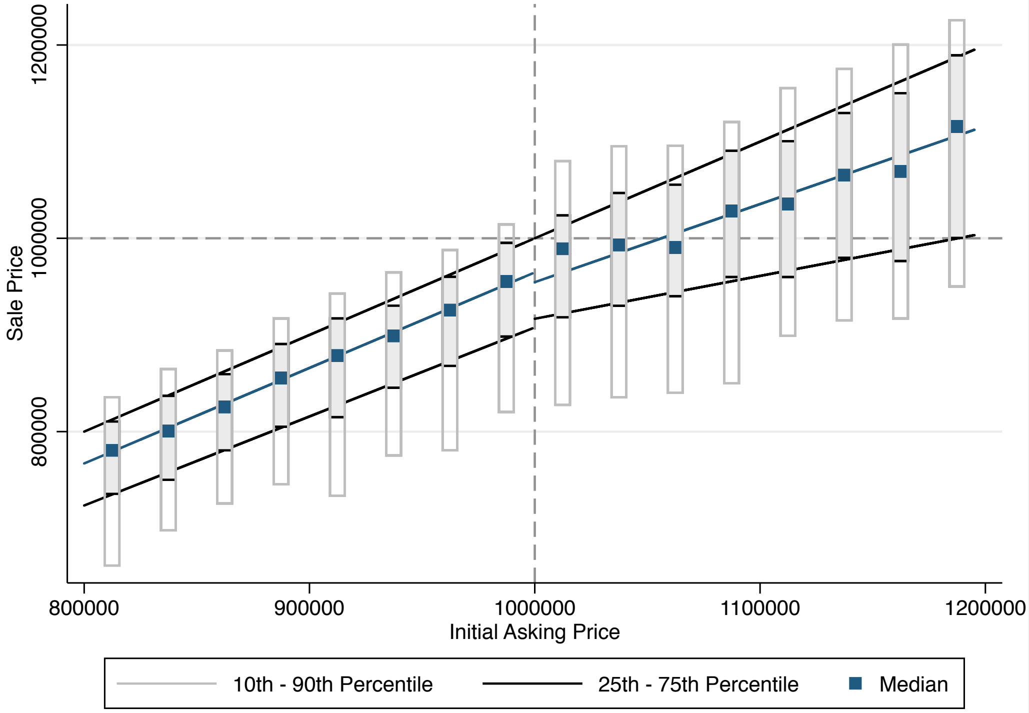

Notes: Plot of the median, 10th, 25th, 75th, and 90th percentiles of sale price per $25,000 initial-asking-price bin. Data from REBNY listings—sample includes all sold REBNY-listed properties (matched to NYC DOF) in the range $800,000–1,200,000. Lines represent quantile regressions for the given range ($800k–$990k and $1M – $1.2M).

Figure A.9: Probability that Listed Property Sells by Initial Asking Price

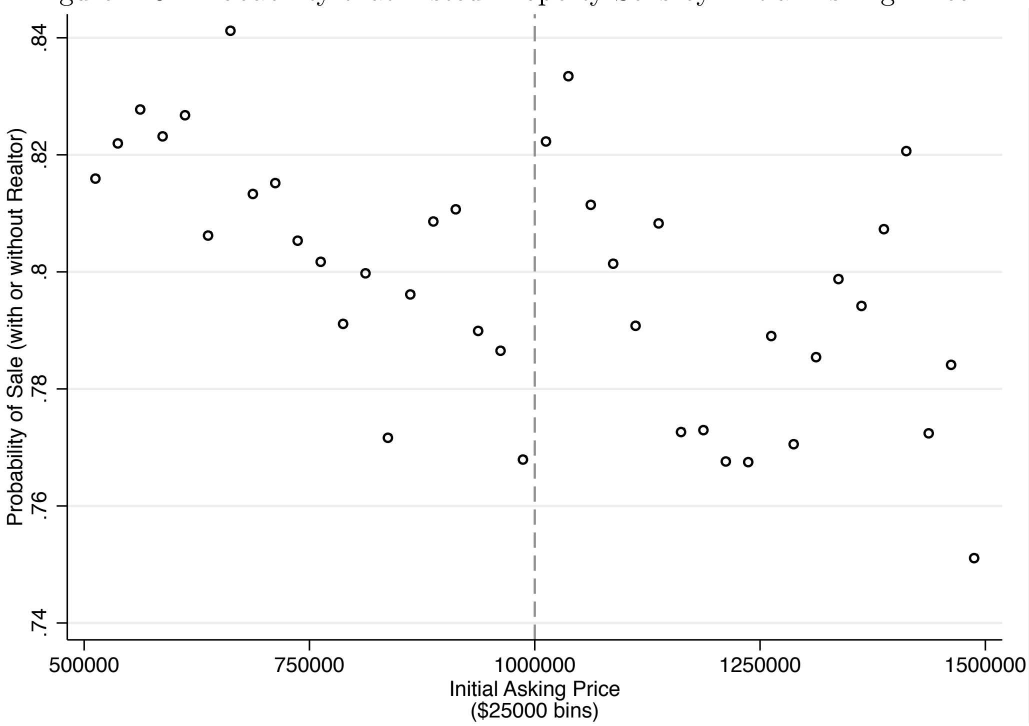

Notes: Plot of the share of REBNY-listed properties that close or are matched to a NYC DOF sale per $25,000 bin. Data from REBNY listings—sample includes all listed properties in the range $500,000–1,500,000. "Sold" defined as any property with a final listing status of "closed" or any listing that matches to NYC DOF sales.

