Property Transaction Taxes and the Housing Market: Evidence from Notches and Housing Stimulus in the UK*

Michael Best & Henrik Jacobsen Kleven†

Preliminary Draft

December 8, 2012

Abstract

We provide evidence on the effect of property transaction taxes (stamp duty) on house prices, the timing of house transactions, and the volume of house transactions. To address these questions, we exploit administrative data covering the universe of stamp duty tax returns in the United Kingdom from 2005-2011 along with compelling sources of identifying variation. First, discontinuous jumps in the stamp duty at threshold property prices—notches—allow us to estimate the effect of the tax on house prices. Second, anticipated and unanticipated changes in the tax schedule allow us to estimate the dynamics of price responses and timing effects on house transactions. Third, a temporary exemption of some properties from the tax—a stamp duty holiday—aimed at stimulating the housing market during recession allows us to provide micro evidence on macro stimulus policy. We find that the effect of transaction taxes on house prices is large (often larger than the tax itself) and that dynamic adjustment to changes in transaction taxes is very fast. We also find that the timing of house transactions responds sharply to anticipated tax increases. Finally, temporary cuts in transaction taxes successfully stimulate housing market activity in the short run—a 1% cut in the tax achieves a stimulus effect of 10% additional transactions at its peak—but the temporary boost in activity is followed by a slump in activity after the policy is withdrawn. The cumulative effect of a 16-month stimulus program was fully neutralized only 10 months after the end of the program.

*We would like to thank Tim Besley, Raj Chetty, Camille Landais, Attila Lindner, Emmanuel Saez and Joel Slemrod for helpful suggestions. We would also like to thank the staff at Her Majesty's Revenue & Customs' (HMRC) datalab for access to the data and their support of this project. The HMRC data used in this paper is Crown Copyright. The research datasets used may not exactly reproduce HMRC aggregates. The use of HMRC data in this work does not imply the endorsement of HMRC in relation to the interpretation or analysis of the information. All results have been screened by HMRC to ensure confidentiality is not breached. All remaining errors are the authors'.

†Best: Department of Economics, London School of Economics (m.c.best@lse.ac.uk). Kleven: Department of Economics, London School of Economics (h.j.kleven@lse.ac.uk)

1 Introduction

Government intervention in the housing market is pervasive in most countries. Many of these interventions aim to promote homeownership through various tax and mortgage subsidies, while other interventions aim to regulate housing market activity in the short run, motivated by macroeconomic stabilization. One policy that partially offsets these pro-homeownership policies is the imposition in many countries—including the US and the UK—of substantial transaction taxes in connection with the buying and selling of property. In the UK, for example, property transaction taxes have raised around 0.6% of GDP over recent years,[1] making them one of the UK government's major sources of revenue.[2] Several countries have recently experimented with temporary cuts in transaction taxes (or transaction subsidies) in an attempt to stimulate the housing market in response to the global recession, including a transaction tax holiday in the UK and a home buyer tax credit in the US. Despite the ubiquity of property transaction taxes and their potential importance for both short-run activity and long-run allocation, there is relatively little empirical analysis of their implications. This paper takes a step towards filling this gap by analyzing the UK version of this policy, known as the Stamp Duty Land Tax.

Our paper is based on a context and dataset that offer important methodological advantages. One advantage is our unique access to administrative stamp duty records covering the universe of property transactions in the UK from 2005–2011, roughly 9 million property transactions, with rich tax return information on each transaction. Another advantage is the presence of compelling quasi-experimental variation in the UK stamp duty, allowing us to credibly estimate how prices and transaction volumes respond to transaction taxes as well as the dynamics of such responses over short-term and medium-term horizons. The fact that the data cover several years before and after the onset of the current recession makes it possible to estimate the effects of transaction taxes in both good and bad times.

Two features of the UK stamp duty are crucial for our analysis. First, the tax liability depends on the transaction price according to a progressive schedule that produces discontinuous jumps in tax liability—*notches*—at cutoff prices.[3] For example, the tax rate jumps from 1% to 3% of the entire transaction price at a cutoff of £250,000 (about $400,000), creating an increase in tax liability of £5,000 (about $8,000) as the house price crosses this cutoff. Such notches create very strong incentives for transactions to take place at a price just below the cutoff rather than at prices above the cutoff, thereby creating a hole in the price distribution on the high-tax side and excess bunching in the price distribution on the low-tax side of the notch. Using the approach developed by Kleven & Waseem (2012) in the context of labor supply, we use notches to non-parametrically identify price responses to transaction taxes and we use changes in notches to identify the dynamics

[1] Revenues are strongly pro-cyclical and hence have fluctuated a lot over the last decade and were hit hard by the slump in the housing market.

[2] see http://www.hmrc.gov.uk/stats/tax_receipts/tax-receipts-and-taxpayers.pdf, last accessed on 24 October 2012.

[3] Notches are not uncommon in the taxation of property transactions and feature for example in both the US and Canadian tax systems.

of such responses.

Second, a *stamp duty holiday* lasting 16 months eliminated transaction taxes in a certain price range. The beginning of this holiday (tax cut) was fully unanticipated, whereas the end of the holiday (tax hike) was fully anticipated and therefore created a *time notch* that can be used to investigate short-term timing responses to tax changes. Moreover, the fact that the holiday was implemented in a specific price range allows us to combine the notches strategy with a differences-in-differences strategy to investigate the causal effect of fiscal stimulus on housing market activity.

Our main findings and contributions are the following. First, we set out a theoretical framework that will guide the empirical analysis and provide some predictions that we test empirically. This framework characterizes the effect of transaction taxes on the market value of houses ("house prices") and on the number of houses traded.[4] Second, we show empirically that there is large and sharp bunching below notch points and that there are large holes above notch points in the distribution of house prices. Our bunching estimates imply that transaction taxes reduce house values by more than the tax itself, which in turn implies that these house value responses are not simply incidence effects, but also contain real housing demand responses. Third, we find that the dynamic adjustment of bunching and holes to changes in notches is very fast, with a new steady state emerging in about 3 months for unanticipated changes and almost immediately for anticipated changes. Fourth, we find very strong short-term timing responses to the anticipated tax change created by the end of the tax holiday.

Fifth and finally, we show that temporary housing stimulus successfully boosts short-run activity as trading volumes in the treatment group clearly diverge from trading volumes in a variety of possible control groups during the stamp duty holiday. At its peak, a 1% cut in transaction taxes increases market activity by 10%. However, the effect is offset by a drop in activity after the removal of the stimulus policy. The cumulative increase in house purchases as a result of the 16-month tax holiday is fully neutralized by the post-holiday slump after only 10 months, consistent with the qualitative findings of Mian & Sufi (2012) on the effects of a much shorter stimulus on auto purchases in the US. We argue that our finding of quick reversal is particularly important because it is based on a stimulus program that was relatively long-lasting and aimed at a market seen to be at the heart of the macroeconomy as well as the recent financial and economic crisis.

This paper contributes to several literatures. First, a voluminous literature has studied the distortionary effects of standard property taxes and income tax subsidies to owner-occupied housing (such as the deductibility of mortgage interest and the exemption of imputed rental income), including for example Rosen (1979a,b, 1985) and Poterba (1984, 1992). More recently, Adelino *et al.* (2012) study the effects of changes in credit availability and terms as houses move across the conforming loan limit. Since transaction taxes such as the one studied here also affect the user cost of housing but are structured very differently, they offer an alternative way to address some of the same fundamental questions. Second, a large empirical literature has examined the impact of capital

[4]Strictly speaking, house values and house prices are not the same thing because the former is, in principle, equal to the price per unit of housing times the quantity of housing services. Since the per-unit price is never observed, we use the standard terminology "house price" to refer to the total value of the dwelling.

gains taxation on asset sales and asset prices (e.g. Feldstein *et al.*, 1980; Auerbach, 1988; Burman & Randolph, 1994) and some of this work has focused specifically on the taxation of housing capital gains (Cunningham & Engelhardt, 2008; Shan, 2011). Capital gains taxes and transaction taxes share the feature that tax liability is triggered by a transaction, with the key difference being that transaction taxes fall on the entire value of the asset and not just on the appreciation of the asset. This difference in tax base combined with the notched tax structure in the UK create much larger tax variation than typically created by capital gains taxes. Third, our paper is related to recent work using micro-level variation to estimate the effect of macro stimulus programs (Johnson *et al.*, 2006; Agarwal *et al.*, 2007; Mian & Sufi, 2012).

Fourth, a few recent papers consider the implications of property transaction taxes in different settings (Besley *et al.*, 2011; Dachis *et al.*, 2012). In contrast to these papers, in our paper we are able to simultaneously exploit a large dataset of administrative tax records and multiple sources of quasi-experimental tax variation, allowing for sharp non-parametric identification of several margins of response. Most closely related, in contemporaneous work, Kopczuk & Munroe (2012) and Slemrod *et al.* (2012) study the effects of transaction tax notches in New York and Washington DC, respectively. However, while these papers focus on effects at the higher end of the price distribution, our setting features notches throughout the price distribution, allowing us to shed light on general effects on all buyers and sellers. Moreover, we are able to study the dynamics of responses to policy changes encompassing both short term adjustments that shed light on the extent of transaction frictions in this market as well as medium term timing responses elucidating the efficacy of policies that seek to stimulate this key market during economic downturns. Finally, our paper is methodologically related to recent bunching approaches developed in the literatures on labor supply and taxable income responses (Saez, 2010; Chetty *et al.*, 2011; Kleven & Waseem, 2012), but applied to a different market that is arguably as important.

The paper is organized as follows. Section 2 presents our theoretical framework, section 3 describes the context and data, section 4 develops the bunching methodology, section 5 presents the empirical results, and section 6 concludes.

2 Theoretical Framework

2.1 A Competitive Model of Behavioral Responses to Transaction Taxes

This section sets out a static model of a competitive housing market that generates the key predictions that we investigate empirically. The theoretical framework is deliberately unrealistic in some dimensions as our goal is to build the most parsinomious model possible that is still general enough to demonstrate the key empirical effects. The next section generalizes these predictions to a housing market with matching frictions and price bargaining between buyers and sellers, a feature which is notably absent in the current framework. However, as section 2.2 shows, the same qualitative predictions emerge.

Agents choose whether or not to become homeowners (extensive margin) and how much housing to buy conditional on owning (intensive margin). Letting c denote units of a numeraire consumption good and h denote units of quality-adjusted housing stock, we consider the following parametrization of preferences

$$u(c,h) = c + \frac{A}{1+1/\alpha}\left(\frac{h}{A}\right)^{1+1/\alpha} - q \cdot 1(h>0), \tag{1}$$

where A, α are parameters characterizing housing preferences and q is a fixed cost of entering the owner-occupied market including both transaction costs (search costs, broker fees, etc.) and the utility from renting instead of owning. We will allow for heterogeneity in all of these parameters captured by a smooth density distribution $f(A,\alpha,q)$. The quasi-linear utility function conveniently eliminates income effects on housing demand as we will focus purely on the price effect on housing demand.

As a baseline, consider a flat transaction tax rate t on the value of housing purchased. Denoting the price per unit of housing by p and income by y, the budget constraint is given by

$$c + (1+t)\,ph = y. \tag{2}$$

Conditional on owning ($h > 0$), maximizing utility (1) with respect to the budget constraint (2) yields the following housing demand function

$$h^* = A\left((1+t)\,p\right)^{\alpha}, \tag{3}$$

where α is the price elasticity of housing demand. Indirect utility conditional on $h > 0$ and exclusive of the fixed cost q can be defined as $v((1+t)\,p, y) \equiv u(c^*, h^*) + q$, while indirect utility conditional on $h = 0$ is given by $u(y, 0) = y$. The agent then enters the owner-occupied housing market iff

$$q \leq v((1+t)\,p, y) - y \equiv q^* \tag{4}$$

Total housing demand is then given by

$$D((1+t)\,p) = \int_A \int_\alpha \int_0^{q^*} h^* f(A,\alpha,q)\, dq d\alpha dA. \tag{5}$$

We will be agnostic about the details of the supply side and denote housing supply by $S(p)$. The equilibrium condition $D((1+t)\,p) = S(p)$ determines the equilibrium price p as a function of $1+t$.

Now consider the introduction of a discrete jump Δt in the average transaction tax rate—a notch—at a cutoff property value. Denoting property value by $h_v \equiv ph$, the notched tax schedule can be written as $T(h_v) = t \cdot h_v + \Delta t \cdot h_v \cdot 1\left(h_v > \overline{h}_v\right)$ where $\overline{h}_v$ is the cutoff and $1(\cdot)$ is an indicator for being above the cutoff. Figure 1 illustrates the implications of this notch in a budget set diagram (Panel A) and density distribution diagrams (Panels B-D). The budget set diagram (depicted in (h_v, c)-space) illustrates intensive responses among individuals with heterogeneous housing preferences A, but a specific demand elasticity α. The notch creates bunching at the cutoff $\overline{h}_v$ by all individuals in a preference range $\left(\overline{A}, \overline{A} + \Delta\overline{A}\right)$, who would have bought houses on the segment

$\left(\overline{h}_v, \overline{h}_v + \Delta\overline{h}_v\right)$ in the absence of the notch. The marginal bunching individual at $\overline{A}+\Delta\overline{A}$ is indifferent between the notch point $\overline{h}_v$ and the best interior location $\bar{h}_v^I$. No individual is willing to locate between $\overline{h}_v$ and $\bar{h}_v^I$, and hence this range is completely empty. The density distribution of property values corresponding to the budget set diagram (all A, one specific α) is shown in Panel B. Since the behavioral response in Panels A-B depends on the size of the demand elasticity α (and converges to zero for completely price inelastic buyers), the density distribution in the full population (all A, α) can be illustrated as in Panel C where some individuals are willing to buy just above the notch point.[5]

As shown by Kleven & Waseem (2012), the relationship between bunching and the demand elasticity can be characterized by considering the marginal bunching individual who is indifferent between the notch point and her best interior location. This indifference condition along with the first-order condition for the no-notch location $\overline{h}_v+\Delta\overline{h}_v$ implies a relationship $\Delta\overline{h}_v/\overline{h}_v = k\left(\alpha, \Delta t/\left(1+t\right)\right)$ where $k\left(\cdot\right)$ is monotonically increasing in both arguments.[6] Conversely, given the width of the bunching segment $\Delta\overline{h}_v$ (the estimation of which will be described later) and the tax parameters $\overline{h}_v$ and $\Delta t/\left(1+t\right)$, this condition gives a unique demand elasticity α.

In addition to intensive responses, the notch creates extensive responses above the cutoff by individuals close to being indifferent between buying and not buying (with $q \approx q^*$). However, such extensive responses will be negligible *just* above the cutoff. This can be seen by considering an individual who prefers a location on the segment $\left(\overline{h}_v, \overline{h}_v + \Delta\overline{h}_v\right)$ without the notch and therefore prefers the cutoff $\overline{h}_v$ with the notch (conditional on buying). For such an individual, the change in the threshold fixed cost Δq^* induced by the notch is given by

$$\Delta q^* = u\left(\overline{c}, \overline{h}_v/p\right) - u\left(c^*, h^*\right), \tag{7}$$

where $\overline{c}, \overline{h}_v/p$ is the consumption bundle obtained by the bunching individual in the presence of the notch. As the preferred point absent the notch h^* converges to the cutoff $\overline{h}_v/p$ from above (and hence c^* converges to $\overline{c}$), Δq^* converges to zero and extensive responses disappears. Intuitively, if in the absence of the notch, an individual would choose to buy a house slightly above $\overline{h}_v$, then in the presence of the notch, she will be better off by buying a house at $\overline{h}_v$ (which is almost as good) rather than not buying at all. This reasoning implies that extensive responses affect the density distribution as illustrated in Panel D of Figure 1.

