

Housing Market Responses to Transaction Taxes: Evidence From Notches and Stimulus in the U.K.

MICHAEL CARLOS BEST

Columbia University

and

HENRIK JACOBSEN KLEVEN

Princeton University

First version received February 2015; Editorial decision October 2016; Accepted March 2017 (Eds.)

We investigate housing market responses to transaction taxes using administrative data on all property transactions in the U.K. from 2004 to 2012 combined with quasi-experimental variation from tax notches and tax stimulus. We present two main findings. First, transaction taxes are highly distortionary across a range of margins, causing large distortions to the price, volume, and timing of property transactions. Secondly, temporary transaction tax cuts are an enormously effective form of fiscal stimulus. A temporary elimination of a 1% transaction tax increased housing market activity by 20% in the short run (due to both timing and extensive responses) and less than half of the stimulus effect was reversed after the tax was reintroduced (due to re-timing). Because of the complementarities between moving house and consumer spending, these stimulus effects translate into extra spending per dollar of tax cut equal to about 1. We interpret our empirical findings in the context of a housing model with downpayment constraints in which leverage amplifies the effects of transaction taxes.

Key words: Transaction taxes, Property taxes, Fiscal stimulus.

JEL Codes: H20, H21, H31, E62

1. INTRODUCTION

Transaction taxes on assets are widely discussed but remain understudied (Campbell and Froot, 1994; Poterba, 2002; Matheson, 2011; European Commission, 2013). We analyse transaction taxes on housing, a policy that is very common and raises substantial amounts of revenue in many countries.¹ Our empirical analysis focuses on the U.K. property transaction tax, known

1. As of 2010, twenty-seven countries in the OECD imposed property transaction taxes: Austria, Australia, Belgium, Canada, Chile, the Czech Republic, Denmark, Estonia, Finland, Germany, Greece, Hungary, Ireland, Iceland,

as the the *Stamp Duty Land Tax* (SDLT). The research design relies on administrative stamp duty records on the universe of property transactions in the U.K. from 2004 to 2012 combined with rich quasi-experimental variation from notches and stimulus policy. We find that transaction taxes cause large distortions to house prices, transaction volume, and timing, and that temporarily cutting such taxes can provide very effective stimulus during recession.

Two sources of tax variation are central to our analysis. First, the U.K. stamp duty features large discontinuities in tax liability—notches—at threshold property prices. For example, the tax rate jumps from 1% to 3% of the entire transaction price at a threshold of £250,000 (about \$400,000), creating an increase in tax liability of £5,000 (about \$8,000) as the house price crosses this threshold. Such notches create strong incentives for bunching below the threshold, allowing for non-parametric identification of house price responses using bunching techniques (see Kleven (2016) for a review). Secondly, the U.K. stamp duty features substantial time variation, including permanent and temporary tax changes that affect specific price brackets but not others. We focus on a *stamp duty holiday* that lasted 16 months and eliminated transaction taxes in a particular price range in order to provide stimulus to the housing market during the Great Recession. This provides an ideal setting for a difference-in-differences approach to estimating extensive responses (whether or not to buy a house) and timing responses (when to buy a house) to temporary stimulus.

To facilitate interpretation of the empirical results, we develop a theoretical model of housing market responses to transaction taxes. Given our focus on short-run stimulus policy, we consider a setting with a fixed aggregate housing stock in which agents are trading existing houses. A central aspect of housing transaction taxes is that they require large cash payments up front that cannot be financed by borrowing, which may be particularly harmful in a market where agents are liquidity constrained and highly leveraged. Motivated by this observation, our model builds on the seminal housing model with downpayment constraints by Stein (1995). We show that transaction taxes can have large effects on both prices and transaction volumes, especially if there are many downpayment-constrained and highly leveraged households. The intuition turns on multiplier effects driven by the fact that transaction taxes reinforce downpayment constraints while the resulting fall in prices reinforces those constraints even further. We argue that this model provides a natural way of rationalizing the large effects we find.

Our empirical findings can be divided into four main categories. First, the distribution of house prices features large and sharp bunching below notch points combined with large holes above notch points. Our bunching estimates imply that house prices respond by a factor of 2–5 times the size of the tax increase at the notch, with larger effects at the bottom than at the top of the price distribution. Bunching in house prices, or rather in the values of transacted houses, represent nominal demand responses rather than price responses *per se*. As we show, in a world with downpayment constraints, such demand responses depend not only on the structural demand elasticity, but also on the amount of leverage among constrained buyers. For example, a downpayment-constrained buyer with a loan-to-value ratio of 75% (the modal value in the U.K.) will respond by a factor of 4 times the tax increase.²

Israel, Italy, Japan, the Netherlands, Norway, Poland, Portugal, South Korea, Spain, Sweden, Turkey, the U.K., and the U.S. (Andrews *et al.*, 2011). Within the U.S., thirty-eight states had a property transaction tax in 2012 (Lincoln Institute of Land Policy, 2014). Outside the OECD, countries such as Hong Kong, India, Pakistan, and Singapore impose property transaction taxes.

2. In addition, in a world with matching frictions and price bargaining, house price responses to transaction taxes may also be affected by search and bargaining effects (see Best and Kleven 2014; Besley *et al.* 2014; Kopczuk and Munroe 2015). There may also be an element of tax evasion and avoidance as we discuss below. Hence, estimating house price responses to taxes is conceptually related to the large literature estimating taxable income responses (see Saez *et al.* 2012), which combines real labour supply, wage bargaining, evasion/avoidance, etc.

Secondly, we consider the dynamics of house price responses using anticipated and unanticipated changes in the location of notches due to the stamp duty holiday. The dynamic adjustment of bunching and holes to movement of notches is very fast, with a new steady state emerging in about 3–4 months for unanticipated changes and almost immediately for anticipated changes. The remarkable sharpness of our dynamic findings suggests that agents in the housing market are less affected by optimization frictions (inattention, inertia, etc.) than, for example, agents in the labour market (Chetty *et al.*, 2011; Chetty, 2012; Kleven and Waseem, 2013; Gelber *et al.*, 2015; Kleven, 2016).