Figure A.10: Median Days to Sale by Initial Asking Price

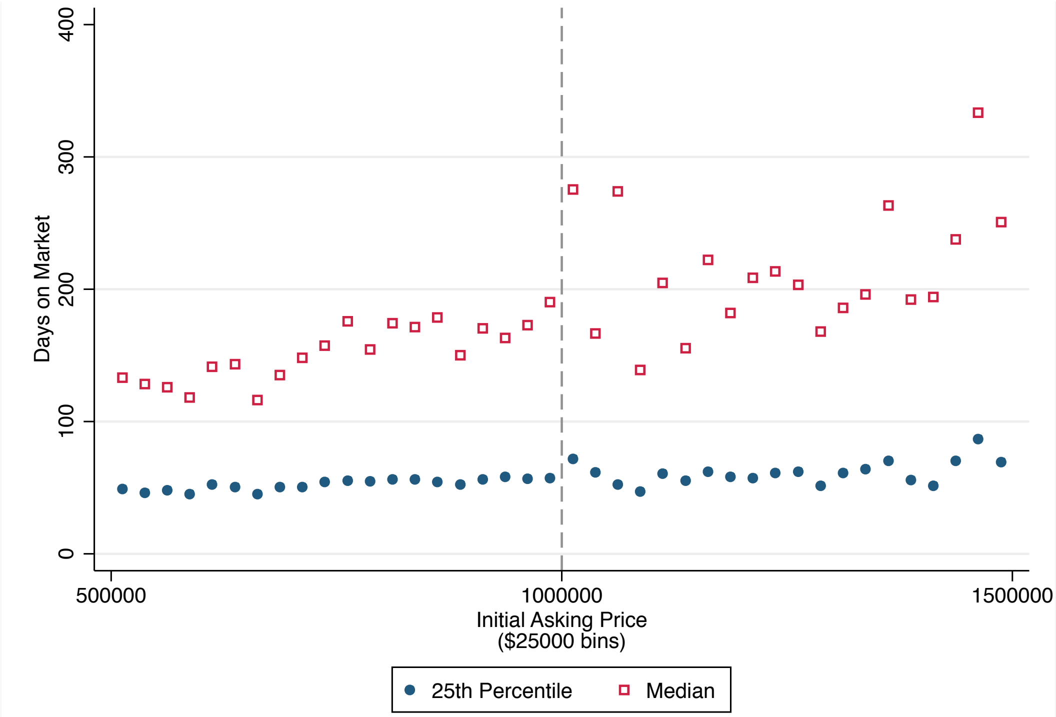

Notes: Plot of the median and 25th percentile of days to sale per $25,000 initial-asking-price bin. Data from REBNY listings—sample includes all REBNY-listed properties in the range $500,000–1,500,000. Days to sale defined as the number of days between initial listing of the property and buyer and seller entering into contract (defined as final status = "in contract"). Unsold properties are assigned a value of 999 days.