[5] Notice that the tax-induced change in housing demand affects the equilibrium price p, which by itself will shift indifference curves in Panel A (as they are depicted in (h_v, c)-space) and hence shift the density distribution of property values. The qualitative characterization above holds for any arbitrary price and therefore also for the new equilibrium price. The key insight is that, in this competitive model, price incidence occurs at the *market level* and therefore does not contribute to bunching and holes locally around notches. The next section considers a bargaining model where price incidence occurs at the *match level* in which case price incidence does create bunching and holes.

[6] Under the specific parametrization in (1), the relationship $\Delta\overline{h}_v/\overline{h}_v = k\left(\alpha, \Delta t/\left(1+t\right)\right)$ is implicitly defined by the following condition

$$\frac{1}{1+\Delta\overline{h}_v/\overline{h}_v} - \frac{1}{1+1/\alpha}\left[\frac{1}{1+\Delta\overline{h}_v/\overline{h}_v}\right]^{1+1/\alpha} - \frac{1}{1+\alpha}\left[1+\frac{\Delta t}{1+t}\right]^{1+\alpha} = 0. \tag{6}$$

Empirically, we uncover intensive responses (changes in house values conditional on trading) using bunching at transaction tax notches and extensive responses (changes in the number of houses traded) using a tax reform that temporarily eliminated transaction taxes in a certain price range. In a one-period model where buying a house and owning a house are indistinguishable, the extensive margin corresponds to the choice between owning and not owning (renting). In a more realistic dynamic model, the extensive margin would include both the choice between owning or renting and the choice by incumbent owners to buy a different house or stay in their current house (the so-called lock-in effect). However, the essence of the effects of transaction taxes at the intensive and extensive margins carry over to a dynamic setting.

2.2 A Bargaining Model of Price Responses to Transaction Taxes

A key feature of the competitive housing market model presented above is that excess bunching and holes around notch points reflect real demand responses (as opposed to price incidence) and therefore reveal the elasticity of real housing demand. This section shows that the same qualitative effects on the house price distribution can be generated by bargaining between buyers and sellers in a model with matching frictions. In this model, bunching responses reveal the bargaining power of buyers versus sellers.

Consider a specific match where the buyer has valuation B_v and the seller has valuation S_v of the property. Considering a flat transaction tax t (remitted by the buyer), the buyer's surplus from trading at the before-tax house price h_v is equal to $B_v - (1+t)\, h_v$ and the seller's surplus is equal to $h_v - S_v$. The necessary and sufficient condition for a trade to take place is that there exists a price such that both traders obtain a positive surplus, i.e. we must have $S_v \leq \frac{B_v}{1+t}$.

The buyer and seller engage in Nash bargaining with bargaining power β for the buyer and $1-\beta$ for the seller. The agreed before-tax price h_v maximizes $W = [B_v - (1+t)\, h_v]^{\beta} \, [h_v - S_v]^{1-\beta}$, which yields

$$h_v = \beta S_v + (1-\beta)\, \frac{B_v}{1+t}. \tag{8}$$

Hence, conditional on trading, the transaction tax reduces the house price h_v, with the strength of the price effect being proportional to the bargaining power of the seller $1-\beta$. This means that we can characterize extensive and intensive responses to the transaction tax t in the following way. House transactions that were desirable to the buyer and seller in the absence of transaction taxes but sufficiently close to the indifference margin for both $(B_v/\,(1+t) < S_v \leq B_v)$ will no longer occur (extensive response). House transactions that continue to be desirable in the presence of transaction taxes $(S_v \leq B_v/\,(1+t))$ will occur at lower prices according to equation (8) (intensive response). Assuming a smooth distribution of matches S_v, B_v and bargaining power β, captured by a density distribution $f\,(S_v, B_v, \beta)$, there will be a smooth distribution of traded house prices under the flat transaction tax t.

Consider now the introdution of a notch Δt in the transaction tax at the cutoff house price $\overline{h}_v$. Under the notched tax schedule and Nash bargaining between the buyer and seller, the agreed house price h_v is picked to maximize

$$W = \left[B_v - \left(1 + t + \Delta t \cdot 1\left(h_v > \overline{h}_v\right)\right) h_v\right]^{\beta} \left[h_v - S_v\right]^{1-\beta}. \tag{9}$$

In general, solving this bargaining problem requires us to solve for the best price point within each tax bracket (below and above $\overline{h}_v$) and then pick the candidate solution that yields the largest welfare W. Trades that would occur below $\overline{h}_v$ under the baseline flat tax are clearly unaffected by the notch and continue to feature house prices given by (8). On the other hand, trades that would occur above $\overline{h}_v$ under the baseline flat tax are affected by the notch. To see how these trades are affected, note first that any trade occurring stricly above the cutoff must satisfy the interior pricing condition (8) with the $1 + t$ replaced by $1 + t + \Delta t$. This allows us to distinguish between three cases.

First, transactions just above $\overline{h}_v$ under the baseline tax rate t would have an interior solution below $\overline{h}_v$ under the larger tax rate $t+\Delta t$ (based on eq. (8) at tax rate $1+t+\Delta t$). This is inconsistent with an interior solution in either bracket, and so these transactions bunch at the cutoff. Second, transactions that would be just above $\overline{h}_v$ under an interior solution at the new tax rate $t+\Delta t$ (again based on eq. (8) at tax rate $1+t+\Delta t$) also bunch at the cutoff. For such transactions, a small move to the cutoff provides a discrete gain to the buyer and only a marginal loss to the seller, yielding a larger value of W than at the interior location. Given a smooth distribution of matches (S_v, B_v), there will be a marginal bunching transaction such that welfare at the cutoff $\overline{h}_v$ is precisely equal to welfare at the best interior location above the notch $\bar{h}_v^I$. No transaction takes place in the interval $\left(\overline{h}_v, \bar{h}_v^I\right)$, and so we get a hole in the price distribution there. The width of this hole depends on bargaining power and converges to zero as the bargaining power of the buyer β converges to zero. Third and finally, transactions above $\bar{h}_v^I$ under an interior solution at the new tax rate $t + \Delta t$ are associated with a larger W at the new interior solution than at the cutoff. For those transactions, we get a downward price shift within the upper bracket.

The characterization above is analogous to the characterization for the competitive market model, with the bargaining power parameter β in the bargaining model playing the role of the demand elasticity α in the competitive model. A graphical illustration similar to Figure 1 is also possible. Figure 2 shows the direct analog of panel A of Figure 1 for the case of the bargaining model, and shares all of its qualitative features. The density diagrams in panels B-D of Figure 1 can also be reinterpreted in terms of the bargaining model, with Panel B depicting the effect on the house price distribution for a given β (and a distribution of matches S_v, B_v) and Panel C depicting the effect on the house price distribution for the full distribution of βs (some of which will have $\beta = 0$ and locate just above the notch). Finally, this characterization applies only to matches for which a trade is still beneficial. The notch will also create extensive responses above the cutoff as house transactions that were desirable to the buyer and seller under the flat tax but close enough to the indifference margin for both $(B_v/(1 + t + \Delta t) < S_v \leq B_v/(1 + t))$ will no longer take place.

3 Context and Data

3.1 Taxation of Property Transactions in the UK: Notches and Reforms

The UK property transaction tax—Stamp Duty Land Tax (SDLT)—is imposed on anything of economic value, money as well as transfers of debts or any other works and services, which is given in exchange for land or property, known as the "chargeable consideration".[7] The statutory incidence of the SDLT falls on the buyer, who is required to file a stamp duty return and remit tax liability to HMRC within a few weeks of the completed transaction. The SDLT has become a significant source of government revenue in the UK, much more so than other wealth transfer taxes such as inheritance taxation and capital gains taxation. The SDLT raises revenue of around 0.6% of GDP and the political debate in the UK suggests that future rates (on highly priced properties) are more likely to go up than down.

A central aspect of the stamp duty is that it features discrete jumps in tax liability—notches—at threshold property prices. Tax liability is calculated as a proportional tax rate times the transacted property price, with different tax rates in different price brackets. Hence, as the purchase price crosses a bracket threshold, a higher tax rate applies to the entire amount and not just the portion that falls above the cutoff as in standard graduated schedules. Figure 3 illustrates the stamp duty schedule for residential property in tax year 2010-11 (the last year of our data).[8] The schedule features three notches as the proportional tax rate jumps from zero to 1% at a price of £125,000 (a tax liability jump of £1,250), from 1% to 3% at a price of £250,000 (a tax liability jump of £5,000), and finally from 3% to 4% at a price of £500,000 (a tax liability jump of £5,000).[9] Schedules are different for residential property purchased by first-time buyers or in certain disadvantaged areas (where the first bracket threshold is at a higher price) as well as for non-residential property. It should also be noted that stamp duty schedules are not indexed for inflation, which creates "bracket creep" as property price inflation pushes houses into higher stamp duty brackets. This was particularly important during the housing boom that lasted until 2007.

Another important aspect of the stamp duty is that it has been subject to a great deal of policy experimentation over the years. During our data period, the location of the first notch point has moved several times as shown in Table 1. The most important change was the so-called *stamp duty holiday* between 3 September 2008 and 31 December 2009, which moved the first notch point from £125,000 to £175,000 and thereby eliminated stamp duty in a £50,000 range. Two features of the stamp duty holiday are important for our analysis. First, the beginning of the holiday was *unanticipated* as it was announced suddenly by the then Chancellor Alistair Darling on the day before its introduction. Although there was some media speculation about the possibility of

[7] The chargeable consideration includes the buildings and structures on the land as well as fixtures and fittings (such as in bathrooms and kitchens) but excludes freestanding furniture, carpets or curtains. If such extras are included in the sale, the buyer and seller are to agree on the market value of these extras and subtract it from the chargeable consideration. See http://www.hmrc.gov.uk/sdlt/calculate/value.htm for details.

[8] The UK tax year for personal taxes runs from April 6 in one year to April 5 the next year.

[9] After our data period, additional notches have been implemented at £1,000,000 (a tax liability jump of £10,000) from tax year 2011-12 and at £2,000,000 (a tax liability jump of £40,000) from tax year 2012-13.

a stamp duty holiday in the month leading up to the announcement, the details and start date of such a holiday were unknown. Second, the end of this holiday was *anticipated*. The initial announcement was that the holiday would last for one year (until September 2009), but in April 2009 this was extended until the end of 2009 and the government committed to no further extensions (and indeed did not grant any extensions). The sudden announcement of the stamp duty holiday and the preannounced committment to its end date allow us to compare the effects of expected and unexpected changes in tax policy. In particular, the pre-announced end date creates a *time notch* in a specific price range allowing us to analyze short-term timing effects. Furthermore, as the stamp duty holiday applied only to properties in a certain price range, we are able to study the stimulus effects of the policy by looking for increased activity in this price range – its intended effect.

The UK stamp duty land tax is characterized by strong enforcement and high compliance. According to HMRC estimates, the so-called *tax gap*—the difference between true taxes owed and actual taxes paid on a timely basis—in the stamp duty land tax is 4.8% of true tax liability in the tax year 2010-11. This is lower than the tax gap estimates for most other taxes in the UK. It is perhaps not surprising that tax evasion is a minor issue for this tax when considering the following points. First, almost all property transactions in the UK are facilitated by licensed real estate agencies, implying that stamp duty tax evasion requires collusion between a buyer, a seller and a real estate agency (typically with multiple employees). Such evasion collusion involving many agents is unlikely to be sustainable (Kleven *et al.*, 2009). Second, the scope for tax evasion is further reduced by the existence of a considerable lag between agreeing on a house price and completing the contract.[10] If the house price reported to tax authorities is lower than the true house price, the buyer must make a side payment to the seller. If the buyer makes the side payment at the time of agreeing on the house price, the seller would be able renege before completing the contract and it would be difficult for the buyer to recoup the payment. If instead the buyer promises to make the payment at the time of completing the contract, the seller would take his property off the market with no credible commitment from the buyer that he would not renege later when the bargaining position of the seller may be weaker. Hence, such side payments would be associated with substantial risk for either the buyer or the seller or both. Finally, the tax base is defined in an extremely comprehensive manner meaning that the scope for misvaluation and/or misclassification of specific features of the property to avoid the tax is extremely limited. For these reasons, we believe that house prices reported on stamp duty tax returns reflect true house prices in the overwhelming majority of cases.

3.2 Data

The empirical analysis is based on administrative data covering the universe of stamp duty (SDLT) returns in the UK from April 2005 to March 2011. Since most property transactions require the filing of an SDLT return (the main filing exemption being for property transactions under £40,000), our data is close to the universe of property transactions in the UK. The full dataset contains almost 9 million transactions. The dataset contains rich tax return information for each transaction, but

[10] This lag is about 2 months on average in the UK housing market (Besley *et al.*, 2011).

no information outside the tax return. The variables that we currently use are the following: date of transaction, sales price, type of property (residential, non-residential or mixed use), geographical area, and whether any stamp duty relief applies (such as the special schemes that reduce stamp duty on some transactions by first-time buyers and in disadvantaged areas).

The housing market has seen substantial turmoil during the period we consider. Figure 4 shows the monthly number of house transactions (Panel A) and the monthly average property price (Panel B) in all of the UK and in London alone. The figure shows nominal prices (real prices give the same qualitative picture) and normalizes both the price and the number of transactions to one at the start of the period. We make the following observations. First, housing market activity collapses between 2007 and 2009 as the the number of transactions falls by almost two-thirds. There has been some recent recovery, but activity is still very far from pre-recession levels. Second, property prices also fall between 2007 and 2009, but the price drop is less dramatic and the subsequent recovery much stronger. Third, property prices (though not activity) in London have evolved differently than in the rest of the UK during the recession. While UK-wide property prices have recovered only partially in the past couple of years, London property prices are back on their pre-recession trend. Finally, the recovery in house prices and activity throughout 2009 coincides with the stamp duty holiday, which has been used as an argument that the policy had the desired effect. We will take a quasi-experimental approach to evaluate how much of the recovery (if any) can indeed be explained by the stamp duty holiday.

4 Bunching Methodology

As elucidated in the theoretical frameworks of section 2, we expect a transaction tax notch at the cutoff property price $\bar{h}_v$ to induce excess bunching at the cutoff by properties that would have been sold at prices between $\bar{h}_v$ and $\bar{h}_v + \Delta\bar{h}_v$ absent the notch. In the competitive model of section 2.1 this effect was driven by real responses governed by the demand elasticity α, while in the bargaining model of section 2.2 the effect was driven by price incidence governed by the bargaining power β of buyers relative to sellers. In both cases, these effects generated an excess mass of

$$B\left(\bar{h}_v\right) = \int_{\bar{h}_v}^{\bar{h}_v+\Delta\bar{h}_v} g_0\left(h_v\right) dh_v \approx g_0\left(\bar{h}_v\right)\Delta\bar{h}_v, \tag{10}$$

where $B\left(\bar{h}_v\right)$ is excess mass at the cutoff and $g_0\left(h_v\right)$ is the counterfactual density of house values (i.e. the density that would prevail absent the notch). The approximation is accurate to the extent that the counterfactual is approximately uniform around the notch. Based on equation (10), it is possible to recover the house price response $\Delta\bar{h}_v$ based on estimates of the counterfactual distribution $g_0\left(h_v\right)$ and bunching $B\left(\bar{h}_v\right)$.