Thirdly, we find strong evidence of short-term timing responses to anticipated tax changes that create time notches at cutoff dates. In the 2 weeks leading up to the end of the stamp duty holiday, activity levels in the housing market increased by around 150%. Our sharp, non-parametric evidence on timing responses in the housing market contributes to previous findings that short-term timing responses may far exceed medium- or long-term responses (Auerbach, 1988; Burman and Randolph, 1994; Goolsbee, 2000).

Fourthly, we estimate medium-term timing and extensive margin responses to temporary tax changes using the stamp duty holiday. This stimulus programme suspended a 1% transaction tax in a specific price range at the lower end of the distribution. We find that the programme successfully raised housing market activity as transaction volumes in the treated price range clearly diverged from transaction volumes in untreated price ranges during the holiday. Market activity in the treated price range increased by as much as 20% due to the stimulus. This effect combines a *timing effect* (intertemporal substitution by those who would have purchased a house anyway) and an *extensive margin effect* (house purchases that would not have taken place absent the tax holiday). We can separate the two effects by comparing treatments and controls following the removal of the programme. Consistent with a timing effect, activity levels in the treatment group dropped by about 8% compared to the control group in the first year after the holiday, with no further reversal in the second year after the holiday. However, the total reversal effect due to re-timing was less than half of the total stimulus boost, in contrast to Mian and Sufi (2012) who find complete reversal within 1 year of a U.S. stimulus programme in the car market.

As in our model, these stimulus effects are most naturally interpreted as increased transactions of existing houses as the aggregate housing stock cannot respond to taxes in the short run. While more trading of existing houses do not add mechanically to GDP, they can have important real effects because trading house typically means moving house.³ Besides the implications of homeowner mobility for housing and labour markets, we highlight that moving house is associated with substantial household spending on repairs, renovations, removals, durable goods, and commissions to agents and lawyers. Using U.K. consumption survey data, we estimate that a house transaction triggers extra spending of about 5% of the house price. Combined with our estimated increase in transaction volume (20%) and the size of the tax cut (1% of the house price), this implies extra household spending per dollar of tax cut of about 1. This captures only the direct stimulus impact; it does not include potential general equilibrium effects (multiplier or price effects). Compared to a large body of evidence on consumer responses to other forms of fiscal stimulus such as tax rebates (*e.g.* Shapiro and Slemrod 2003a,b; Johnson *et al.* 2006; Agarwal *et al.* 2007; Parker *et al.* 2013), our findings suggest that the spending impact of the U.K. housing stimulus programme was considerably larger.

Our stimulus results carry qualitative lessons that are worth highlighting. First, the traditional Keynesian view is that tax multipliers are smaller than spending multipliers, because the direct tax multiplier is given by the marginal propensity to consume out of temporary income and is thus

3. Indeed, our empirical analysis focuses on *residential* property transactions.

significantly smaller than 1. In contrast, our findings suggest that transaction tax stimulus could be as effective as spending stimulus, not because the marginal propensity to consume is necessarily high, but because of large tax elasticities combined with consumption complementarities. Secondly, our article provides a counter-example to the important paper by Mian and Sufi (2012), who find complete and swift reversal following the U.S. Cash-for-Clunkers programme. In the article we discuss the potential reasons why our findings are so different than those from Cash-for-Clunkers. Thirdly, while our findings are based on the elimination of a distortionary tax, the insights may apply more broadly to stimulus policies that reduce transaction costs in the housing market. This includes the homebuyer tax credit introduced by the U.S. Stimulus Bill of 2009.

We contribute primarily to three literatures. First, we contribute to the voluminous literature attempting to estimate the effects of fiscal stimulus, and especially to recent work that uses micro data to address this question (*e.g.* Johnson *et al.* 2006; Agarwal *et al.* 2007; Mian and Sufi 2012; Parker *et al.* 2013). Our analysis of housing investment stimulus has a counterpart in recent work on business investment stimulus from U.S. bonus depreciation policies (*e.g.* House and Shapiro 2008; Zwick and Mahon 2014). In particular, the finding by Zwick and Mahon (2014) that financially constrained firms (such as small firms with low cash holdings) respond more to stimulus than unconstrained firms, which is conceptually related to our argument that financially constrained households (those who are downpayment constrained and highly leveraged) respond more strongly to housing stimulus.

Secondly, our article belongs to a small body of recent and contemporaneous work studying the effects of property transaction taxes on house prices and homeowner mobility (van Ommeren and van Leuvensteijn, 2005; Dachis *et al.*, 2012; Besley *et al.*, 2014; Slemrod *et al.*, 2014; Kopczuk and Munroe, 2015).⁴ In particular, the contemporaneous papers by Slemrod *et al.* (2014) and Kopczuk and Munroe (2015) study house price responses using bunching at U.S. tax notches. Unlike these papers, we analyse the dynamics of house price responses, extensive margin responses, timing and stimulus, and so the only real overlap is in the static bunching analysis. Also contemporaneous with this article, Besley *et al.* (2014) study the incidence of transaction taxes using the U.K. stamp duty holiday and mortgage data on house appraisals. None of the earlier literature considers the role of downpayment constraints and leverage for the impact of housing transaction taxes.

Finally, a larger empirical literature has examined the impact of capital gains taxes on asset prices and asset transactions (*e.g.* Feldstein *et al.* 1980; Auerbach 1988; Burman and Randolph 1994) and some of this work has focused specifically on the taxation of housing capital gains (Cunningham and Engelhardt, 2008; Shan, 2011). Transaction taxes and capital gains 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 is particularly important in the presence of downpayment constraints and high leverage, making behavioural responses to transaction taxes an order of magnitude larger than responses to capital gains taxes, other things equal.

The article proceeds as follows. Section 2 develops our theoretical model, section 3 describes the context and data, section 4 estimates house price responses using notches, section 5 estimates timing and extensive margin responses to stimulus, and section 6 concludes.

4. Besides these empirical papers, theoretical work by Lundborg and Skedinger (1999) has analysed the implications of housing transaction taxes in a search model.

2. THEORETICAL FRAMEWORK WITH TRANSACTION TAXES

This section develops a dynamic model of housing market responses to transaction taxes, building on the seminal housing model with downpayment constraints by Stein (1995). Accounting for the downpayment requirements of housing transactions is important in this context, because transaction taxes represent large upfront payments that cannot be financed by debt and therefore reinforce downpayment requirements. We model the effects of transaction taxes on existing homeowners deciding whether to trade their old home for a new one, holding the aggregate housing stock fixed. To guide the empirical analysis of fiscal stimulus, we focus on the impacts of a temporary transaction tax cut on housing demand and house prices, demonstrating how the effects are amplified by downpayment constraints.⁵ To simplify the exposition, we restrict attention to linear transaction taxes in the main text, while the Online Appendix considers the effects of notched transaction taxes.