Figure A.11: Probability of Selling Without REBNY by Initial Asking Price

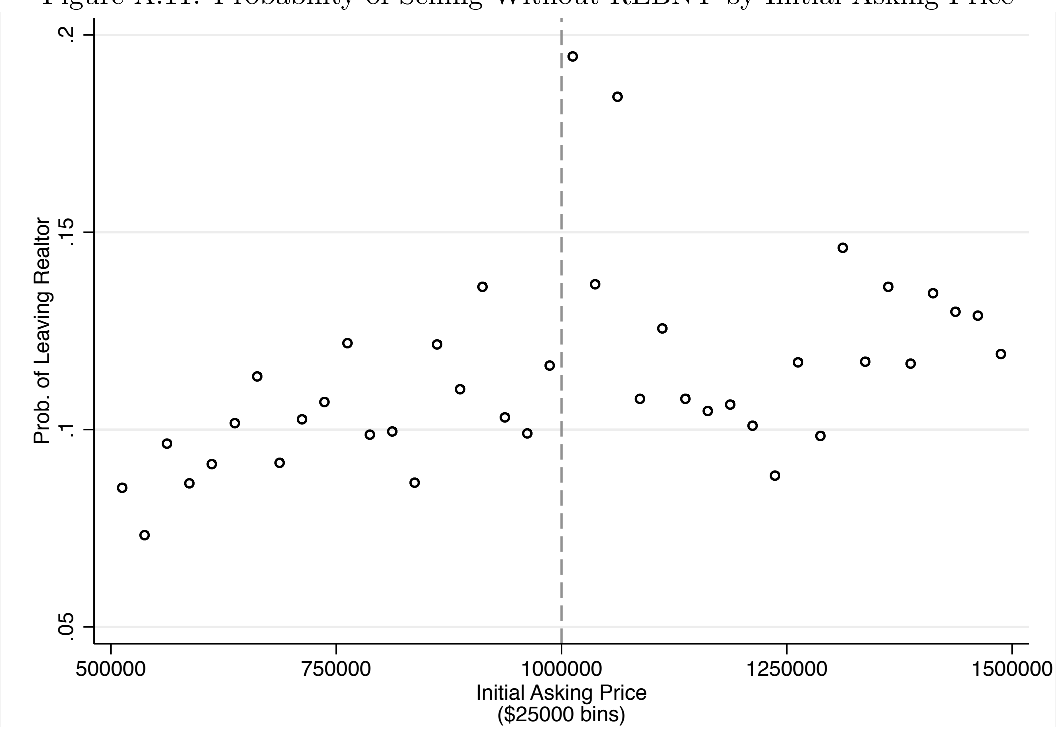

Notes: Plot of the share of REBNY listed properties that are sold in NYC DOF data, but are not listed as closed in the REBNY listing per $25,000 initial-asking-price bin. Data from REBNY listings—sample includes all REBNY-listed properties in the range $500,000–1,500,000.

Figure A.12: Median Price Discount by Final Asking Price

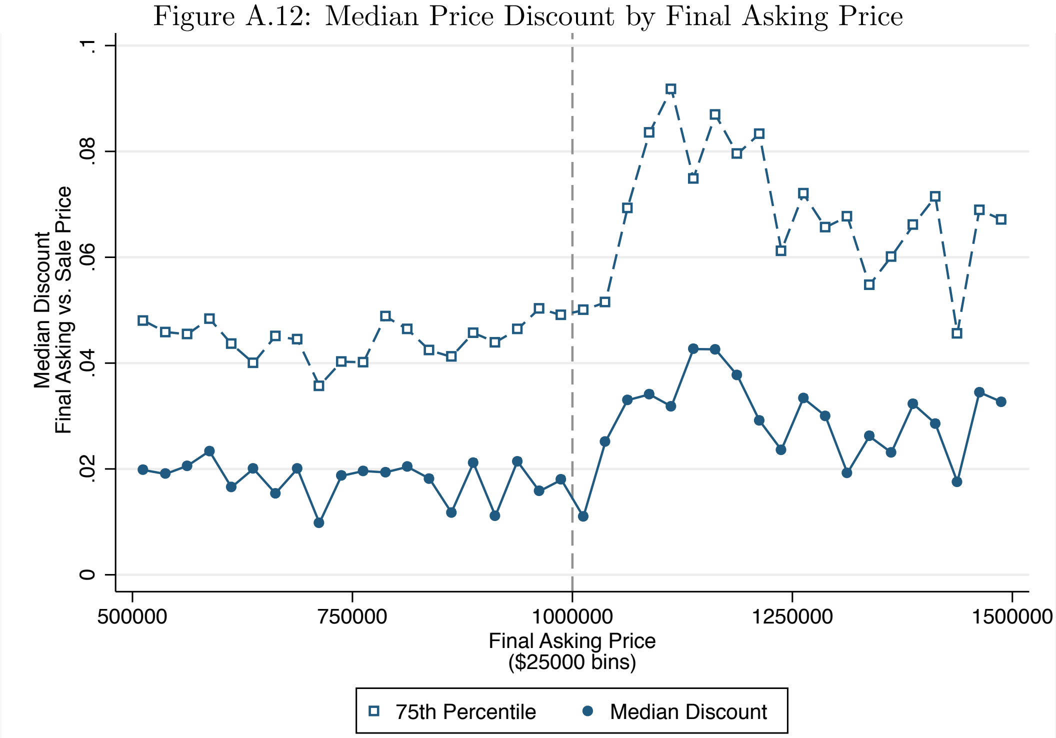

Notes: Plot of the median and 25th percentile discount from final asking price to sale price (= 1 - sale/final) per $25,000 final-asking-price bin. Data from REBNY listings—sample includes all closed REBNY-listed properties in the range $500,000–1,500,000 that match to NYC DOF data.