The relationship (10) implicitly assumes that there is just one bunching segment $\left(\bar{h}_v, \bar{h}_v + \Delta\bar{h}_v\right)$, which amounts to an assumption that the underlying driver of price response (demand elasticity α or bargaining power β) is homogeneous in the population. Our conceptual framework allows for heterogeneity in responses and we can also account for it in the empirical implementation. Denoting the underlying source of heterogeneity by $x = (\alpha, \beta)$, there will be a price response $\Delta \bar{h}_v(x)$ and a counterfactual density $\tilde{g}_0(h_v, x)$ associated with each type x. In this case, equation (10) can be generalized to

$$B\left(\bar{h}_v\right) = \int_x \int_{\bar{h}_v}^{\bar{h}_v + \Delta \bar{h}_v(x)} \tilde{g}_0(h_v, x)\, dh_v dx \approx g_0\left(\bar{h}_v\right) E\left[\Delta \bar{h}_v(x)\right], \tag{11}$$

where $E\left[\Delta \bar{h}_v(x)\right]$ is the average price response across all x. As before, the approximation requires that the counterfactual density is locally uniform in house prices h_v (but not type x) around the notch point. Equation (11) shows that estimates of the counterfactual distribution and bunching allows us to recover the average house price response in the population.

We now describe how the counterfactual density distribution and excess bunching are estimated. Following Chetty *et al.* (2011) and Kleven & Waseem (2012), we estimate the counterfactual density by fitting a high-order polynomial to the observed density, excluding data in a range around the notch. This procedure will yield a valid estimate of the counterfactual distribution under the identifying assumption that all other factors determining the price are distributed smoothly around the notch. That is, it must be the case that nothing else that determines the sale price jumps discontinuously at the notch. A slight wrinkle with this method is the presence in the data of bunching at round numbers. That is, transactions tend to take place at salient, round numbers so that neither the observed nor the counterfactual distribution is completely smooth, but rather features spikes at round numbers. To account for this, we extend the estimated polynomial to include round-number dummies as in Kleven & Waseem (2012).

Grouping transactions into price bins of £100, the regression used to estimate the counterfactual distribution around a notch at price $\bar{h}_v$ is given by

$$c_i = \sum_{j=0}^{q} \beta_j (z_i)^j + \sum_{r \in \mathcal{R}} \eta_r I\left\{\frac{\bar{h}_v + z_i}{r} \in \mathbb{N}\right\} + \sum_{k=\bar{h}_v^-}^{\bar{h}_v^+} \gamma_k I\{i = k\} + \mu_i, \tag{12}$$

where c_i is the number of transactions in price bin i, q is the order of the polynomial, and z_i is the distance between the upper bound of price bin i and the cutoff $\bar{h}_v$. The second term in (12) includes fixed effects for prices that are multiples of the round numbers in the set $\mathcal{R}$, where $\mathcal{R} = \{500, 1000, 5000, 10000, 25000\}$, $\mathbb{N}$ is the set of natural numbers, and $I\{\cdot\}$ is an indicator function. Finally, the third term in (12) excludes a region (h_v^-, h_v^+) around the notch that is distorted by bunching responses to the notch, and μ_i is a residual reflecting misspecification of the density equation.[11] Our estimate of the counterfactual distribution is defined as the predicted bin counts $\hat{c}_i$ from (12) omitting the contribution of the dummies in the excluded range, and excess bunching is estimated as the difference between the observed and counterfactual bin counts in the part of the excluded range that falls below the notch $\hat{B}\left(\bar{h}_v\right) = \sum_{i=\bar{h}_v^-}^{\bar{h}_v} (c_i - \hat{c}_i)$. We may also define

[11] Note that the residual does not reflect sampling error as our data covers the entire universe of property transactions.

an estimate of missing mass (the hole) above the notch as $\hat{M}\left(\bar{h}_v\right) = \sum_{i>\bar{h}_v}^{\bar{h}_v^+}\left(\hat{c}_i - c_i\right)$.[12] Standard errors on these estimates are calculated based on a bootstrap procedure as in Chetty *et al.* (2011) and Kleven & Waseem (2012).

The estimation of equation (12) relies on choosing the excluded range (h_v^-, h_v^+) and the polynomial degree q. As with previous bunching approaches, visual inspection of the empirical distribution will guide those choices. First, because bunching and holes around notch points are very clear in our data, the excluded range can be determined visually without much ambiguity. The lower bound $\bar{h}_v^-$ is set at the point where excess bunching starts, while the upper bound $\bar{h}_v^+$ is set at the point where the hole ends (i.e., where bunching responses to the notch stops). Because the empirical distribution considered below is declining outside regions close to notches, the end point of the hole may be defined as the point where the observed density above the notch changes slope from positive to negative (apart from spikes at round numbers). This corresponds to the situation depicted in the theoretical Figure 1 (Panels C and D). Second, for the polynomial degree, we have experimented with degrees between 4 and 7 to allow for a great deal of flexibility. By and large the estimates are not sensitive to this choice, but in certain instances it matters. In particular, high degrees like 6 and 7 are occassionally *too* flexible and it is visually clear that counterfactual is not right.[13]

We use a similar bunching methodology to study the time notch created by the anticipated tax increase at the end of the stamp duty holiday on 31 December 2009. As described in section 3.1, properties sold at prices between £125,000 and £175,000 paid an additional tax of £1,250-£1,750 if they were sold on 1 January 2010 as opposed to being sold on 31 December 2009. Analogously to the case of a price notch, we therefore expect there to be bunching in the distribution of transaction *dates* in this price range at or just before 31 December 2009 along with a hole in the distribution of transaction dates after 1 January 2010. To evaluate this, we compare the observed distribution of transaction dates to an estimate of the counterfactual distribution absent the time notch.

The counterfactual is estimated in weekly bins by fitting a polynomial to the observed distribution, excluding a time interval around the turn of the year 2009/10. Visual inspection of the observed date distribution reveals that there is a tendency for transactions to occur at the end of each month. This is much like bunching at round numbers in the price distribution, so our solution is the same: we include in our estimation equation a fixed effect for the last week of every month.

[12]This measure of missing mass is, in general, not exactly equal to the amount of excess bunching. This can be seen from Panel D of Figure 1, where the missing mass that corresponds to bunching is the light-grey area above the notch. The estimate $\hat{M}$ may be different from $\hat{B}$ due to extensive responses and/or the (small) shift in the distribution created by intensive responses of those who do not bunch. The measure of missing mass that corresponds precisely to bunching is the expression on the right-hand sides of equations (10) and (11).

[13]Implementing specification (12) is based on two additional choices that are less important. The first is the width of the bins used to estimate the density. We have used a bin width of £100, but have also experimented with bin widths of £500 and £1,000 and the results are not sensitive to this. The other is the set of round numbers in $\mathcal{R}$. Here, we have chosen to a fairly rich set of salient round numbers to ensure that we are able to fully capture the anatomy of rounding behavior in the data.

The estimation equation is given by

$$c_w = \sum_{j=0}^{q} \beta_j (z_w)^j + \eta I \{w \in \text{end of month}\} + \sum_{k=\bar{w}^-}^{\bar{w}^+} \gamma_k I \{w = k\} + \mu_w, \tag{13}$$

where c_w is the number of transactions in week w and z_w is the distance of week w from the end of 2009. The second term is the fixed effect for weeks at the end of the month, while the third term excludes weeks in a range $(\bar{w}^-, \bar{w}^+)$. In the empirical analysis below, the excluded range is set to include the last 3 weeks of 2009 and the first 10 weeks of 2010. Our estimate of excess mass in transaction acticty just before the notch is then $\hat{B}(\bar{w}) = \sum_{w=\bar{w}^-}^{\bar{w}} (c_w - \hat{c}_w)$, where $\bar{w}$ is week 52 of 2009 and $\hat{c}_w \equiv \sum_{j=0}^{q} \hat{\beta}_j (z_w)^j$. We do not include the end-of-the month fixed effect $\hat{\eta}$ in the counterfactual $\hat{c}_w$ just before the notch (December), because in that particular month the rush is usually in the penultimate week of the year in the lead up to Christmas. Instead, to deal with the Christmas rush, we compare bunching estimates in the treated group (transactions in the affected price range and year) to bunching estimates in control groups (other price ranges and years)—a difference-in-bunching strategy—as described in section 5.3.

5 Empirical Findings

5.1 Static Price Notches: Bunching and Holes

This section presents static results using price notches during periods when they are stable. We consider residential property transactions that do not qualify for stamp duty relief since a different tax schedule applies to that class of transactions.[14] Figure 5 considers the two upper notches located at cutoff prices of £250,000 (Panel A) and £500,000 (Panel B), both of which have remained in place throughout the period of our data. Each panel shows the empirical distribution of house values (blue dots) as a histogram in £5,000 bins and the counterfactual distribution (red line) estimated as a fifth-order polynomial with round-number fixed effects as in specification (12). The excluded range (h_v^-, h_v^+) in the estimation of the counterfactual is demarcated by vertical dashed lines; the lower bound h_v^- is set at the point where excess bunching starts and the upper bound h_v^+ is set at the point where the hole ends (where the empirical distribution above the cutoff changes slope from positive to negative). Each panel shows estimates of excess bunching below the notch scaled by the counterfactual frequency at the notch (b), the size of the hole (missing mass) above the notch scaled by the counterfactual frequency at the notch (m), the difference between these two ($m - b$), the average house value change created by the notch, and the tax liability change at the notch.

Our main findings are the following. First, both notches create large and sharp bunching below the cutoff. Excess bunching is 1.32 and 0.91 times the height of the counterfactual distribution at £250,000 and £500,000, respectively, and is strongly significant in each case. Second, both notches are associated with a large hole in the distribution above the cutoff. The size of the hole is

[14] Results for non-residential property are qualitatively similar, but more noisy due to much fewer observations.

larger than the size of excess bunching, although the difference between the two is not statistically significant from zero. Third, the hole in the distribution spans a £25,000 range above each cutoff, implying that the most responsive agents reduce their transacted house value by five times as much as the jump in tax liability of £5,000. Finally, the average house value response is £6,610 at the £250,000 notch and £4,562 at the £500,000 notch. That is, the transaction tax notches reduce the average house price by about the size of the tax jump.

We now turn to the lower notch, the location of which has changed several times during the period under consideration. The cutoff was located at £120,000 until 22 March 2006, at £125,000 between 23 March 2006 and 2 September 2008, at £175,000 between 3 September 2008 and 31 December 2009, and again at £125,000 from 1 January 2010 onwards. This section takes a static approach by considering bunching responses within each of these four periods separately, while the next section investigates dynamic adjustment paths around the reform episodes. Figure 6 shows results for the four periods in separate panels, each of which is constructed as in the preceding figure. The findings for the lower notch are qualitatively consistent with those for the upper notches, with a clear and statistically significant bunching response to the tax notch in each period. The size of the bunch and the hole is smaller at the lower notch than at the upper notches, but so is the size of the notch. The effect of the notch on the average transacted house value is between £2,700 and £4,200, or about 2–3 times the size of the tax liability jump, which is a qualitatively larger effect than at the upper notches.

Reported house values in our data can be described by $h_v \equiv p \cdot h - e$, where p is the real price per unit of quality-adjusted housing, h is the amount of quality-adjusted housing, and e is stamp duty evasion. This means that, in general, our estimates of reported house value responses conflate real price changes (incidence), real demand changes and tax evasion. This is conceptually similar to the way in which the literature on taxable income responses (e.g., Saez *et al.*, 2012) conflates real wage changes, real labor supply changes and tax evasion. For reasons discussed in section 3.1, tax evasion is not a very significant issue in the UK stamp duty, and we can therefore think of our estimates as combining mostly price incidence and real house demand.[15] Given that the price incidence effect is bounded by the size of the tax change, the fact that we find house value responses that are larger than the jumps in tax liability implies that there must be real demand responses as well. What is more, the size of the real response can be bounded by considering the extremes of zero incidence on the buyer (incidence change in house value equals tax liability change) and full incidence on the buyer (incidence change in house value equals zero).

[15] In the competitive housing market model in section 2.1, bunching is driven purely by real demand responses Δh as price incidence effects Δp shift the entire distribution and do not create bunching by themselves. Conversely, in the model with matching frictions and price bargaining in section 2.2, bunching is driven purely by price bargaining effects. It is possible to extend the bargaining model to include real responses in which case bunching can be driven by a combination of price incidence and real effects.

5.2 Moving Price Notches: Dynamics of Bunching

This section investigates the dynamics of behavioral adjustment to the changes in the position of the lower notch that were mentioned above. When considering dynamic adjustments, it is important to keep in mind that there is always a lag between agreeing on a purchase price and completing the housing contract. In the UK housing market, this lag is under 90 days for most transactions and about 60 days on average (Besley *et al.*, 2011). Since the official transaction date in our data refers to contract completion, the time it takes for the market to settle into a new equilibrium is bounded from below by about 3 months.

Figure 7 considers the movement of the lower notch from £120,000 to £125,000 on 23 March 2006. Each panel shows the empirical and counterfactual distributions in a given month between February 2006 and September 2006. The two vertical lines demarcate the £120,000 and £125,000 cutoffs and are either solid green (for the cutoff that applies to the month in question) or dashed black (for the cutoff that does not apply). April 2006 is the first full month where the new cutoff is in place. The figure shows very clearly how the bunch moves over time in response to the changed location of the notch. Most of the adjustment has occurred after four months (in July 2006) and a new equilibrium has been reached after 6 months (in September 2006). Hence, most of the lag in the adjustment to the new equilibrium can be explained by the administrative lag between contract exchange and contract completion.

The next three figures consider the movement of the lower notch from £125,000 to £175,000 on 3 September 2008 (the start of stamp duty holiday) and the subsequent movement back to £125,000 on 1 January 2010 (at the end of stamp duty holiday). When interpreting the findings, it is worth keeping in mind that the start of the holiday was unanticipated while the end of the holiday was anticipated (see section 3.1). Figure 8 shows monthly bunching graphs over a 12-month period around the beginning of the holiday. It is constructed like the preceding figure, except that we now add estimates of excess bunching b around the two cutoffs in each month. The main findings are the following. First, it takes 3-4 months for bunching at the old £125,000 cutoff to disappear (bunching becomes statistically insignificant for the first time in December 2008), corresponding roughly to the lag between contract agreement and completion. Second, it takes about 3 months for bunching at the new £175,000 cutoff to build up and reach a steady state (bunching b is around 0.9 from November 2008 onwards). Third, although bunching at £175,000 in the winter months of 2008/09 is smaller in absolute terms than bunching at £125,000 in the summer months before the holiday, bunching in proportion to the counterfactual distribution (b)—the right measure of responsiveness—is in fact slightly larger at £175,000. The presence of smaller absolute bunching at £175,000 is a result of seasonality in the housing market with fewer house transactions in the winter than in the summer.[16] The presence of larger relative bunching b at £175,000 is consistent with the fact that this notch is larger than the previous one at £125,000 (tax liability jumps of £1,750 and £1,250 respectively).