2.1. *The model*

We consider a 2-period model. Households enter period 1 endowed with 1 unit of housing, and an outstanding debt of k_1 denominated in units of a numeraire consumption good. We allow k_1 to be heterogeneous in the population. Households can trade houses at a price of p_1 per unit of quality-adjusted housing h_1 . If they move house, they sell their existing house, repay their debt k_1 , and incur a fixed cost of moving q_1 . This gives them liquid assets of $p_1 - k_1 - q_1$. To buy a house of value $v_1 \equiv p_1 \cdot h_1$, households must make a downpayment of $\gamma_1 \cdot v_1$ financed from their liquid assets, while they borrow the remainder. They must also pay a transaction tax $\tau_1 v_1$ financed from their liquid assets. These financing requirements imply that a household can buy any house satisfying the liquidity constraint

$$v_1 \leq \frac{p_1 - k_1 - q_1}{\gamma_1 + \tau_1} \equiv v_1^c. \quad (1)$$

Households choose whether or not to move (extensive margin) and how much housing to buy conditional on moving (intensive margin). After making housing choices, they receive income y_1 , giving them consumption of

$$c_1 = y_1 + I_{m_1} \cdot [p_1 - k_1 - q_1 - (\gamma_1 + \tau_1)v_1], \quad (2)$$

where I_{m_1} is an indicator for moving in period 1. If a household moves, its outstanding mortgage debt equals $v_1(1 - \gamma_1)$, while otherwise it equals k_1 . The household's net debt carried forward to period 2 is therefore given by $k_2 = k_1 + I_{m_1} \cdot [v_1(1 - \gamma_1) - k_1]$. Denoting the gross interest rate by R , households enter period 2 with Rk_2 units of debt.

Households face a similar housing choice in period 2, but enter the period with h_1 units of housing (where $h_1 = 1$ if they did not move in period 1). If they move in period 2, they will have liquid assets of $p_2 h_1 - Rk_2 - q_2$ with which to finance a downpayment of $\gamma_2 v_2$ and a transaction tax of $\tau_2 v_2$. Hence the period-2 liquidity constraint requires households to buy a house satisfying

$$v_2 \leq \frac{p_2 h_1 - Rk_2 - q_2}{\gamma_2 + \tau_2} \equiv v_2^c. \quad (3)$$

5. For the sake of simplicity, we focus on a competitive housing market and therefore abstract from the search frictions that some have argued are significant in the housing market (see *e.g.* Wheaton 1990; Krainer 2001; Piazzesi and Schneider 2009; Ngai and Tenreyro 2014). In Best and Kleven (2014), we extend the model to allow for matching frictions and price bargaining, but in a version without downpayment constraints.

After making housing decisions, households receive income y_2 and repay all debts, giving them consumption of

$$c_2 = y_2 - Rk_2 + I_{m_2} \cdot (p_2 [h_1 - (1 + \tau_2)h_2] - q_2). \quad (4)$$

We normalize $y_2 = Ry_1$ and $y_1 = (1 + k_1)/2$ for convenience.⁶

Households maximize lifetime utility $u_1(c_1, h_1) + \delta u_2(c_2, h_2)$, where the per-period utility functions are given by

$$u_t(c_t, h_t) = c_t + \frac{A_t}{1 + 1/e_t} \left(\frac{h_t}{A_t} \right)^{1 + 1/e_t} \quad t \in \{1, 2\}. \quad (5)$$

Here $A_t > 0$ and $e_t < 0$ are parameters characterizing housing preferences, which we allow to be heterogeneous in the population and to vary over time. The quasi-linear utility function eliminates income effects on housing demand as well as consumption-smoothing motives for savings, both of which we abstract from for simplicity.⁷

Households move whenever moving to the best house they can afford, $h_t^* \leq v_t^c/p_t$, gives them higher utility than staying in their current house. In period 2, households move whenever

$$u_2(c_2^*, h_2^*) \geq u_2(y_2 - Rk_2, h_1), \quad (6)$$

where c_2^* is the consumption resulting from housing choice h_2^* . In period 1, households anticipate the effect that their housing choices will have on their utility in period 2. We can summarize a households' indirect utility in period 2 as a function of their period-1 housing choice and their resulting debt k_2 as $V(h_1, Rk_2) \equiv \max \{u_2(c_2^*, h_2^*), u_2(y_2 - Rk_2, h_1)\}$. Households therefore move in period 1 whenever

$$u_1(c_1^*, h_1^*) + \delta V(h_1^*, Rk_2^*) \geq u_1(y_1, 1) + \delta V(1, Rk_1). \quad (7)$$

This dynamic model features households with four types of moving patterns: Those who move in period 1, but not in period 2 ($h_1 = h_1^*$ & $h_2 = h_1$); those who move in both period 1 and period 2 ($h_1 = h_1^*$ & $h_2 = h_2^*$); those who do not move in period 1, but move in period 2 ($h_1 = 1$ & $h_2 = h_2^*$); and those who never move ($h_1 = 1$ & $h_2 = 1$). We denote the share of households who move in period t and are not downpayment constrained by U_t , the share of households who move in period t but are downpayment constrained by C_t , and the share of all movers in period t by $M_t = U_t + C_t$. Movers sell their existing house and buy a new house as characterized above. Equilibrium in the market requires that excess demand in each period D_t equals zero, that is,

$$D_1 = U_1 E[h_1^u - 1 | U_1] + C_1 E[h_1^c - 1 | C_1] = 0, \quad (8)$$

$$D_2 = U_2 E[h_2^u - h_1 | U_2] + C_2 E[h_2^c - h_1 | C_2] = 0, \quad (9)$$

where $h_t^c = v_t^c/p_t$ and h_t^u denote housing demand in period t among constrained and unconstrained movers, respectively. In equations (8) and (9), expectations are taken over household types $(A_1, A_2, e_1, e_2, \gamma_1, \gamma_2, k_1, q_1, q_2)$.