Figure A.13: Decomposition of Mean Price Discounts by Initial Asking Price

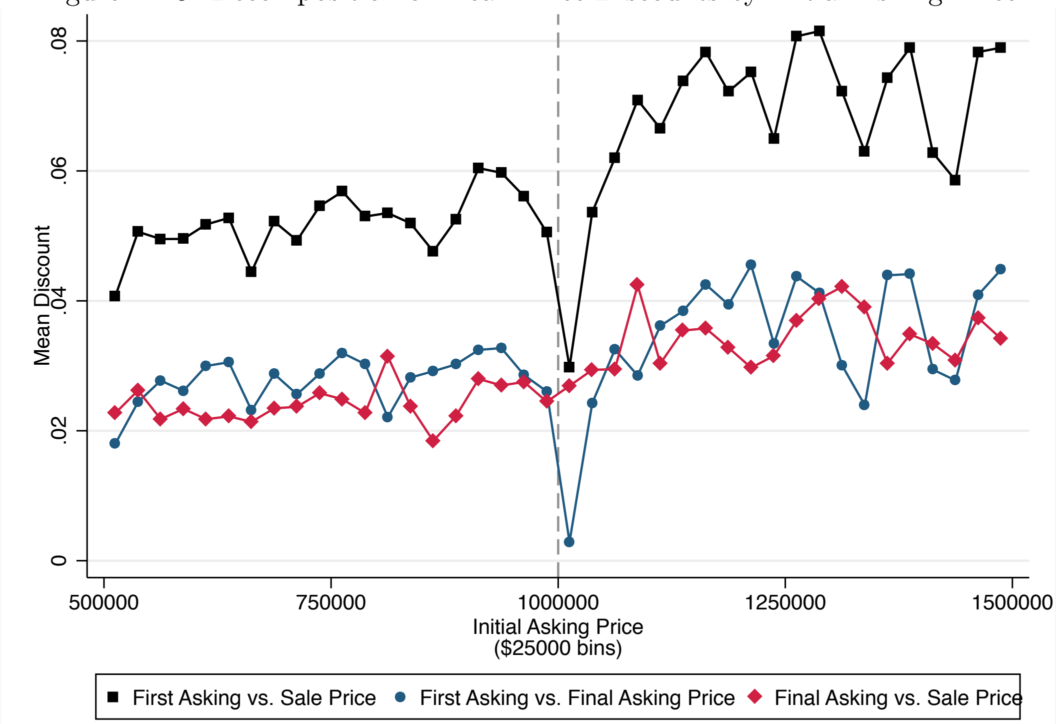

Notes: Plot of the average discount from initial asking to sale price (= 1 - sale/initial), initial asking to final asking price (= 1 - final/initial), and final asking price to sale price relative to initial asking price (= (final - sale)/initial) per $25,000 initial-asking-price bin. Data from REBNY listings—sample includes all closed REBNY-listed properties in the range $500,000–1,500,000 that match to NYC DOF data.

Figure A.14: Dispersion of Sale Price, Conditional on First Asking Price

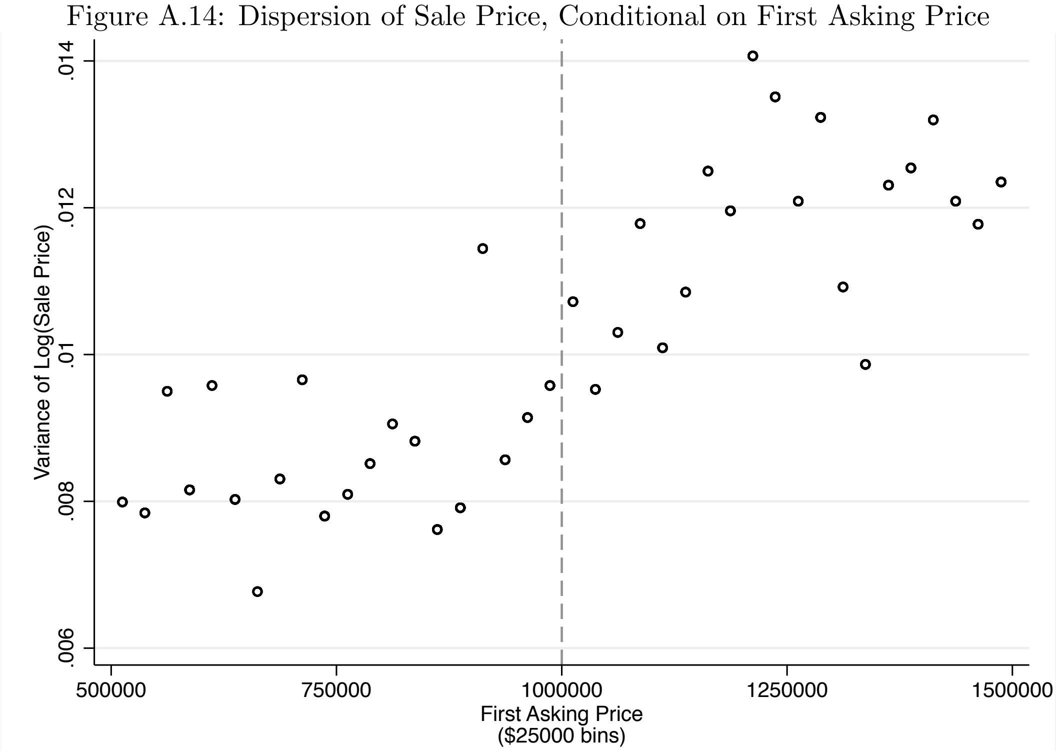

Notes: Plot of the variance of the log of sale price by $25,000 inital asking price bin. Data from REBNY listings—sample includes all sold REBNY-listed properties (matched to NYC DOF) in the range $500,000–1,500,000. Observations with price discounts (first asking price to sale price) in the 1st or 99th percentile are omitted.