[16]Seasonality in the housing market is a well-known phenomenon that has been studied in the macro literature (e.g. Ngai & Tenreyro, 2012).

Figure 9 turns to the 12-month period around the end of the holiday on 1 January 2010 and is constructed exactly as the preceding figure. It is interesting to see the difference in the speed of adjustment to a tax change that is fully anticipated. First, the bunching at £175,000 vanishes immediately in January of 2010 when this cutoff is no longer a notch point. This shows that traders did indeed anticipate the end of the holiday and made sure to complete their housing contracts before the end of December 2009. We see such behavior in the graph for December 2009: there is a large upward shift in the December distribution between £125,000 and £175,000 (even though this is normally a low-season month) and an increase in excess bunching at £175,000. The next section investigates such short-term timing behavior in greater detail. Second, it takes about 2 months for bunching at the new £125,000 cutoff to build up and reach a stable equilibrium (b is roughly constant from February 2010 onwards). While this is faster adjustment than at the start of the holiday, it is not as fast as the disappearance of bunching at the end of the holiday. The implication is that, while traders were rushing to *complete* agreed housing contracts below the the £175,000 notch just before the end of the holiday (immediate disappearance of old bunching), they did not to the same degree *agree* (but not complete) housing contracts below the £125,000 notch just before the end of the holiday (slower emergence of new bunching).

Figure 10 summarizes the evidence in the preceding figures by showing the monthly bunching estimate b from January 2007 to January 2011 at the £125,000 cutoff (blue dots) and the £175,000 cutoff (orange crosses) with 95% confidence intervals around each series. The solid vertical lines demarcate the beginning and end of the stamp duty holiday, while the dashed vertical line demarcates the de facto time at which the holiday took full effect given the lag between agreed and completed house purchases. The figure highlights just how sharply house prices react to tax notches and to changes in tax notches even at the monthly level. The level of bunching at the £125,000 cutoff is remarkably constant on each side of the holiday, while the level of bunching at the £175,000 cutoff is constant during the holiday. The steady state level of bunching at £175,000 ($b \approx 0.9$) is larger than at £125,000 ($b \approx 0.6$) as the former notch is larger. Once we account for the built-in sluggishness due to the time it takes to complete a housing contract, the market adjusts to a new stable equilibrium remarkably quickly. We also do not see any difference in price responsiveness during good times and bad times (compare early part of 2007 to the rest of the period). Compared to recent bunching evidence from labor markets (e.g. Saez, 2010; Chetty *et al.*, 2011; Kleven & Waseem, 2012), the remarkable sharpness of our evidence suggests that behavioral responses in the housing market are much less affected by optimization frictions such as inattention, inertia, etc. Our evidence suggests that agents in the housing market respond precisely and quickly to tax incentives.

5.3 Time Notch: Short-Term Timing Effects

This section analyzes timing responses in house purchases, using the time notch created by the anticipated jump in the transaction tax rate on 1 January 2010 for houses between £125,000 and £175,000. Before discussing the empirical results, we make two remarks. First, a notch on 1 January is effectively a notch just before Christmas due to the fact that the housing market almost shuts down between Christmas and New Year. Hence, agents should respond to the notch by moving the date of purchase from the early weeks of 2010 to the third week of December 2009. Second, the existence of the Christmas holiday (with or without a tax notch) may in itself lead to a piling up of house transactions in the third week of December. This means that we cannot analyze the time notch using a "pure" bunching strategy as observed bunching in transactions before Christmas 2009 may overstate the response to the tax notch. We therefore pursue a difference-in-bunching strategy by comparing bunching in the treated group (transactions between £125,000–£175,000 in December 2009) to bunching in control groups (other years and/or other price ranges).

Figure 11 shows the weekly number of transactions around New Year in different price ranges and different years. The top panels consider the difference-in-bunching analyses of interest, while the bottom panels show placebo difference-in-bunching analyses. In particular, Panel A compares the treated price range £125,000–£175,000 in the treated period 2009/10 to the surrounding price ranges in same period. The treated group features very strong bunching just before the notch and a large hole after the notch. The control groups also feature bunching and a hole (Christmas effect), but to a much smaller extent. Furthermore, the shutdown of activity between Christmas and New Year is less extreme in the treated group than in the control groups.

To evaluate the timing response, we estimate excess bunching in each distribution during the last three weeks of the year based on the difference in density mass between the empirical distribution and a counterfactual distribution (obtained from specification (13)), scaled by the average height of the counterfactual distribution in the bunching range. The timing response is then given by the difference between bunching in the treated range and average bunching in the surrounding control ranges (D-i-Bunching in the figure). We find that excess mass induced by the time notch is almost 3 times the height of the counterfactual and strongly significant, implying that the average timing response to the notch is 3 weeks. Panel B is constructed in the same way, except that it compares the treated price range £125,000-£175,000 in the treated period 2009/10 to the same price range in other periods (one year earlier or two years earlier). The results are very similar, with estimated excess mass before the notch being somewhat larger and still strongly significant. Finally, the placebo tests in the bottom panels repeat the strategy in Panel A (comparing different price ranges), but one year or two years earlier. In each case, the timing effect is close to zero and statistically insignificant. Overall, this provides very compelling evidence of timing responses to notches, and the findings are consistent with the sharpness of price responses to notches that we discussed above.

5.4 Stamp Duty Holiday: Housing Stimulus Effects

The previous analysis shows that agents respond to the price and time notches introduced by the stamp duty holiday by changing house prices and transaction dates in the expected fashion. This is very useful for the understanding of behavior in the housing market, but it does not directly shed light on the key macro question: did the stamp duty holiday have the desired effect of stimulating housing market activity during the recession? This question relates not just to the specific housing stimulus policy considered here, but to fiscal stimulus policy more generally. We now turn to this important issue.

The evidence is presented in Figure 12 in four panels, all of which pursue a difference-in-differences approach by comparing the evolution over time of trading volumes in the treated price range and the surrounding prices ranges. Panel A represents a "naive" baseline that we refine later. This panel compares the log monthly number of transactions in the treated range £125,000-£175,000 (blue dots) and the surrounding £50,000 ranges combined (orange crosses). We have normalized the log number of transactions in each month by the average log number of transactions in the pre-treatment period (the 12 months leading up to the holiday). The solid vertical lines mark the beginning and the end of the stamp duty holiday, while the dashed vertical line marks three months after the official beginning of the holiday to account for the built-in administrative inertia described above. The two series display completely parallel trends leading up to the holiday and then begin to diverge precisely when the holiday starts. The positive effect of housing stimulus in the treated range increases during the holiday, features a short-term spike at the end, and then reverses after the holiday as activity in the previously treated range falls below activity in the untreated ranges. Taken at face value, this graph implies that housing stimulus gave a *large boost* to housing market activity during the policy along with a subsequent *smaller slump* in activity when the policy was withdrawn. However, we argue that this overstates the positive impact of the stimulus policy.

The issue with the analysis in panel A is that treatment assignment (whether a transaction takes place in the £125,000–£175,000 price bracket) is endogenous to movements across bracket cutoffs. The stamp duty holiday creates an incentive to move into the treated price bracket from both sides. At the upper end of the bracket, the holiday creates a new notch at £175,000 that induces traders to move from a region above the cutoff to a point just below the cutoff (bunching). We have shown above that bunching responses at £175,000 do indeed occur, and this increases activity in the treated range compared to the surrounding ranges. At the lower end, the holiday eliminates the notch at £125,000 and therefore induces bunchers at this cutoff to move back into the hole above the cutoff. We have shown that the disappearance of bunching at £125,000 also occurs, and this further increases activity in the treated range compared to the surrounding ranges. Hence, the positive effect of housing stimulus in panel A combines the true effect on overall activity levels with endogenous price changes that are simply the result of the change in the location of the notch.

This endogeneity issue can be dealt with in two different ways. A crude way (Panel B) is to expand the treated range in order to ensure that any price manipulation around notches occurs *within* the treated range and so does not affect measured activity levels in the treated range. This requires us to increase the treated range to include the pre-holiday and post-holiday bunching range below the £125,000 cutoff (£115,000–£125,000 according to figure 6) as well as the within-holiday hole range above the £175,000 cutoff (£175,000–£195,000 according to figure 6). This strategy results in effects of housing stimulus that are qualititatively similar, but much smaller. While the boost in activity during the holiday is still visually clear, the effect is modest and it begins to emerge only a few months after the introduction of the policy. The delayed effect makes sense given that the reform was fully unanticipated and the time it takes to complete housing contracts. The simple empirical strategy in Panel B is likely to underestimate the effect of the reform as it includes untreated transactions in the treatment group (making it an intent-to-treat effect), particularly transactions just below £125,000 or just above £175,000 that do not respond to notches (or the elimination of notches) by moving into the treated price range. Hence, we conclude that while this panel has dealt with the endogenous treatment issue, this panel understates the effect of the stimulus policy.

Our preferred way of dealing with the selection into treatment, shown in panels C and D, is to exploit the fact that we have monthly estimates of price responses to notches and can therefore directly control for it. That is, we may consider the number of transactions in different price brackets adjusted for the effect of bunching behavior in each month. To be precise, the number of transactions in the £125,000-£175,000 range is adjusted by removing excess mass just below £175,000 and adding excess mass just below £125,000 in every month before, during and after the holiday (these adjustments increase activity in the treated range *outside* the holiday and reduce it *during* the holiday). The number of transactions in the surrounding price ranges is adjusted by adding excess bunching below £175,000 and removing excess bunching below £125,000 in every month (these adjustments reduce activity in the control group *outside* the holiday and increase it *during* the holiday). The findings in panel C are qualitatively similar to panels A and B and quantitatively in the middle of those two panels. The findings show that temporary housing stimulus successfully increases activity above its counterfactual level during the policy, with the effect being around 10% additional housing transactions in the later months of the stimulus. However, the stimulus also reduces activity below its counterfactual level after the policy. This suggests that the stimulus effect is entirely a short- to medium-term timing phenomenon. To clarify this finding, Panel D shows the *cumulative* log number of transactions based on the flows in Panel C. This graph shows that the entire stimulus effect is cancelled out only 10 months after the holiday.

These stimulus findings—especially the clear evidence of post-stimulus reversal—are consistent with the important findings by Mian & Sufi (2012) based on a temporary (2 months) auto purchase subsidy scheme in the US. One contribution of our stimulus analysis is to show similarly strong reversal for a program that was larger and longer-lasting (16 months), and which was aimed at a market viewed by many as pivotal to macroeconomic fluctuations. Of course, since the housing market is seen to be central to the macroeconomy and a key channel through which spillovers between markets occur, this timing effect may still have had a salutory overall effect at a point when the economy was in recession and so this need not be an indictment of the policy as a whole. Indeed, this suggests that the policy may have been withdrawn too soon.

6 Conclusions

This paper has studied the impact of a property transaction tax on the property market, using unique administrative data on every property transaction in the UK from 2005-2011 and compelling sources of quasi-experimental variation created by the notched structure of the tax schedule and sharp changes in the tax schedule over time. We have presented evidence on the effects of transaction taxes on house prices and house transaction volumes, including an analysis of the dynamics of adjustment to both anticipated and unanticipated tax changes. The overall finding is that prices and activity in the housing market respond sharply and quickly to transaction taxes in the way that economic theory predicts. Our study of transaction taxes in the property market could also have implications for the potential effects of transaction taxes in other asset markets, including transaction taxes on financial assets that have been discussed widely in recent years.[17]

Our findings from the 2008–2009 stamp duty holiday contribute to the scant micro evidence of the effectiveness of fiscal stimulus and, in particular, present some of the first evidence on the effectiveness of using temporary changes in tax policy to stimulate the housing market during economic downturns. The 16-month stamp duty holiday was successful in stimulating activity, with its effect peaking at around 10% extra monthly transactions (disregarding the 50% boost in activity in the final month of the holiday, due mostly to short-term timing effects). However, consistent with previous findings on the effect of a much shorter-term fiscal stimulus of the car market by Mian & Sufi (2012), the cumulative effect of the stamp duty holiday was neutralized only 10 months after the end of the stamp duty holiday. The quick reversal effect for such a large and long-lasting stimulus program has potentially important implications for macro stabilization policy.

[17] The UK also imposes stamp duty on stock transactions, but this stamp duty schedule is structured very differently and does not offer the same kind of quasi-experimental variation.

Table 1: Residential Property Tax Notches

<table>
<tr><th>Date Range / Price Range</th><th>17 Mar 2005
to
22 Mar 2006</th><th>23 Mar 2006
to
2 Sep 2008</th><th>3 Sep 2008
to
31 Dec 2009</th><th>1 Jan 2010
to
5 Apr 2011</th></tr>
<tr><td>0 - £120K</td><td>0</td><td rowspan="2">0</td><td rowspan="3">0</td><td rowspan="2">0</td></tr>
<tr><td>£120K - £125K</td><td rowspan="3">1</td></tr>
<tr><td>£125K - £175K</td><td rowspan="2">1</td><td rowspan="2">1</td></tr>
<tr><td>£175K - £250K</td><td>1</td></tr>
<tr><td>£250K - £500K</td><td>3</td><td>3</td><td>3</td><td>3</td></tr>
<tr><td>£500K - ∞</td><td>4</td><td>4</td><td>4</td><td>4</td></tr>
</table>

Notes: The table shows how the stamp duty land tax schedule for residential property has varied over time. Each column represents a time period during which the tax schedule was constant. The rows represent price ranges, and the entry in each cell is the tax rate that applies to that price range in the time period.