6. These normalizations imply that the income streams of all households with initial debt k_1 are the same. Relaxing these would require us to keep track of the net present value (NPV) and time profile of households' income when studying their housing decisions, adding complexity without changing any of the qualitative conclusions. Notice also that we assume that income in both periods y_1, y_2 is received after making housing and consumption decisions, which is the reason why current income does not relax current liquidity constraints.

7. While the quasi-linear utility formulation eliminates the standard consumption-smoothing motive to save, households could still have an incentive to save to relax future liquidity constraints. We rule out such liquidity-driven savings motives by assumption. We do not use these assumptions on savings and income effects in our empirical analysis.

2.2. Effects of temporary stimulus

Now let us consider a reduction of the transaction tax rate in period 1, keeping the transaction tax rate in period 2 unchanged. This should be interpreted as an unanticipated tax cut (as it happens in the first period of the model) and is known by taxpayers to be temporary, corresponding to the stimulus policy analysed in our empirical application below. We start by characterizing the effects of a tax stimulus in partial equilibrium, that is, taking the price levels p_t as given. The stimulus has three effects on the housing market as summarized in the following proposition:

Proposition 1. (Effects of Tax Stimulus on Demand). *A reduction of the period-1 transaction tax rate τ_1 , keeping τ_2 fixed, has three effects on demand in the housing market:*

- (i) *an **intensive margin effect** as households who move demand more housing;*
- (ii) *a **timing effect** as some households are induced to shift their house purchases from period 2 to period 1;*
- (iii) *an **extensive margin effect** as some households are induced to make an additional house purchase in period 1, without changing their planned house purchases in period 2. Because this increase in activity in period 1 is not offset by a reduction in period 2, there is incomplete reversal of the stimulus impact.*

Proof. See Online Appendix A.2. ||

We now consider how leverage (as captured by the loan-to-value ratio $1 - \gamma_1$) amplifies demand responses to transaction taxes. This can be characterized as follows:

Proposition 2. (Leverage Amplifies Demand Effects of Taxes). *Leverage $1 - \gamma_1$ amplifies the intensive margin, timing, and extensive margin effects of a temporary reduction in the transaction tax rate τ_1 among downpayment constrained households.*

Proof. See Online Appendix A.3. ||

Proposition 2 implies that small transaction taxes can generate large behavioural responses amongst liquidity constrained and leveraged households even if the structural demand elasticity e_1 is modest. It is straightforward to show that the intensive margin elasticity amongst downpayment constrained buyers is equal to $\varepsilon_{\tau_1}^{c_1} \equiv -\frac{dh_1^c}{d(1+\tau_1)} \frac{1+\tau_1}{h_1^c} = \frac{1+\tau_1}{\gamma_1+\tau_1}$. This implies that a household with a loan-to-value ratio of 75% (the modal value in the U.K. in our sample period, see Financial Conduct Authority, 2014) would have an elasticity of approximately 4, much larger than what might be a reasonable value for the structural elasticity e_1 . Similarly, extensive margin responses are stronger for more leveraged buyers. For highly leveraged buyers, even small changes in the transaction tax are leveraged into large changes in the house that they can afford, and so can have large impacts on transaction levels in the housing market.

We now turn from partial equilibrium demand effects (taking house prices as given) to the analysis of house price effects of tax stimulus.

Proposition 3. (Effects of Tax Stimulus on House Prices). *The elasticity of house prices in periods 1 and 2 with respect to the transaction tax rate in period 1 can be written as*

$$\left(\frac{dp_1/p_1}{d\tau_1/(1+\tau_1)} \right) = - \left(\frac{\frac{\partial D_1}{\partial p_1} p_1}{\frac{\partial D_2}{\partial p_1} p_1} \frac{\frac{\partial D_1}{\partial p_2} p_2}{\frac{\partial D_2}{\partial p_2} p_2} \right)^{-1} \left(\frac{\frac{\partial D_1}{\partial (1+\tau_1)} (1+\tau_1)}{\frac{\partial D_2}{\partial (1+\tau_1)} (1+\tau_1)} \right), \quad (10)$$

where the derivatives are given by the following averages of demand-weighted intensive and extensive margin elasticities:

$$\begin{aligned}
\frac{\partial D_1}{\partial(1+\tau_1)}(1+\tau_1) &= U_1 E[\varepsilon_{\tau_1}^{u_1} h_1^u | U_1] + C_1 E[\varepsilon_{\tau_1}^{c_1} h_1^c | C_1] + U_1 \eta_{\tau_1}^{u_1} E[h_1^u - 1 | U_1] + C_1 \eta_{\tau_1}^{c_1} E[h_1^c - 1 | C_1] \\
\frac{\partial D_2}{\partial(1+\tau_1)}(1+\tau_1) &= -U_2 E[\varepsilon_{\tau_1}^1 h_1 | U_2] + C_2 E[\varepsilon_{\tau_1}^{c_2} h_2^c - \varepsilon_{\tau_1}^1 h_1 | C_2] + U_2 \eta_{\tau_1}^{u_2} E[h_2^u - h_1 | U_2] \\
&\quad + C_2 \eta_{\tau_1}^{c_2} E[h_2^c - h_1 | C_2] \\
\frac{\partial D_1}{\partial p_1} p_1 &= U_1 E[\varepsilon_{p_1}^{u_1} h_1^u | U_1] + C_1 E[\varepsilon_{p_1}^{c_1} h_1^c | C_1] + U_1 \eta_{p_1}^{u_1} E[h_1^u - 1 | U_1] + C_1 \eta_{p_1}^{c_1} E[h_1^c - 1 | C_1] \\
\frac{\partial D_1}{\partial p_2} p_2 &= U_1 E[\varepsilon_{p_2}^{u_1} h_1^u | U_1] + U_1 \eta_{p_2}^{u_1} E[h_1^u - 1 | U_1] + C_1 \eta_{p_2}^{c_1} E[h_1^c - 1 | C_1] \\
\frac{\partial D_2}{\partial p_1} p_1 &= -U_2 E[\varepsilon_{p_1}^1 h_1 | U_2] + C_2 E[\varepsilon_{p_1}^{c_2} h_2^c - \varepsilon_{p_1}^1 h_1 | C_2] + U_2 \eta_{p_1}^{u_2} E[h_2^u - h_1 | U_2] \\
&\quad + C_2 \eta_{p_1}^{c_2} E[h_2^c - h_1 | C_2] \\
\frac{\partial D_2}{\partial p_2} p_2 &= U_2 E[\varepsilon_{p_2}^{u_2} h_2^u - \varepsilon_{p_2}^1 h_1 | U_2] + C_2 E[\varepsilon_{p_2}^{c_2} h_2^c - \varepsilon_{p_2}^1 h_1 | C_2] + U_2 \eta_{p_2}^{u_2} E[h_2^u - h_1 | U_2] \\
&\quad + C_2 \eta_{p_2}^{c_2} E[h_2^c - h_1 | C_2]
\end{aligned}$$