Figure A.15: Dispersion of Sale Price, Conditional on Last Asking Price

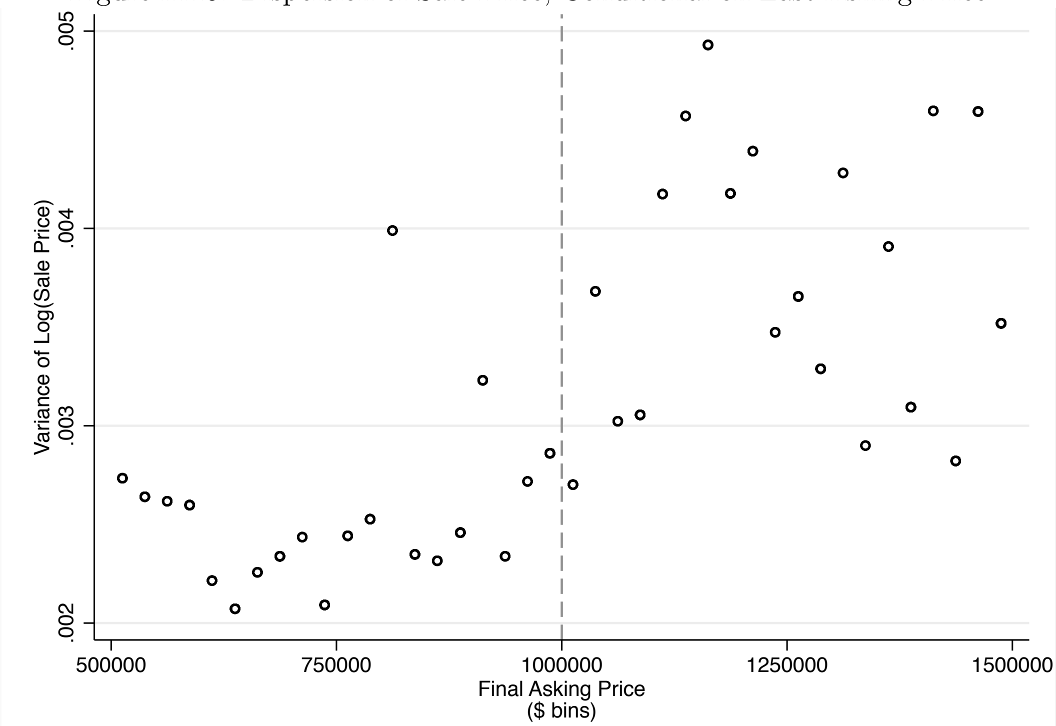

Notes: Plot of the variance of the log of sale price by $25,000 final price bin. Data from REBNY listings—sample includes all sold REBNY-listed properties (matched to NYC DOF) in the range $500,000–1,500,000. Observations with price discounts (first asking price to sale price) in the 1st or 99th percentile are omitted.

Figure A.16: Predicted Dispersion of Log of Sale Price (Median Regression in First Stage)