Figure 1: Theoretical Figures

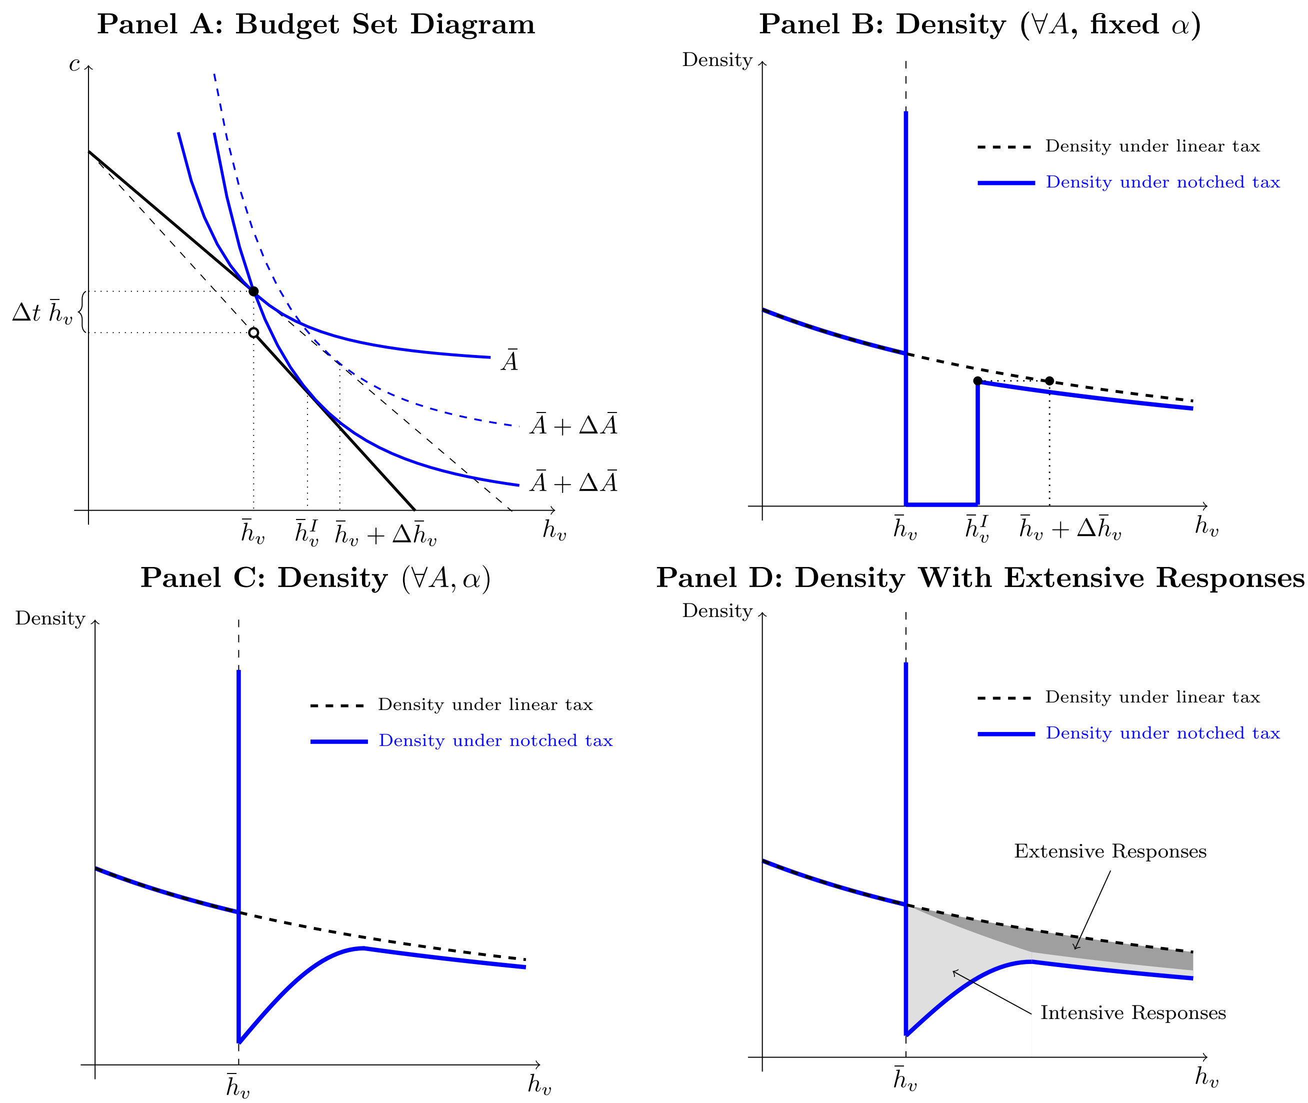

Notes: Figure 1 illustrates the implications of a notched transaction tax schedule in a budget set diagram (Panel A) and density distribution diagrams (Panels B-D). The budget set diagram in panel A (depicting preferences as in equation (1) and the budget set given by equation (2) in (h_v, c)-space) illustrates intensive responses among individuals with heterogeneous housing preferences A, but a specific demand elasticity α. The notch creates bunching at the cutoff $\overline{h}_v$ by all individuals in a preference range $(\overline{A}, \overline{A}+\Delta\overline{A})$, who would have bought houses on the segment $(\overline{h}_v, \overline{h}_v+\Delta\overline{h}_v)$ in the absence of the notch. The marginal bunching individual at $\overline{A}+\Delta\overline{A}$ is indifferent between the notch point $\overline{h}_v$ and the best interior location $\bar{h}_v^I$. No individual is willing to locate between $\overline{h}_v$ and $\bar{h}_v^I$, and hence this range is completely empty. The density of property values corresponding to the budget set diagram (all A, one specific α) is shown in Panel B. Since the behavioral response in Panels A-B depends on the size of the demand elasticity α (and converges to zero for completely price inelastic buyers), the density in the full population (all A, α) can be illustrated as in Panel C where some individuals are willing to buy just above the notch point. In addition to intensive responses, the notch creates extensive responses above the cutoff by individuals close to the indifference point between buying and not buying ($q \approx q^*$, where q^* is defined in equation (4)). However, such extensive responses will be negligible *just* above the cutoff. Intuitively, if an individual prefers buying a house slightly above $\overline{h}_v$ in the absence of the notch, then he will be better off by buying a house at $\overline{h}_v$ (which is almost as good) than not buying at all in the presence of the notch. This reasoning implies that extensive responses affect the density as illustrated in Panel D.

Figure 2: Budget Set Diagram for Bargaining Model

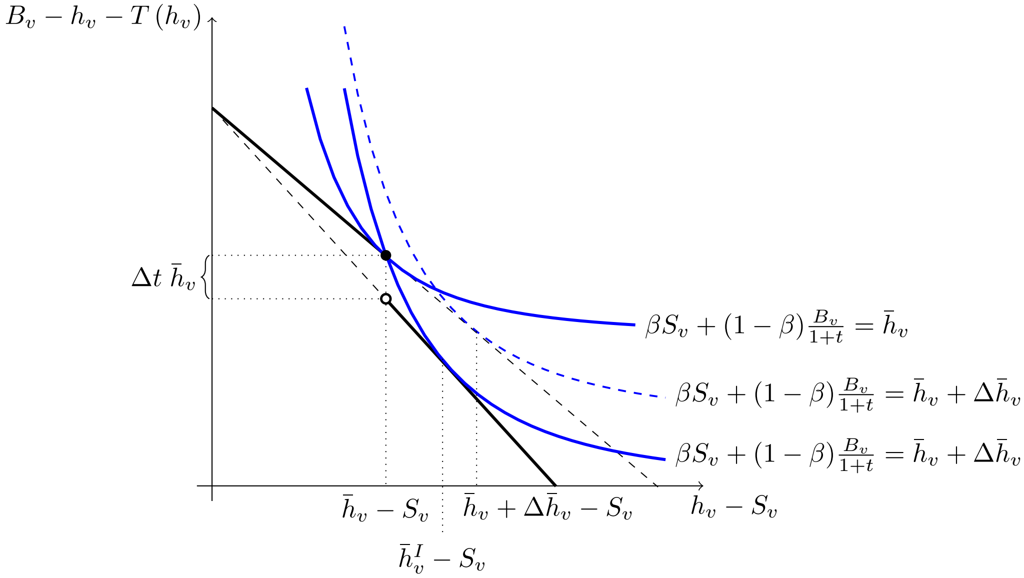

Notes: The budget set diagram depicts the Nash product as in equation (9) and the budget set of feasible allocations under the notched tax schedule in the space of net of tax surpluses (i.e. $(B_v - h_v - T(h_v), h_v - S)$-space) and illustrates intensive responses among individuals with heterogeneous valuations $\{B_v, S_v\}$, but a specific bargaining power β. The notch creates bunching at the cutoff $\bar{h}_v$ by all individuals in a preference range $\beta S_v + (1-\beta)\frac{B_v}{1+t} \in [\bar{h}_v, \bar{h}_v + \Delta\bar{h}_v]$, who would have bargained prices on the segment $[\bar{h}_v, \bar{h}_v + \Delta\bar{h}_v]$ in the absence of the notch. The marginal bunching match is indifferent between the notch point $\bar{h}_v$ and the best interior location $\bar{h}_v^I$. No individual is willing to locate between $\bar{h}_v$ and $\bar{h}_v^I$, and hence this range is completely empty. This figure is the direct analog of panel A of figure 1, and shares all its qualitative features.

Figure 3: Tax Schedule in 2010

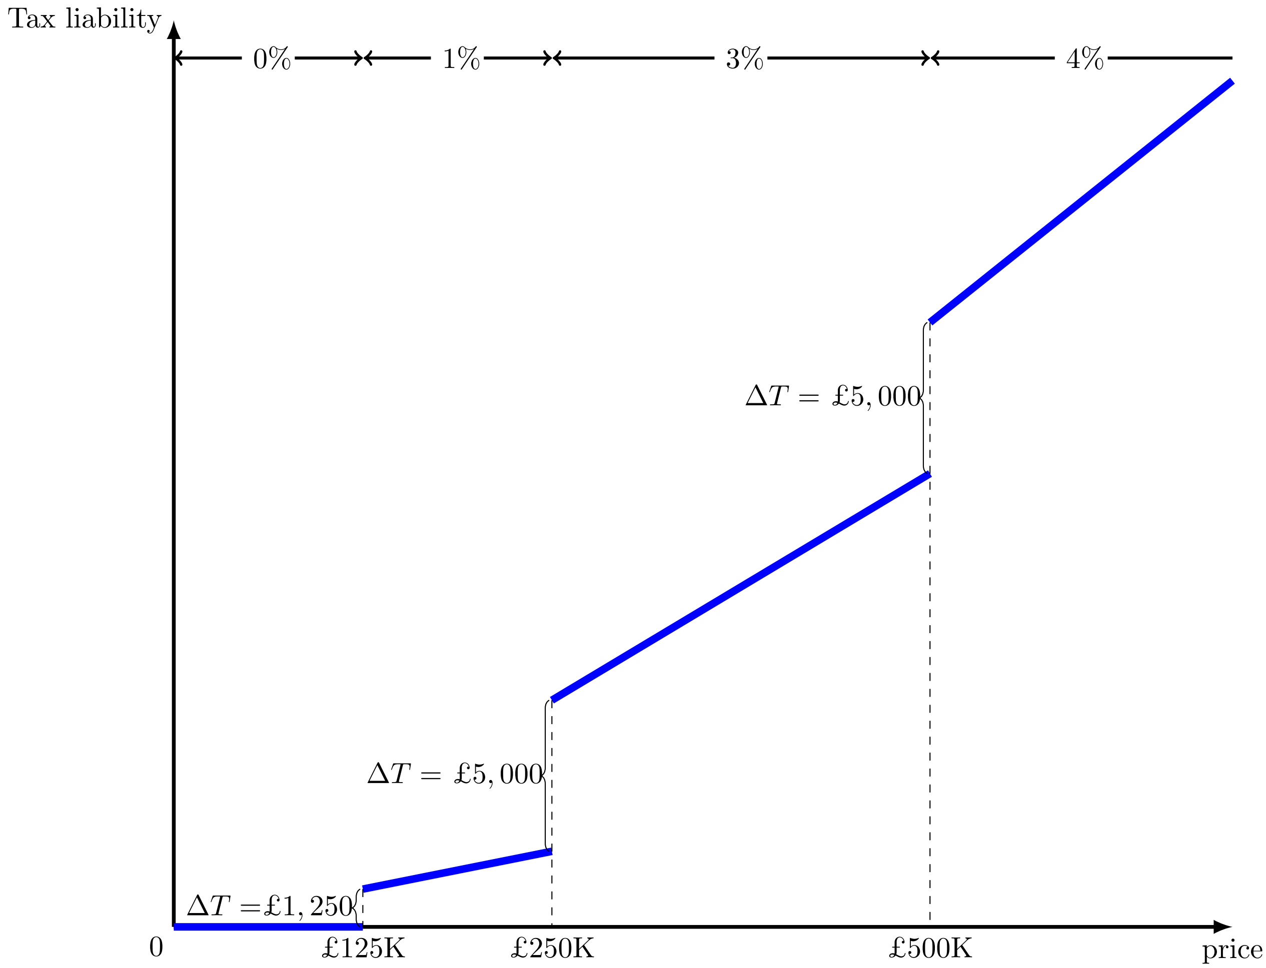

Notes: Figure 3 shows the stamp duty land tax schedule for residential properties in place in 2010 graphically as the solid blue line. The tax liability jumps discretely at the notches at £125,000, £250,000 and £500,000. Within the brackets defined by these notches, the tax rate is constant, and applied to the whole transaction price at the rates shown along the top of the figure.

Figure 4: Summary Statistics

Panel A: Number of Transactions

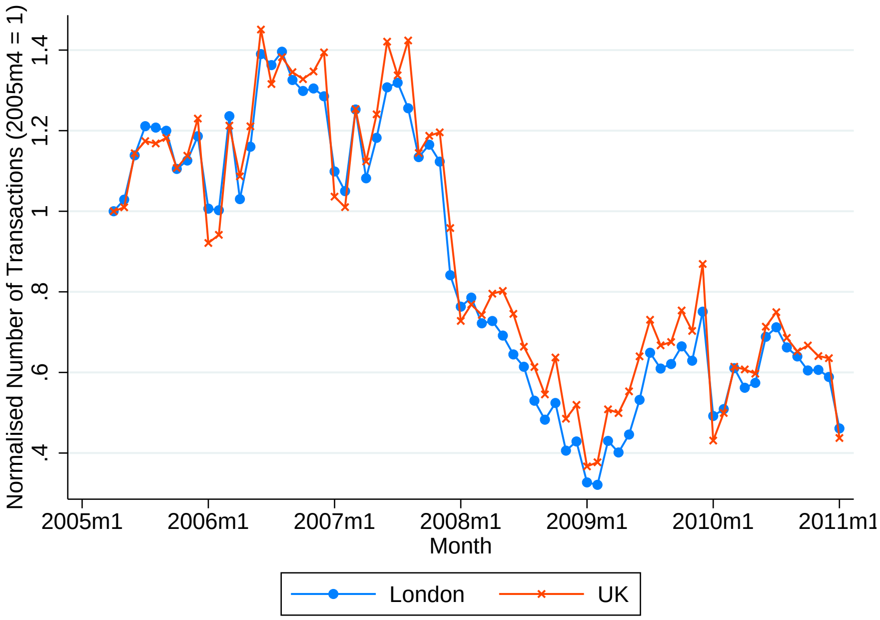

Panel B: Average Price

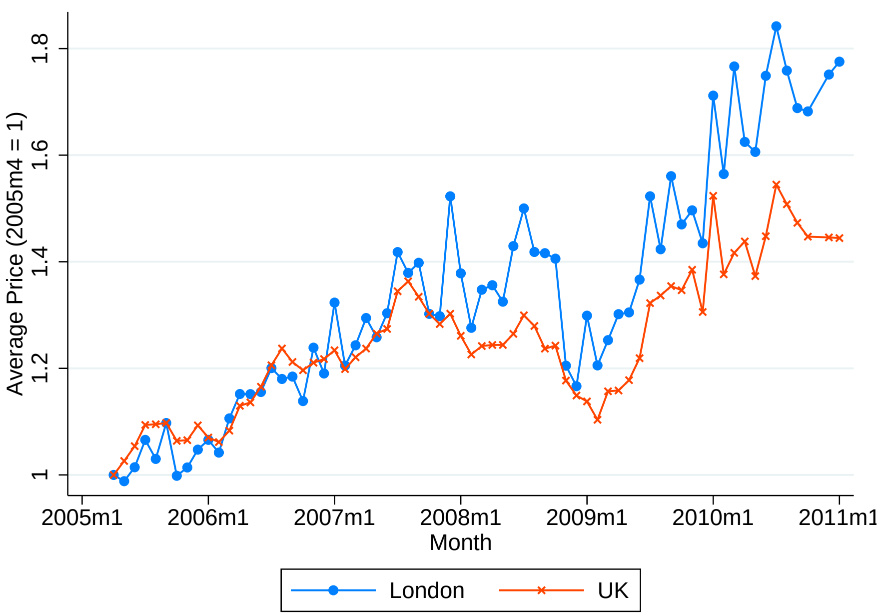

Notes: Panel A shows the monthly average price of property transactions relative to the average price in April 2005 (the beginning of our data period) in London (blue circles) and the U.K. (orange crosses). The average price of property transactions in London during the period April 2005 - January 2011 was £327,717 and the average price in the U.K. during our data period was £189,888. Panel B shows the monthly total number of property transactions relative to the number that took place in April 2005 (the beginning of our data period) in London (blue circles) and the U.K. (orange crosses). The average monthly number of property transactions in London during the period April 2005 - January 2011 was 12,352 while the average monthly number of property transactions in this period in the U.K. was 110,730.