in which $\varepsilon_{x_t}^{u_s} \equiv \frac{\partial h_s^u}{\partial x_t} \frac{x_t}{h_s^u}$ is the intensive margin elasticity of unconstrained housing demand in period s with respect to $x_t \in p_t, (1+\tau_t)$; $\varepsilon_{x_t}^{c_s} \equiv \frac{\partial h_s^c}{\partial x_t} \frac{x_t}{h_s^c}$ is the intensive margin elasticity of constrained housing demand; $\varepsilon_{x_t}^1$ is the average intensive margin elasticity of period-1 housing demand; $\eta_{x_t}^{u_s} \equiv \frac{\partial U_s}{\partial x_t} \frac{x_t}{U_s}$ is the extensive margin elasticity of unconstrained movers; and $\eta_{x_t}^{c_s} \equiv \frac{\partial C_s}{\partial x_t} \frac{x_t}{C_s}$ is the extensive margin elasticity of constrained movers.

Proof. See Online Appendix A.4. \parallel

This result characterizes dynamic tax incidence in a housing market with downpayment constraints. The formula is based on a set of demand-weighted average elasticities at the intensive and extensive margins. To understand the proposition, it is instructive to start by considering the perfect market case without downpayment constraints ($\gamma_1 = \gamma_2 = 0$), transaction taxes ($\tau_1 = \tau_2 = 0$), or fixed costs of moving ($q_1 = q_2 = 0$). In this case there are no constrained movers ($C_1 = C_2 = 0$) and no extensive margin responses ($\eta_{x_t}^{c_s} = \eta_{x_t}^{u_s} = 0$). As a result, equation (10) reduces to $\frac{dp_1/p_1}{d\tau_1/(1+\tau_1)} = -1$ and $\frac{dp_2/p_2}{d\tau_1/(1+\tau_1)} = 0$, that is, the standard incidence result in a world with fixed supply.⁸ The simple reason for the standard result is that, for demand to be unchanged in equilibrium, it must be the case that $(1+\tau_1)p_1$ is unchanged.

8. As this incidence result is based on $\tau_1 = 0$, it represents the incidence of *small* transactions taxes. The reason why we need to start from $\tau_1 = 0$ to obtain full pass-through (whereas in standard models with fixed supply, full pass-through holds for any τ) is that in our model transaction taxes create liquidity constraints. Furthermore, the reason why we also need $q_1 = q_2 = 0$ to obtain full pass-through is that otherwise the standard extensive margin elasticities with respect to taxes and prices, respectively, are different: $\eta_{x_t}^{u_1} \neq \eta_{p_1}^{u_1}$. This difference arises because in our model, higher house prices have a positive wealth effect on movers but not on stayers (as the house depreciates to zero at the end of life) which makes moving decisions depend differently on taxes and prices.

While equation (10) is a very involved expression, the key force introduced by downpayment constraints comes from the intensive margin demand effects of prices and taxes on the constrained movers. The effect of prices on the constrained movers ($\varepsilon_{p_1}^{c_1}, \varepsilon_{p_2}^{c_2}$) is positive: higher prices *increase* demand by relaxing downpayment constraints. This effect was precisely the focus of the seminal Stein (1995) paper, but in a different context than tax incidence. Other things being equal, this effect tends to offset the standard demand effect on the unconstrained movers and makes aggregate excess demand less sensitive to prices. At the same time, the effect of transaction taxes on the constrained movers ($\varepsilon_{\tau_1}^{c_1}$) is negative: higher taxes reduce demand by reinforcing downpayment constraints. This reinforces the standard demand effect ($\varepsilon_{\tau_1}^{u_1}$) and makes aggregate excess demand more sensitive to taxes. These two effects combined make prices more sensitive to taxes, potentially by a lot if there are many constrained and highly leveraged households. The basic intuition is that higher taxes reduce demand by more in the presence of liquidity constraints, while lower prices are less effective at bringing demand back up. The end result is also affected by the extensive margin terms, the signs of which are ambiguous depending on whether marginal movers sell more or less housing than they buy, but the intensive margin terms create the possibility of strong leverage multiplier effects that make prices excessively sensitive to taxes. In special cases, one can show analytically that leverage unambiguously amplifies price responses to the transaction tax. For example,

Proposition 4. (Leverage Amplifies Price Effects of Taxes). *Starting from a situation in which the only distortion is a small downpayment constraint, leverage $1 - \gamma_1$ amplifies the house price effect of a temporary reduction in the transaction tax rate τ_1 .*

Proof. See Online Appendix A.5. ||

This proposition implies that transaction tax cuts can have large house price effects with highly leveraged households. Importantly, while these strong price responses re-establish equilibrium in the housing market, they do *not* undo the effect of transaction taxes on transaction levels (*i.e.*, the number of movers) and may actually reinforce it. The reason is that higher prices tend to reduce the number of constrained movers by relaxing downpayment constraints. This implies potentially strong multiplier effects of tax stimulus on transaction volumes: lower taxes directly increase transaction levels (more so if leverage $1 - \gamma_1$ is high) in response to which prices rise and potentially increase transaction levels even further (and again, more so if leverage $1 - \gamma_1$ is high). Hence, transaction taxes can have very large effects on both prices and transaction volumes in the housing market, consistent with our empirical findings below.

3. INSTITUTIONAL BACKGROUND AND DATA

3.1. *The U.K. property transaction tax: notches and reforms*

The U.K. property transaction tax—Stamp Duty Land Tax—is imposed on the transaction value of land and any construction on the land, known as the “chargeable consideration”.⁹ This is defined in the broadest possible terms to include anything of economic value given in exchange for land or property, including money, goods, works or services, and transfers of debts. The statutory incidence of the SDLT falls on the buyer, who is required to file a stamp duty return and remit tax

9. 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/sdlr/calculate/value.htm> for details.