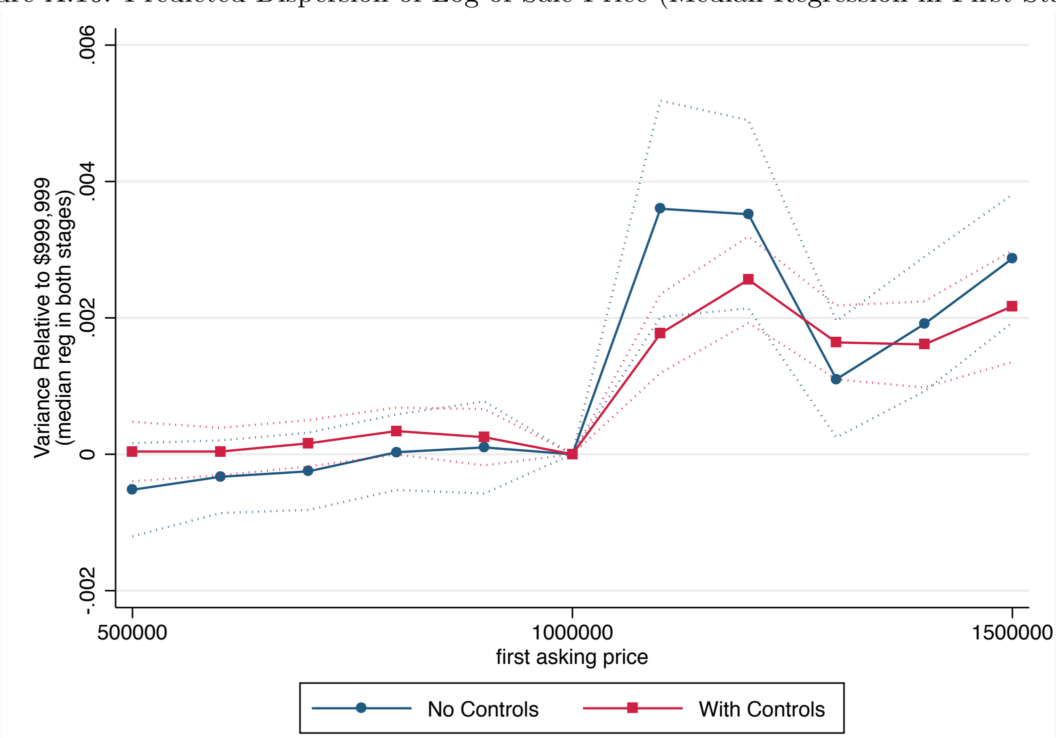

Notes: Plots of the difference between predicted values at given initial asking price and predicted value at $1,000,000 from the following procedure. The log of sale price is regressed (median regression) on a linear spline in the log of initial asking price with $100,000 knots between $500,000 and $1,500,000. Squared residuals from this first stage are then regressed on a linear spline in log of initial asking price (using median regression; results are sensitive to outliers). Controls in the indicated results include year of sale, zipcode, building type, whether the sale is of a new unit, and the log of years since construction. Dashed lines represent 95% confidence intervals from 999 wild bootstrap replications of the two-stage procedure, resampling residuals in the first stage by asking-price clusters.

Table A.1: Mansion Tax, NYC

	Specification	Incidence	Std. Error	$\hat{Z}$	Std. Error	n
1.	Baseline: 3rd Order, Omit $990k – $1.155M	21542.098	1124.005	43861.766	3990.977	102493
2.	1st Order	24095.152	475.318	106372.697	4807.543	102493
3.	2nd Order	22910.857	965.353	46988.518	4118.622	102493
4.	4th Order	21115.344	1160.988	37069.861	5734.794	102493
5.	5th Order	18444.127	1490.541	33652.867	5868.148	102493
6.	Omit $990k – $1.01M	16308.129	974.380	.	.	108766
7.	Omit $990k – $1.025M	16676.799	967.177	.	.	108378
8.	Omit $990k – $1.050M	18168.244	1062.597	2311.782	1280.552	107635
9.	Omit $990k – $1.1M	20163.154	1134.501	15145.428	2117.299	105549
10.	Omit $990k – $1.2M	21978.363	1079.676	51725.536	5165.912	100846
11.	Omit $990k – $1.255M	22122.063	1127.693	53596.361	8617.068	97673
12.	Omit $990k – $1.3M	21917.696	1119.279	25178.003	10285.031	96094
13.	Omit $980k – $1.155M	24980.343	957.747	43954.765	4051.863	100706
14.	Omit $970k – $1.155M	26742.767	1634.857	44726.198	3810.847	99741
15.	Omit $960k – $1.155M	37687.334	2374.628	49452.601	4938.454	98915
16.	3rd Order for LHS, 1st Order for RHS	20034.887	1372.811	32006.872	3891.654	102493
17.	3rd Order for LHS, 2nd Order for RHS	20034.887	1372.811	37243.921	9670.122	102493
18.	No Discontinuity at Threshold	21601.728	1048.292	44349.086	2316.284	102493
19.	OLS, $5000 Bins	19539.143	2093.029	.	.	102493
20.	OLS, $10000 Bins	20956.556	1539.757	.	.	102493

Notes: Estimates by MLE as described in text (or OLS with binned data where indicated) using data from NYC Department of Finance Rolling Sales file for 2003–2011. Sample is restricted to all taxable sales (single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos) with prices between $510,000 and $1,500,000. We do not estimate $\hat{Z}$ for specifications 6 and 7 since our method mechanically restricts the missing mass/gap to the range between $1M and the top of the omitted region. Since the omitted regions above the notch for these two specifications are small, the estimates of the missing mass in the gap are artificially small. Specification 15 estimates the counterfactual separately using data below the omitted region of $990k–$1,155,422 to estimate bunching and incidence, and data above the omitted region to estimate the gap and $\hat{Z}$.

Table A.2: NYC Mansion Tax: Placebos

Cutoff	Incidence	Std. Error	$\hat{Z}$	Std. Error	n
Commercial	-634.223	69.369	2203.190	512.087	5616
600,000	-689.324	32.958	2430.942	386.221	74477
700,000	-708.755	30.548	1257.927	258.632	86026
800,000	-787.246	29.686	360.535	280.648	93148
900,000	-732.660	32.098	-985.343	197.414	98003
1,100,000	-937.686	141.524	3457.587	559.861	106344
1,200,000	-669.331	42.662	1346.163	591.304	103660

Notes: Data from NYC Department of Finance Rolling Sales file for 2003–2011. Commercial sales are defined as any transaction of at least one commercial unit and no residential units or a NYC tax class of 3 (utility properties) or 4 (commercial or industrial properties) and are not subject to the mansion tax. Placebo estimates are found using the baseline MLE procedure around the given cutoff.