Figure 5: Bunching and Holes Around the Upper Notches

Panel A: Notch at £250,000

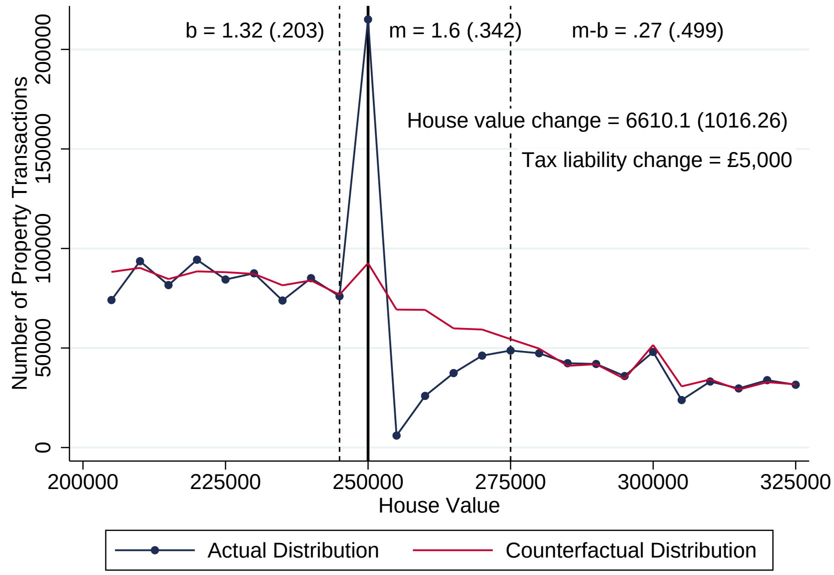

Panel B: Notch at £500,000

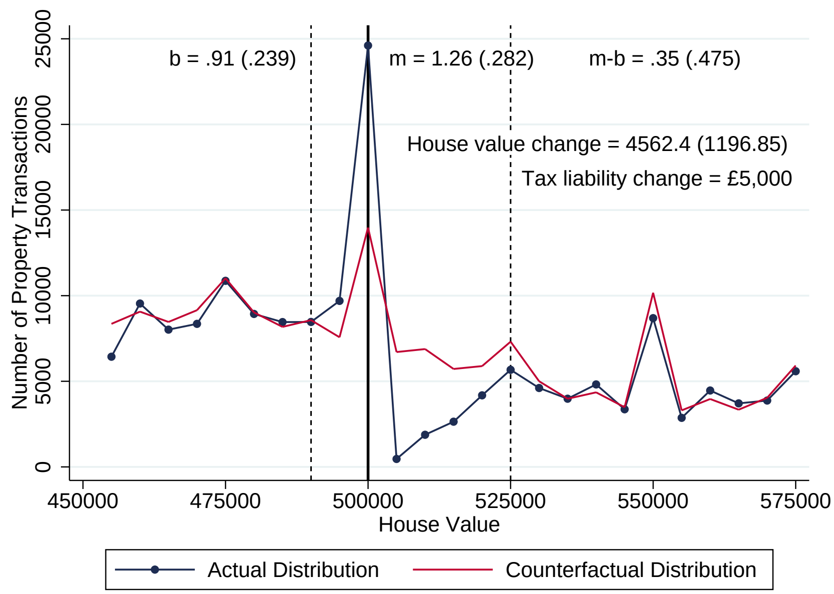

Notes: The figure shows the observed density of property transactions (blue dots) and our estimated counterfactual density (red line) around the notch at £250,000 where the tax liability jumps by £5,000 (from 1% to 3% of the transaction price) in panel A and around the notch at £500,000 where the tax liability jumps by £5,000 again (from 3% to 4% of the transaction price). The data used for these estimates excludes transactions that claim relief from the stamp duty land tax as the regular tax schedule does not apply to these transactions. The counterfactual density is estimated as in equation (12), using bins £100 pounds wide and a polynomial of order 5. The vertical dashed lines denote the upper and lower bounds of the excluded region around the notch. The upper bound of the excluded region is chosen as the point where the observed density changes slope from positive to negative. The estimate of equation (12) controls for round number bunching at multiples of £500, £1,000, £5,000, £10,000, £25,000 and £50,000. Both the empirical and the counterfactual density are shown aggregated up to bins £5,000 wide. b is our estimate of the excess mass just below the notch scaled by the counterfactual frequency at the notch, with its standard error shown in parentheses. m is our estimate of the missing mass above the notch scaled by the counterfactual frequency at the notch, with its standard error shown in parentheses. $m - b$ is our estimate of the difference between the missing mass and the bunching mass, again with its standard error in parentheses. The figures also show the average house value change created by the notch, and the tax liability change at the notch. All standard errors are obtained by bootstrapping the procedure 200 times.

Figure 6: Bunching and Holes Around the Lower Notch

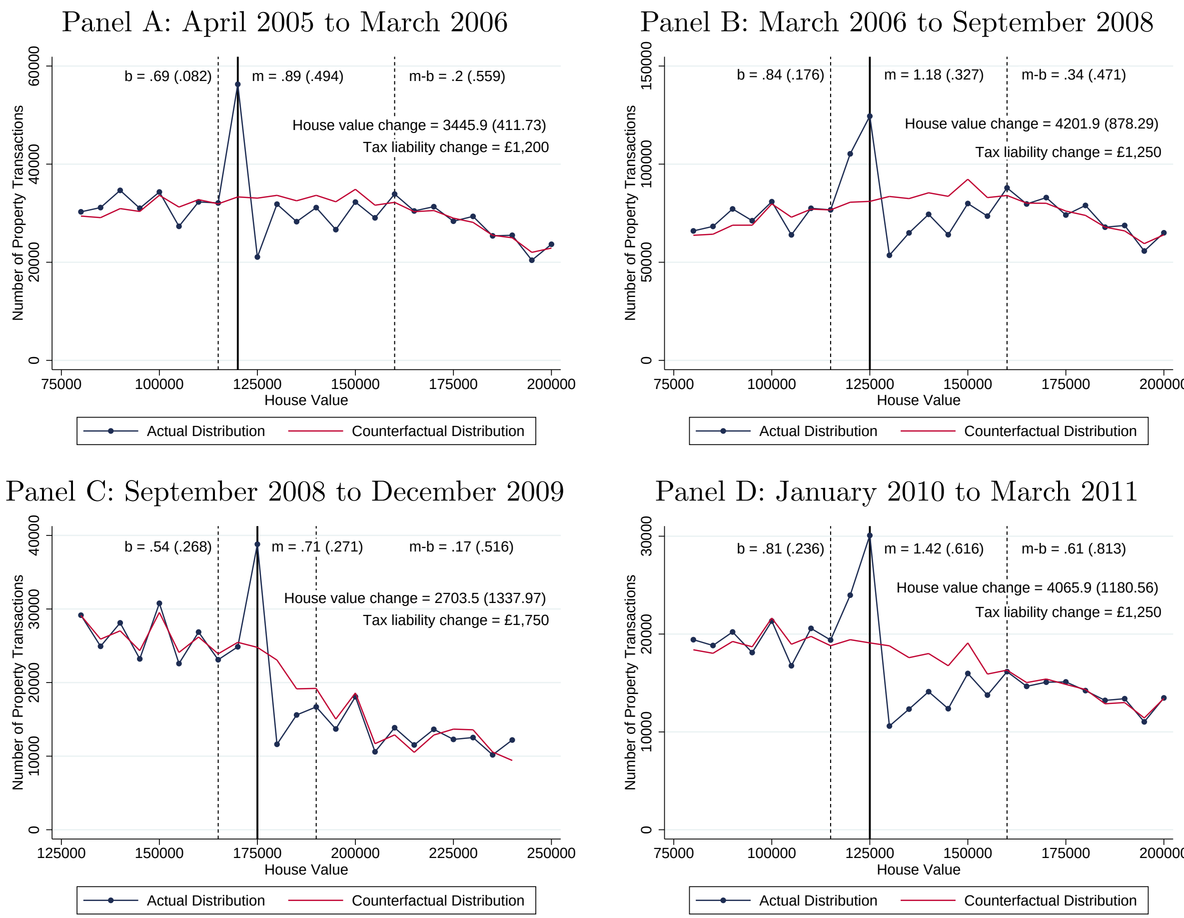

Notes: The figure shows the observed density of property transactions (blue dots) and our estimated counterfactual density (red line) around the lower notch in the residential property tax schedule where the tax liability jumps from 0 to 1% of the transaction price. Panel A shows the period 1 April 2005 to 22 March 2006 when the notch was at £120,000. Panel B shows the period 23 March 2006 to 2 September 2008 when the notch was at £125,000. Panel C shows the period 3 September 2008 to 31 December 2009 when the notch was at £175,000. Panel D shows the period 1 January 2009 to 31 March 2011 when the notch was at £125,000. The data used for these estimates excludes transactions that claim relief from the stamp duty land tax as the regular tax schedule does not apply to these transactions. The counterfactual density is estimated as in equation (12), using bins £100 pounds wide and a polynomial of order 5 in panels A and B and of order 7 in panels C and D. The vertical dashed lines denote the upper and lower bounds of the excluded region around the notch. The upper bound of the excluded region is chosen as the point where the observed density stops increasing and becomes decreasing (apart from spikes at round numbers). The estimate of equation (12) controls for round number bunching at multiples £500, £1,000, £5,000, £10,000, £25,000 and £50,000. Both the empirical and the counterfactual density are shown aggregated up to bins £5,000 wide. b is our estimate of the excess mass just below the notch scaled by the counterfactual density at the notch, with its standard error shown in parentheses. m is our estimate of the missing mass above the notch scaled by the counterfactual density at the notch, with its standard error shown in parentheses. $m - b$ is our estimate of the difference between the missing mass and the bunching mass, again with its standard error in parentheses. The figures also show the average house value change created by the notch, and the tax liability change at the notch. All standard errors are obtained by bootstrapping the procedure 200 times.

Figure 7: Dynamics of Bunching at Bottom Notch around March 2006

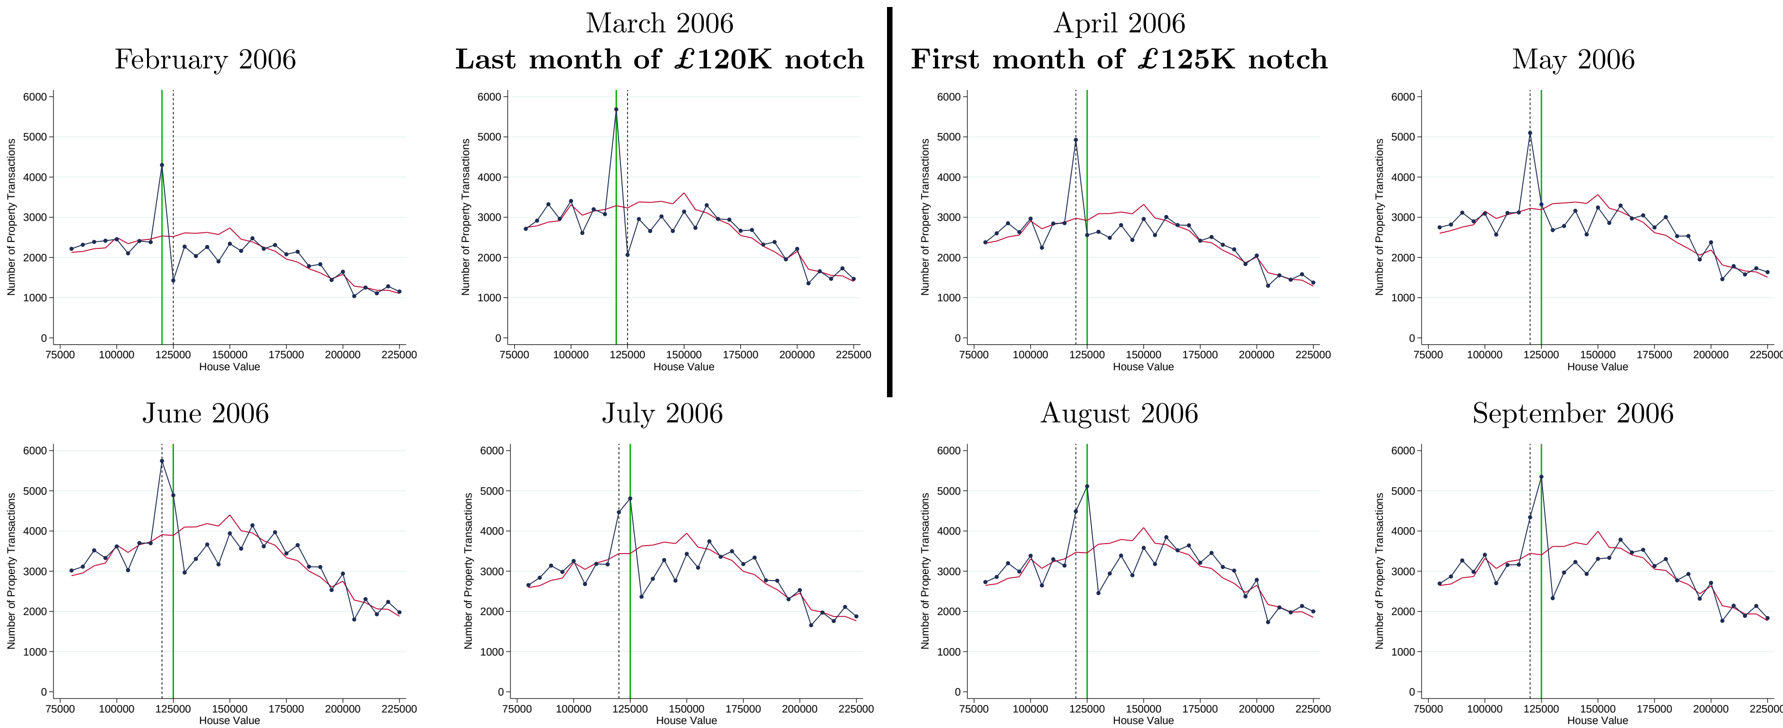

Notes: The figure shows the observed density of property transactions (blue dots) and our estimated counterfactual density (red line) in the region £75,000 – £225,000 separately for each month. On 23 March 2006, the bottom notch moved from £120,000 to £125,000. The estimation of the counterfactual is as described in section 4 and in the notes to figures (5 & 6). The estimation excludes data in the regions £115,000 - 160,000 and £170,000 - £190,000 and uses a polynomial of order 6.