FIGURE 1

Stamp duty schedule in 2012–13

Notes: The figure shows the stamp duty land tax schedule for residential properties in the tax year from 6 April 2012 to 5 April 2013. The tax liability jumps discretely at the notches at £125,000, £250,000, £500,000, £1,000,000, and £2,000,000. 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.

liability to Her Majesty's Revenue & Customs' (HMRC) within a few weeks of the completed transaction. The SDLT is a significant source of government revenue in the U.K., much more so than other wealth transfer taxes such as inheritance taxation and capital gains taxation. The SDLT has raised revenue of around 0.6% of GDP over recent years.¹⁰

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 threshold as in standard graduated schedules. Figure 1 illustrates the stamp duty schedule for residential property in the tax year 2012–13.¹¹ The schedule features five notches as the proportional tax rate jumps from 0 to 1% at a price of £125,000; from 1% to 3% at a price of £250,000; from 3% to 4% at a price of £500,000; from 4% to 5% at a price of £1,000,000; and

10. See http://www.hmrc.gov.uk/stats/tax_receipts/tax-receipts-and-taxpayers.pdf.

11. After this paper was first written, the notched stamp duty schedule has been replaced by a standard kinked schedule.

TABLE 1
Residential property tax notches

Date range Price range	1 December 2003 to 16 March 2005	17 March 2005 to 22 March 2006	23 March 2006 to 2 September 2008	3 September 2008 to 31 December 2009	1 Jan 2010 to 5 April 2011	6 April 2011 to 21 March 2012	22 March 2012 to April 2013
0 - 60K	0	0	0	0	0	0	0
£60K-£120K	1	1	1	1	1	1	1
£120K-£125K							
£125K-£175K							
£175K-£250K							
£250K-£500K	3	3	3	3	3	3	3
£500K-£1000K	4	4	4	4	4	4	4
£1000K-£2000K						5	5
£2000K-∞							7

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.

finally from 5% to 7% at a price of £2,000,000.¹² The schedule is different for residential property in certain disadvantaged areas (where the first bracket threshold is at a higher price) as well as for non-residential property. It is important to note that the buyer cannot mortgage the SDLT liability, it must be financed from savings, and so we should expect the SDLT to have particularly large effects on downpayment constrained buyers (as shown by our theoretical framework).

Another important aspect of the stamp duty is that it has been subject to a great deal of policy experimentation over the years. As shown in Table 1, the main policy experiments during our data period have been (1) changes in the location of the lower notch and (2) the introduction of new notches at £1,000,000 in April 2011 and at £2,000,000 in March 2012. It is worth describing the specific features of some of those policy changes as they will be important for the empirical analysis.

For the lower notch, the most salient change was the *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. The motivation for the programme was to provide housing stimulus during the Great Recession.¹³ The following 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 a stamp duty holiday in the month leading up to the announcement, the details and start date of such a holiday were unknown. Secondly, the end of this holiday was *anticipated*. The initial announcement was that the holiday would last for 1 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

12. At the £2,000,000 notch, the stamp duty rate jumps to 15% if the residential dwelling is purchased by certain “non-natural persons” such as corporations and collective investment schemes.

13. Another stimulus programme was implemented specifically for first-time buyers between 25 March 2010 and 24 March 2012. This programme temporarily abolished the notch at £125,000, thereby eliminating stamp duty in the range between £125,000 and £250,000 for first-time buyers.

commitment 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* (a discrete jump in tax liability at a cutoff date) allowing us to analyse short-term timing effects. Finally, 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 using a difference-in-differences design.

For the top notches, the introduction of a higher stamp duty rate above £1,000,000 was pre-announced a full year in advance, while the higher stamp duty rate above £2,000,000 was confirmed just 1 day before it took effect. Hence, the introduction of the £1,000,000 price notch (but not the £2,000,000 price notch) also creates a time notch that could be used to study anticipatory behaviour.

The U.K. stamp duty appears to be characterized by relatively high compliance. According to HMRC estimates, the so-called tax gap—the difference between taxes owed and taxes paid on a timely basis—is 4–5% of true stamp duty tax liability. This is lower than the tax gap estimates for most other taxes in the U.K. As described above, the tax base is defined in a very comprehensive manner meaning that the scope for shifting or re-classification of specific features of the property to avoid the tax is limited. The one exception is the exclusion from the tax base of freestanding “extras” such as furniture and curtains. If such extras are part of the sale, the buyer and seller are to agree on the market value of these extras and subtract it from the chargeable consideration, which creates an opportunity to evade stamp duty by overvaluing such items (while undervaluing the rest of the property by the same amount). Most stamp duty evasion is likely to occur through this channel, although the extent of this is limited by the fact that large amounts of tax exempt extras is an audit trigger. Whenever stamp duty evasion is not facilitated through these tax exempt extras, there must be monetary side payments between the buyer and the seller in order to cover the difference between reported and true house prices. Such side payments are associated with substantial risk and are therefore likely to be rare.¹⁴ Still, since we cannot rule out evasion responses, our bunching estimates of price responses should be interpreted as combining real responses with potential evasion responses analogously to the literature on taxable income responses (Saez *et al.*, 2012). However, our main findings and interpretations—for example, regarding the impacts of stimulus—do not rely on knowing the decomposition of bunching into evasion and real responses.

3.2. *Data and raw time series evidence*

The empirical analysis is based on administrative data covering the universe of stamp duty returns in the U.K. from November 2004 to October 2012. 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 U.K. The full data set contains about

14. Monetary side payments are very risky for two reasons. First, almost all property transactions in the U.K. are facilitated by licensed real estate agencies, implying that side payments require collusion between a buyer, a seller, and a real estate agency (typically with multiple employees). Such evasion collusion involving many agents is difficult to sustain (Kleven *et al.*, 2016). Secondly, the scope for monetary side payments is further reduced by the existence of a considerable lag between agreeing on a house price and completing the contract (this lag is 2–3 months on average in the U.K.). 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 to 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.

10 million transactions. The data set contains rich tax return information for each transaction, but has very little information outside the tax return.