Table A.3: Local Incidence Over Time

	NYC					NYS					NJ				
Year	Incidence	Std. Error	$\hat{Z}$	Std. Error	n	Incidence	Std. Error	$\hat{Z}$	Std. Error	n	Incidence	Std. Error	$\hat{Z}$	Std. Error	n
1996	.	.	.	.	.	.	.	.	.	.	-832.902	2010.303	591.586	4577.626	1929
1997	.	.	.	.	.	.	.	.	.	.	-846.881	892.57	4672.449	3965.021	2308
1998	.	.	.	.	.	.	.	.	.	.	-516.749	325.587	-7.224	2278.698	3176
1999	.	.	.	.	.	.	.	.	.	.	-1651.362	1641.192	5799.470	5644.423	4193
2000	.	.	.	.	.	.	.	.	.	.	-855.516	219.819	1610.774	751.773	5402
2001	.	.	.	.	.	.	.	.	.	.	-890.483	362.37	4516.891	1772.83	6461
2002	.	.	.	.	.	22203.323	3459.236	54264.664	14737.658	9596	-739.276	105.021	2921.912	818.467	9606
2003	14624.035	4208.772	40388.150	15073.276	7310	19003.617	3607.358	63687.589	12761.623	12462	-636.847	95.723	2995.208	668.838	12953
2004	23429.529	2510.092	19803.960	10841.689	10451	22087.229	2651.167	63869.400	11328.725	17389	-785.303	94.236	3425.059	736.190	10800
2004 (post)	.	.	.	.	.	.	.	.	.	.	5996.987	4145.917	38365.000	16506.710	8442
2005	32614.473	4526.286	64561.575	12876.311	13554	24014.876	1833.177	47688.156	10156.431	21765	24229.456	1233.458	55087.717	10767.801	24921
2006	21603.437	2520.204	47836.366	11784.422	14728	24217.067	1562.105	28400.202	10617.074	17958	22080.594	2567.267	34109.311	9633.650	21963
2007	17487.074	2388.516	42113.910	8840.712	16477	22673.583	10770.933	16409.173	38609.604	948	24119.928	2025.234	27611.911	10546.270	18921
2008	17685.933	2587.644	34771.013	9461.669	12911	23867.938	2571.473	14847.095	11051.018	11143	19047.947	3479.866	31128.821	10877.795	12913
2009	20449.578	3434.685	1338.866	10748.730	8682	23930.513	3294.993	24523.627	14400.244	8370	20673.616	3852.975	24910.197	13266.069	10039
2010	17125.476	3781.441	74753.214	15731.867	8534	22620.920	3170.810	21216.032	12132.485	8831	19520.012	3566.167	37097.427	13552.642	10948
2011	21739.302	2853.967	70566.195	15725.257	9846	.	.	.	.	.	22460.739	5574.445	27266.456	21619.086	3789

Notes: Baseline local incidence estimates by year for given geography. Data for NYC from Department of Finance Rolling Sales files, restricted to taxable sales (single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos) in the given year. Data for NYS from the Office of Real Property Services deeds records, restricted to taxable sales (all single-parcel residential sales of one-, two-, or three-family homes) in the given year. NYS ORPS data covers 2002–2007 and 2009–2010; observations in 2007 are from sales made in 2007, but recorded in 2008–2010. Data for NJ from NJ Treasury SR1A file.

Table A.4: Predicted Price Discounts

	Discount: Initial Asking to Sale Price		Discount: Initial Asking to Final Asking	
	Slope	Value at Lower Knot relative to $1,000,000	Slope	Value at Lower Knot relative to $1,000,000
Prediction at $1,000,000	*0.05117*		*0.02464*	
500000 to 600000	0.12785	-0.00279	0.71910	-0.00355
	(0.6455)	(0.00579)	(0.44164)	(0.00410)
600000 to 700000	0.25113	-0.00151	-0.02462	0.00364
	(0.40788)	(0.00441)	(0.24894)	(0.00270)
700000 to 800000	0.12379	0.00100	0.13650	0.00340
	(0.37161)	(0.00416)	(0.26585)	(0.00266)
800000 to 900000	0.07008	0.00224	0.18161	0.00476
	(0.39359)	(0.00419)	(0.29979)	(0.00282)*
900000 to 1000000	-0.29398	0.00294	-0.65794	0.00658
	(0.46877)	(0.00469)	(0.33941)	(0.00339)*
1000000 to 1100000	1.95390	0	1.02142	0
	(0.57745)***	0	(0.39059)***	0
1100000 to 1200000	0.38174	0.01954	0.77419	0.01021
	(0.68556)	(0.00577)***	(0.47691)	(0.00391)***
1200000 to 1300000	0.29954	0.02336	-0.49885	0.01796
	(0.65681)	(0.00550)***	(0.45059)	(0.00375)***
1300000 to 1400000	-0.58272	0.02635	0.14506	0.01297
	(0.65944)	(0.00536)***	(0.43487)	(0.00350)***
1400000 to 1500000	0.72161	0.02052	0.44734	0.01442
	(0.74929)	(0.00565)***	(0.52392)	(0.00367)***
1500000		0.02774		0.01889
		(0.00626)***		(0.00431)***
Prediction at $1,500,000	*0.07892*		*0.04354*	