Figure 8: Dynamics of Bunching Around the Beginning of Stamp Duty Holiday

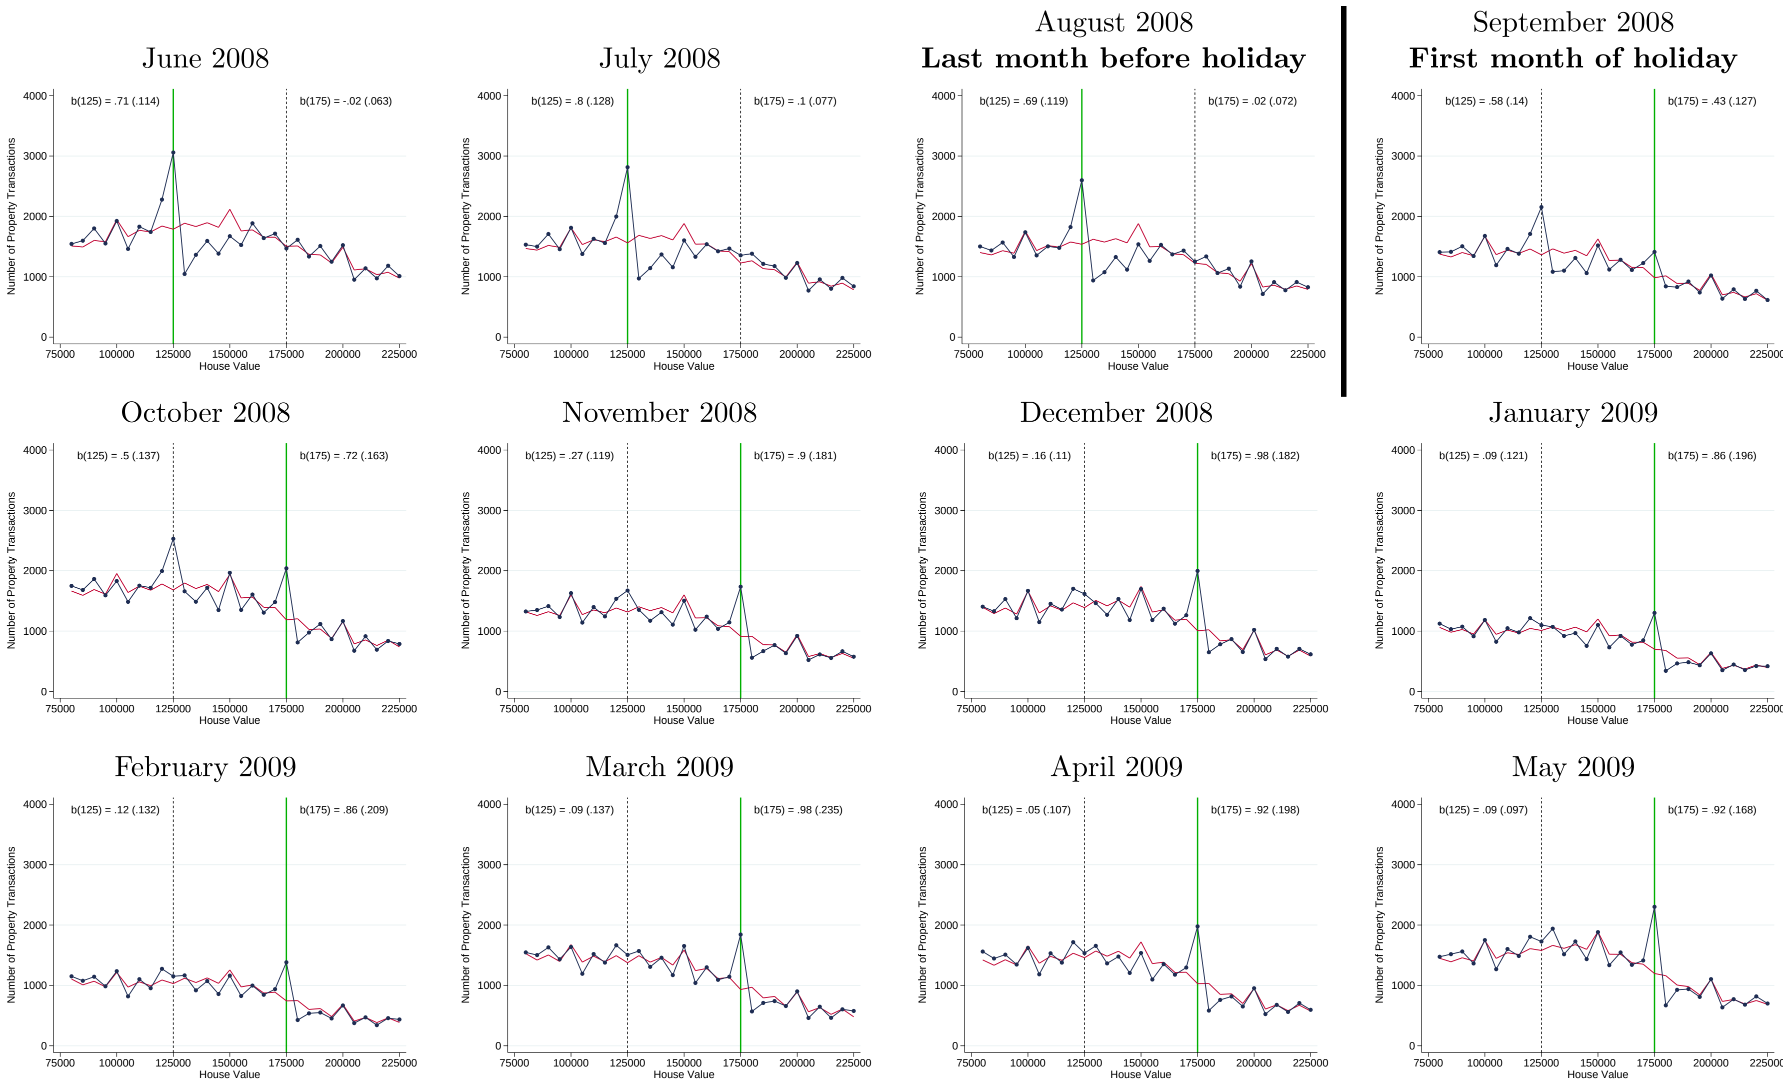

Notes: The figure shows the observed density of property transactions (blue dots) and our estimated counterfactual density (red line) in the region £75,000 – £225,000 separately for each month. On 3 September 2008, the bottom notch was moved unexpectedly from £125,000 to £175,000. The estimation of the counterfactual is as described in section 4 and in the notes to figures (5 & 6). The estimation excludes data in the regions £115,000 – 160,000 and £170,000 – £190,000 and uses a polynomial of order 6.

Figure 9: Dynamics of Bunching Around the End of Stamp Duty Holiday

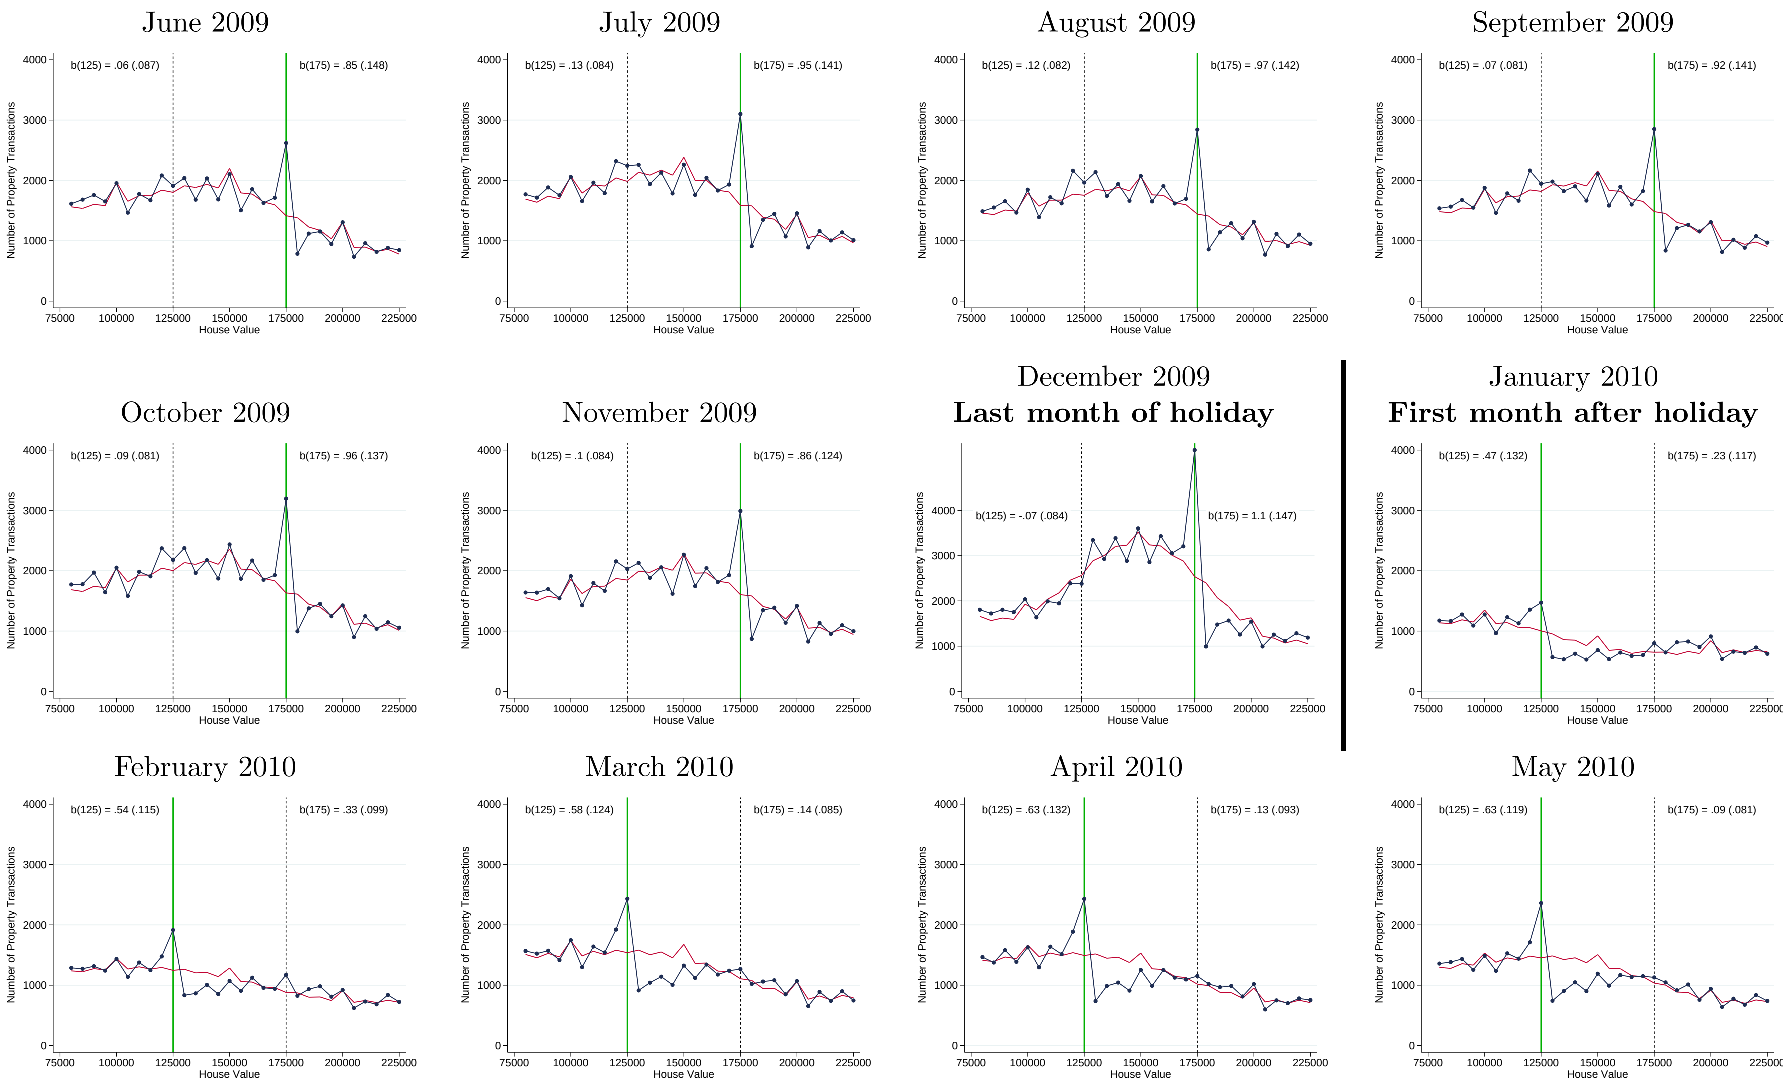

Notes: The figure shows the observed density of property transactions (blue dots) and our estimated counterfactual density (red line) in the region £75,000 – £225,000 separately for each month. On 1 January 2010, the bottom notch was moved back from £175,000 to £125,000 as announced previously. The estimation of the counterfactual is as described in section 4 and in the notes to figures (5 & 6). The estimation excludes data in the regions £115,000 – 160,000 and £170,000 – £190,000 and uses a polynomial of order 6.

Figure 10: Bunching Estimates Over Time

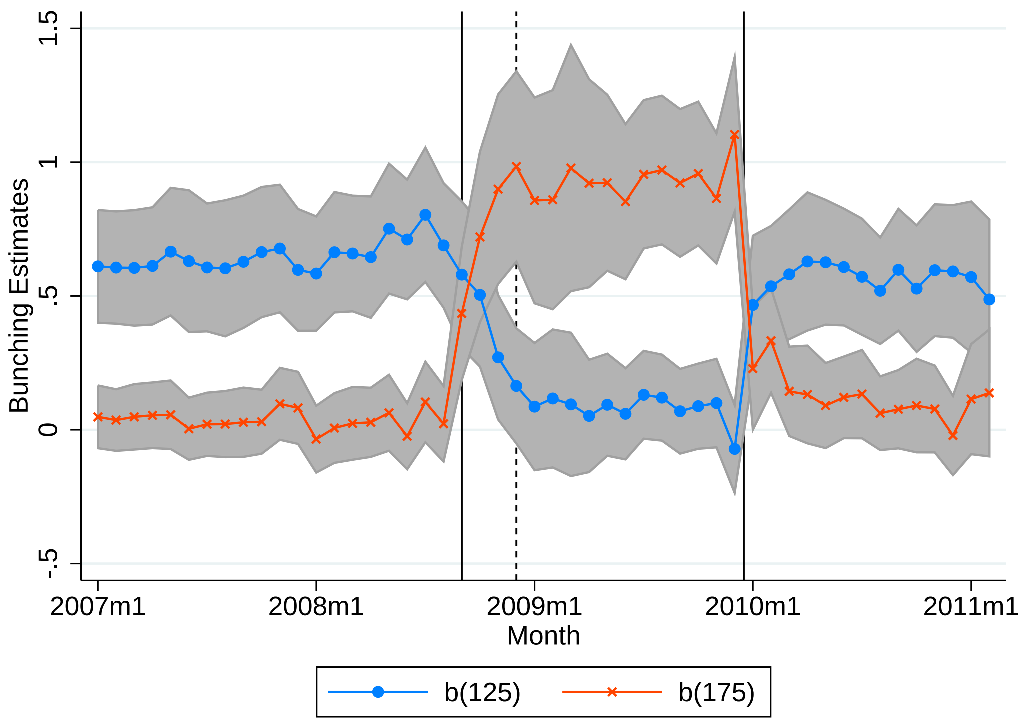

Notes: The figure shows our estimates of $b\left(\bar{h}_v\right)$, the bunching mass just below $\bar{h}_v$ scaled by the counterfactual frequency at $\bar{h}_v$, by month from January 2007 to February 2011 and for two values of $\bar{h}_v$, £125,000 (blue circles) and £175,000 (orange crosses). The first vertical line is at September 2008 when the stamp duty holiday was unexpectedly announced, moving the notch from £125,000 to £175,000. The dashed vertical line is at December 2008 to represent the observation that house transactions take up to 90 days to conclude, and so some inertia in the bunching responses is to be expected. The second vertical line is at December 2009 when the stamp duty holiday came to an end as anticipated, and the notch was moved from £175,000 back down to £125,000. The grey shaded regions are 95% confidence intervals obtained by bootstrapping each month's estimates 200 times.