The housing market has seen substantial turmoil during the period we consider. Figure 2 shows the monthly number of house transactions (Panel A) and the monthly average property price (Panel B) in all of the U.K. 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 1 at the start of the period. We make the following observations. First, housing market activity collapses between late 2007 and early 2009 as the number of transactions falls by around two-thirds. There has been some recent recovery, but activity is still far from pre-recession levels. Secondly, property prices also fall between late 2007 and early 2009, but the price drop is less dramatic and the subsequent recovery much stronger. Thirdly, property prices (though not activity) in London have evolved differently than in the rest of the U.K. during the recession. While U.K.-wide property prices have recovered only partially in the past couple of years, London property prices are almost back on their pre-recession trend. Fourthly, 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 be explained by the stamp duty holiday. Finally, average house prices in London feature a sharp spike in early 2011 and a subsequent dip, which constitutes our first piece of evidence of a behavioural response to stamp duty incentives. This spike reflects excess trading of houses above £1,000,000 just before the pre-announced introduction of the £1,000K stamp duty notch on 6 April 2011 and the dip reflects missing trading of such houses just after the introduction of the notch — a short-term timing response to an anticipated tax change.

4. HOUSE PRICE RESPONSES TO TRANSACTION TAXES: PRICE NOTCHES

4.1. *Static price notches*

This section presents static results using price notches during periods when they are stable. We consider residential property transactions that incur a stamp duty land tax liability.¹⁵ Figure 3 considers the two 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 an estimated counterfactual distribution (red line). Following Chetty *et al.* (2011) and Kleven and Waseem (2013), the counterfactual distribution is estimated by fitting a flexible polynomial to the empirical distribution, excluding data in a range around the notch, and allowing for round-number fixed effects to capture rounding in the price data.¹⁶ The excluded range is demarcated by vertical

15. Results for non-residential property are qualitatively similar, but noisier as we have far fewer observations.

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

$$c_i = \sum_{j=0}^q \beta_j (z_i)^j + \sum_{r \in \mathcal{R}} \eta_r I \left\{ \frac{\bar{v} + z_i}{r} \in \mathbb{N} \right\} + \sum_{k=\bar{h}_i^-}^{\bar{h}_i^+} \gamma_k I \{i=k\} + \mu_i, \quad (11)$$

where c_i is the number of transactions in price bin i , z_i is the distance between price bin i and the cutoff \bar{v} , and q is the order of the polynomial ($q=5$ in Figure 3). The second term in equation (11) includes fixed effects for prices that are multiples of the round numbers in the set \mathcal{R} , where $\mathcal{R} = \{500, 1,000, 5,000, 10,000, 25,000, 50,000\}$, \mathbb{N} is the set of natural numbers, and $I\{\cdot\}$ is an indicator function. Finally, the third term in (11) excludes a region (\bar{v}^-, \bar{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. Our estimate of the counterfactual distribution is defined as the predicted bin counts \hat{c}_i from equation (11) omitting the contribution