Notes: Data from REBNY listings matched to NYC DOF sales records. Estimates from regression of given discount on linear spline in initial asking price ($100k knots between $500k and $1.5M). Slope estimates are for the given interval; predicted values are the difference between the predicted discount at the lower knot of the given interval and the predicted value at $1,000,000.

Table A.5: Predicted Price Dispersion

	Initial Asking		Final Asking		Initial Asking		Final Asking	
Prediction at $1,000,000	0.00333	0.00220	0.00099	0.00061	0.00646	-0.00017	0.00004	-0.00115
500000	-0.00023	-0.00052	0.00009	-0.00019	0.00018	0.00004	-0.00002	0.00001
	(0.00065)	(0.00035)	(0.00022)	(0.00010)	(0.00032)	(0.00022)	(0.00010)	(0.00006)
600000	0.00038	-0.00033	0.00006	-0.00005	0.00041	0.00004	0.00002	0.00003
	(0.00056)	(0.00026)	(0.00020)	(0.00007)	(0.00028)	(0.00018)	(0.00007)	(0.00004)
700000	0.00012	-0.00025	0.00008	-0.00015	0.00032	0.00016	0.00008	0.00002
	(0.00050)	(0.00027)	(0.00023)	(0.00007)**	(0.00025)	(0.00018)	(0.00013)	(0.00004)
800000	0.00051	0.00003	0.00015	-0.00015	0.00042	0.00034	0.00010	0.00002
	(0.00046)	(0.00029)	(0.00025)	(0.00008)*	(0.00023)*	(0.00018)*	(0.00015)	(0.00005)
900000	0.00084	0.00010	-0.00021	0.00013	0.00083	0.00025	0.00002	0.00010
	(0.00071)	(0.00034)	(0.00021)	(0.00011)	(0.00039)**	(0.00022)	(0.00009)	(0.00006)*
1000000	0	0	0	0	0	0	0	0
	0	0	0	0	0	0	0	0
1100000	0.00324	0.00360	0.00184	0.00143	0.00186	0.00177	0.00111	0.00095
	(0.00079)***	(0.00079)***	(0.00067)***	(0.00023)***	(0.00033)***	(0.00030)***	(0.00033)***	(0.00012)***
1200000	0.00478	0.00352	0.00139	0.00083	0.00276	0.00256	0.00065	0.00069
	(0.00058)***	(0.00075)***	(0.00030)***	(0.00016)***	(0.00040)***	(0.00033)***	(0.00016)***	(0.00010)***
1300000	0.00417	0.00110	0.00101	0.00014	0.00292	0.00164	0.00089	0.00027
	(0.00137)***	(0.00041)***	(0.00061)*	(0.00013)	(0.00081)***	(0.00028)***	(0.00049)*	(0.00008)***
1400000	0.00235	0.00191	0.00114	0.00058	0.00201	0.00161	0.00049	0.00056
	(0.00082)***	(0.00052)***	(0.00038)	(0.00013)***	(0.00044)***	(0.00034)***	(0.00015)***	(0.00009)***
1500000	0.00217	0.00287	0.00109	0.00056	0.00236	0.00217	0.00055	0.00049
	(0.00098)**	(0.00046)***	(0.00048)***	(0.00012)***	(0.00045)***	(0.00040)***	(0.00016)***	(0.00012)***
Prediction at $1,500,000	0.00550	0.00508	0.00208	0.00117	0.00882	0.00200	0.00060	-0.00066
First-Stage R^2/Pseudo R^2	0.55660	0.68100	0.59190	0.80230	0.61540	0.71750	0.64260	0.82020
Median Reg. in 1st Stage		X		X		X		X
Property Controls					X	X	X	X

Notes: Data from REBNY listings matched to NYC DOF sales records. The log of sale price is regressed on a linear spline in the log of initial asking price with $100,000 knots between $500,000 and $1,500,000 and a set of controls for property characteristics where indicated (property zipcode, year of sale, property type, and years since construction). Squared residuals from this first stage are then regressed on a linear spline in log of initial asking price and controls where applicable (using median regression; results are sensitive to outliers). Estimates are the difference between the predicted discount at the given knot and the predicted value at $1,000,000. Standard errors found by 999 bootstraps of the two-stage procedure, resampling first-stage residuals, drawing from a moving block of 50 observations.