Figure 11: Time Notch

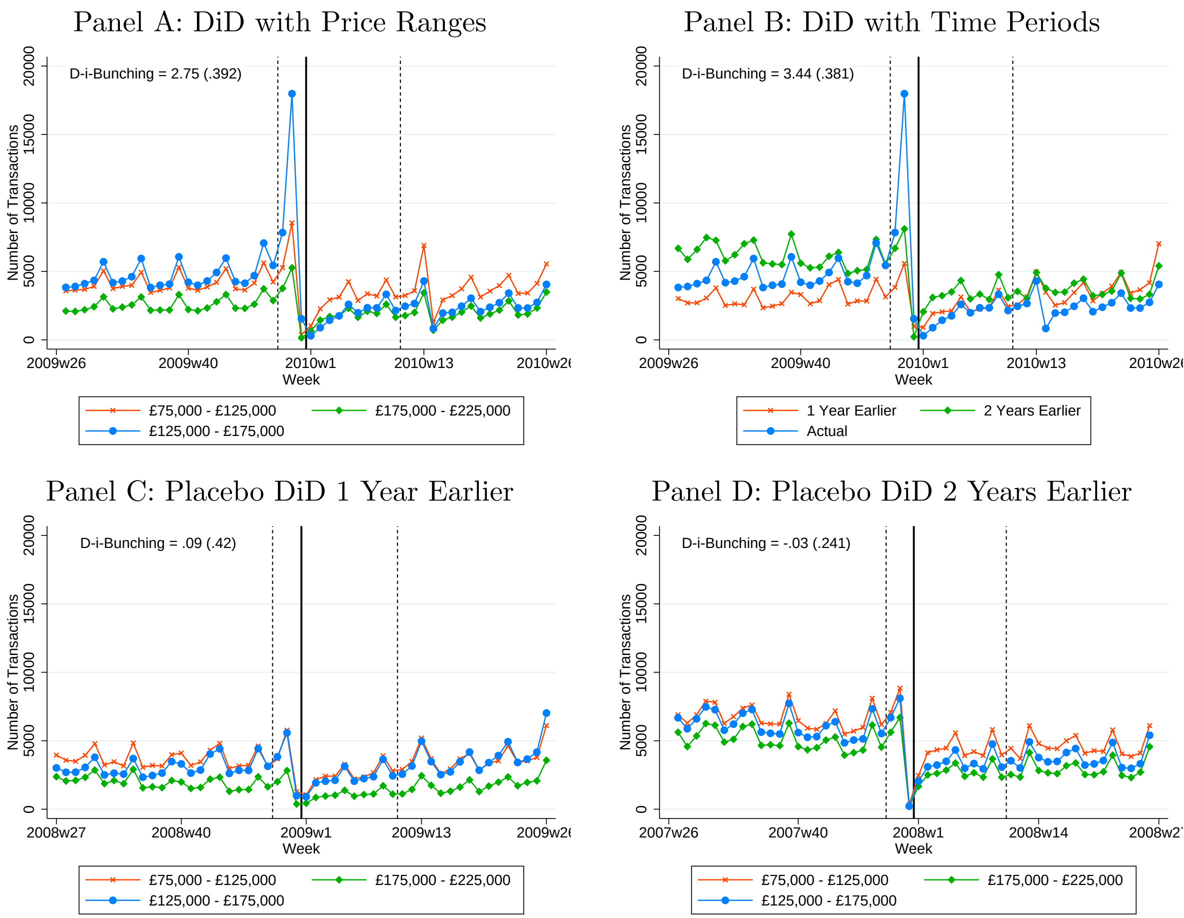

Notes: The figures show counts of the number of transactions in various weeks around the end of the stamp duty holiday on 31 December 2009. Panel A shows the number of transactions taking place between 2009w27 and 2010w26 at prices between £125,000 and £175,000 (blue circles) who as a result of the holiday paid no stamp duty alongside the number of transactions taking place at prices between £75,000 and £125,000 (orange crosses) and between £175,000 and £225,000 (green diamonds). Panel B shows the number of transactions taking place in the affected price range (£125,000 to £175,000) in the year around the end of the stamp duty holiday, 2009w27 to 2010w26 (blue circles) as well as 1 year earlier (orange crosses) and 2 years earlier (green diamonds). Panel C shows the first placebo difference in differences exercise, depicting the same price ranges as in panel A, but using data from 1 year earlier. Similarly, panel D shows the second placebo difference in differences exercise, depicting the same price ranges as in panel A, but using data from 2 years earlier. The solid vertical line is placed at the end of the year (which at the end of 2009 is the end of the stamp duty holiday) and the dashed vertical lines demarcate the last 3 weeks of the year and the first 10 weeks of the year, which are the excluded range for the counterfactual estimates. Overlaid on each picture is the difference-in-bunching estimate corresponding to the choice of treatment (blue circles) and control groups (orange crosses and green diamonds) depicted in the picture. The DiD estimate is the difference between the bunching in the treatment group and the average of the bunching in the two control groups. The counterfactual is estimated as a 7th order polynomial, and includes fixed effects for the last week of each month (except December), when there is a visible spike in each month. We do not include a fixed effect in December because this month is clearly special due to the holiday season, featuring a spike in the penultimate week of the year and a lull in the last week of the year. This seasonality is captured by the bunching estimates in the control groups and the placebo exercises.

Figure 12: Stimulus Effects of the Stamp Duty Holiday

Panel A: £125K to £175K Treatment Range

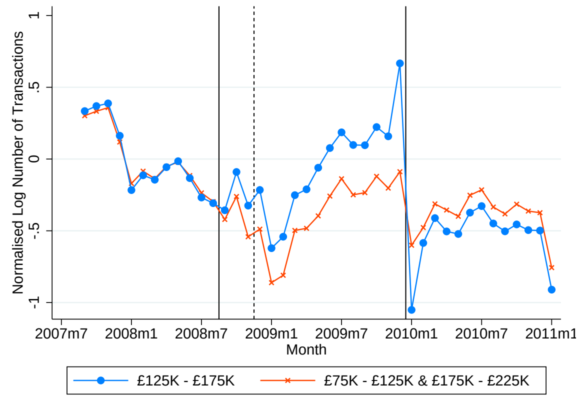

Panel B: £115K to £195K Treatment Range

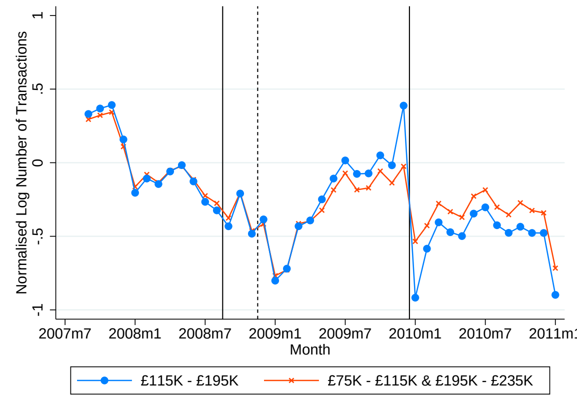

Panel C: £125K to £175K Treatment Range. Adjusting for Bunching

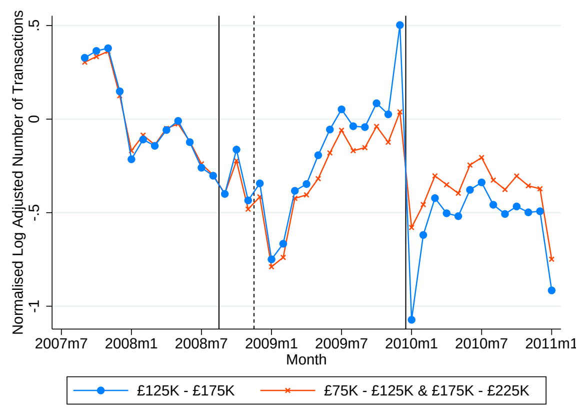

Panel D: £125K to £175K Treatment Range. Adjusting for Bunching. Cumulative Effect

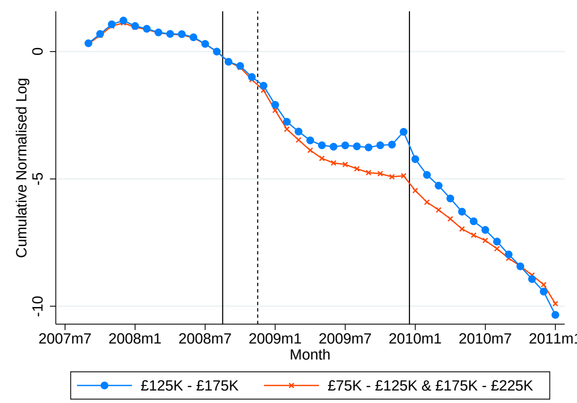

Notes: The figure shows the effect of the stamp duty holiday stimulus on housing market activity using different selections of treatment and control groups. Panel A defines the treated price range as transactions taking place at prices between £125,000 and £175,000 and takes as a control group transactions taking place at prices between £75,000 and £125,000 or between £175,000 and £225,000. Panel B defines the treated price range as transactions taking place at prices between £115,000 and £195,000 and takes as a control group transactions taking place at prices between £75,000 and £115,000 or between £195,000 and £235,000. Panels C & D define the treated price range as transactions taking place at prices between £125,000 and £175,000 and take as a control group transactions taking place at prices between £75,000 and £125,000 or between £175,000 and £225,000 as in Panel A. However, in each month, the number of transactions is adjusted to account for bunching. Specifically the number of transactions taking place between £125,000 and £175,000 is adjusted by removing excess transactions at £175,000 and adding the excess transactions at £125,000, and vice versa for the control group. Panels A-C show the normalised log monthly number of transactions. The Normalised log number of transactions in a month is defined as the log of the number of transactions in that month minus the average of the log of the number of transactions in the 12 months leading up to the start date of the Stamp Duty Holiday (September 2007 - August 2008). Panel D is the cumulative sum of the normalised log counts in panel C. The solid vertical lines mark the beginning and the end of the stamp duty holiday, while the dashed line marks 3 months after the start of the stamp duty holiday, by which time all transactions that started before the holiday should have been concluded.

References

ADELINO, MANUEL, SCHOAR, ANTOINETTE, & SEVERINO, FELIPE. 2012. *Credit Supply and House Prices*. NBER Working Paper no. 17832.

AGARWAL, SUMIT, LIU, CHUNLIN, & SOULELES, NICHOLAS S. 2007. The Reaction of Consumer Spending and Debt to Tax Rebates – Evidence From Consumer Credit Data. *Journal of Political Economy*, **115**, 986–1019.

AUERBACH, ALAN J. 1988. Capital Gains Taxation in the United States: Realizations, Revenue, and Rhetoric. *Brookings Papers on Economic Activity*, **19**, 595–638.

BESLEY, TIMOTHY, MEADS, NEIL, & SURICO, PAOLO. 2011. *The Incidence of a Transactions Tax: Evidence from a Tax Holiday*. Working Paper, London School of Economics.

BURMAN, LEONARD E., & RANDOLPH, WILLIAM C. 1994. Measuring Permanent Responses to Capital-Gains Tax Changes in Panel Data. *American Economic Review*, **84**, 794–809.

CHETTY, RAJ, FRIEDMAN, JOHN, OLSEN, TORE, & PISTAFERRI, LUIGI. 2011. Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records. *Quarterly Journal of Economics*, **126**, 749–804.

CUNNINGHAM, CHRISTOPHER R., & ENGELHARDT, GARY V. 2008. Housing Capital-Gains Taxation and Homeowner Mobility: Evidence From the Taxpayer Relief Act of 1997. *Journal of Urban Economics*, **63**, 803–815.

DACHIS, BEN, DURANTON, GILLES, & TURNER, MATTHEW A. 2012. The Effects of Land Transfer Taxes on Real Estate Markets: Evidence From a Natural Experiment in Toronto. *Journal of Economic Geography*, **12**, 327–354.

FELDSTEIN, MARTIN, SLEMROD, JOEL, & YITZHAKI, SHLOMO. 1980. The Effects of Taxation on the Selling of Corporate Stock and the Realization of Capital Gains. *Quarterly Journal of Economics*, **94**, 777–791.

JOHNSON, DAVID S., PARKER, JONATHAN A., & SOULELES, NICHOLAS S. 2006. Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review*, **96**, 1589–1610.

KLEVEN, HENRIK JACOBSEN, & WASEEM, MAZHAR. 2012. *Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan*. mimeo: London School of Economics.

KLEVEN, HENRIK JACOBSEN, KREINER, CLAUS THUSTRUP, & SAEZ, EMMANUEL. 2009. *Why Can Modern Governments Tax So Much? An Agency Model of Firms as Fiscal Intermediaries*. NBER Working paper 15218.

Kopczuk, Wojciech, & Munroe, David. 2012. *Mansion Tax: The Effect of Transfer Taxes on Residential Real Estate Market*. Mimeo: Columbia University.

Mian, Atif, & Sufi, Amir. 2012. The Effects of Fiscal Stimulus: Evidence From the 2009 Cash For Clunkers Program. *Quarterly Journal of Economics*, **127**, 1107–1142.

Ngai, L. Rachel, & Tenreyro, Silvana. 2012. *Hot and Cold Seasons in the Housing Market*. mimeo: London School of Economics.

Poterba, James M. 1984. Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach. *Quarterly Journal of Economics*, **99**, 729–752.

Poterba, James M. 1992. Taxation and Housing: Old Questions, New Answers. *American Economic Review, Papers and Proceedings*, **82**, 237–242.

Rosen, Harvey S. 1979a. Housing Decisions and the U.S. Income Tax: An Econometric Analysis. *Journal of Public Economics*, **11**, 1–23.

Rosen, Harvey S. 1979b. Owner Occupied Housing and the Federal Income Tax: Estimates and Simulations. *Journal of Urban Economics*, **6**, 247–266.

Rosen, Harvey S. 1985. Housing Subsidies: Effects on Housing Decisions, Efficiency, and Equity. *In: Handbook of Public Economics, vol.1.*

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

Saez, Emmanuel, Slemrod, Joel, & Giertz, Seth. 2012. The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature*, **50**, 3–50.

Shan, Hui. 2011. The Effect of Capital Gains Taxation on Home Sales: Evidence From the Taxpayer Relief Act of 1997. *Journal of Public Economics*, **95**, 177–188.

Slemrod, Joel, Weber, Caroline, & Shan, Hui. 2012. *The Lock-In Effect of Housing Transfer Taxes: Evidence From a Notched Change in D.C. Policy.* mimeo: University of Michigan.