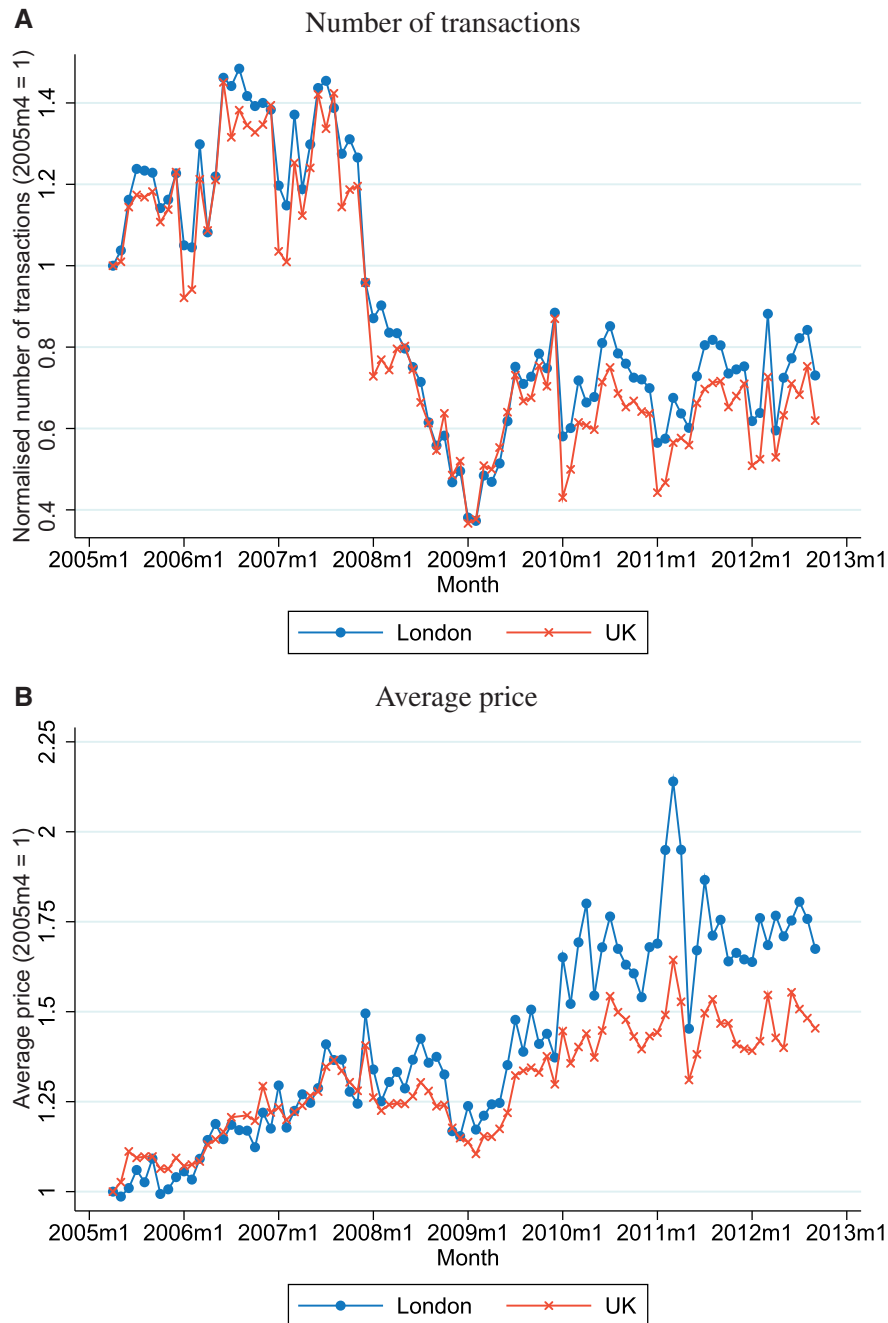


FIGURE 2
Descriptive statistics

Notes: Panel A shows the monthly average price of property transactions relative to the average price in April 2005 in London (blue circles) and the U.K. (orange crosses). The average price of property transactions in London during the period April 2005 to October 2012 was £345,360 and the average price in the U.K. during our data period was £199,479. Panel B shows the monthly total number of property transactions relative to the number that took place in April 2005 in London (blue circles) and the U.K. (orange crosses). The average monthly number of property transactions in London during the period April 2005 to October 2012 was 12,955 while the average monthly number of property transactions in this period in the U.K. was 103,561.

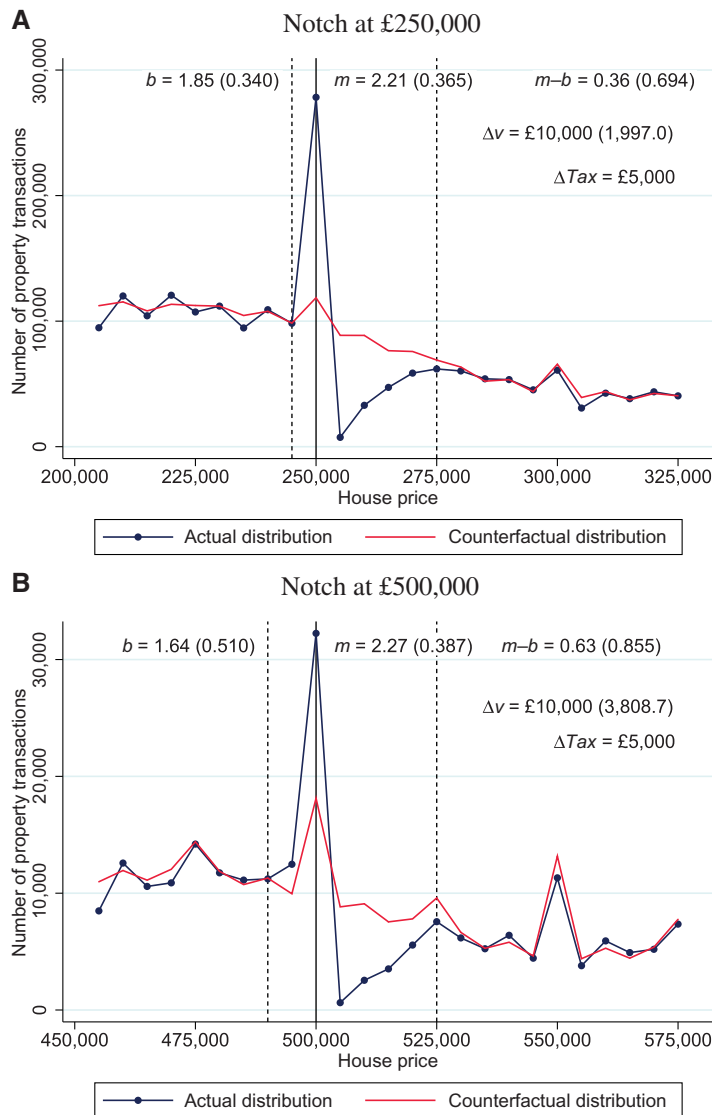


FIGURE 3

Bunching and holes around the notches that remain constant

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 price) in panel B. The data used for these estimates excludes transactions that claim relief from the stamp duty land tax (except for those claiming first-time buyers' relief) as the regular tax schedule does not apply to these transactions. The counterfactual density is estimated as in equation (11), using bins £100 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 polynomial 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 average counterfactual frequency in the excluded range, with its standard error shown in parentheses. m is our estimate of the missing mass above the notch scaled by the average counterfactual frequency in the excluded range, 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 Δv , and the tax liability change at the notch ΔTax . All standard errors are obtained by bootstrapping the procedure 200 times.

dashed lines; the lower bound is set at the point where excess bunching starts and the upper bound is set at the point where the hole ends (where the empirical distribution above the cutoff changes slope from positive to negative).

Before reporting our results, it is worth qualifying the estimated “counterfactual distribution” and what it is supposed to achieve in this context. Two points are worth noting. First, because the estimation uses data above the notch that may be affected by extensive responses, the procedure does not yield the full counterfactual distribution that would prevail if the notch were eliminated. As discussed by Kleven and Waseem (2013), the estimation procedure intends to provide a “partial counterfactual” stripped of intensive responses, but not extensive responses. We may think of the partial counterfactual as the border of the light-grey area in Panel D of Online Appendix Figure A.1.¹⁷ Secondly, even this partial counterfactual distribution may be difficult to estimate precisely in the range above the notch, especially in situations where the hole is large and diffuse (Kleven and Waseem, 2013; Kopczuk and Munroe, 2015). The estimation of missing mass above the notch may be quite sensitive to parametric assumptions such as the order of the polynomial and the width of the excluded range. For this reason, none of our key results in this article will rely on being able to estimate the counterfactual distribution and missing mass above the notch. The estimation of house price responses in this section as well as the estimation of extensive and timing responses in the next section make use only of our bunching estimates. And these bunching estimates are extremely robust due to the fact that observed bunching is very sharp and that the empirical distribution is very flat and smooth below the bunch. Our results are essentially unaffected by the degree of the polynomial, the excluded region, the binwidth, and whether we estimate bunching using only data below the notch. Hence, while we show estimates of the counterfactual distribution and missing mass in the following figures (as they may be of independent interest), it is important to keep in mind that these estimates are not strictly necessary for what we do.

In Figure 3, 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 price response to the notch ($\Delta \bar{v}$), and the tax liability change at the notch (ΔTax). Our main findings are the following. First, both notches create large and sharp bunching below the cutoff. Excess bunching is 1.85 and 1.64 times the height of the counterfactual distribution at £250,000 and £500,000, respectively, and is strongly significant in each case. Secondly, both notches are associated with a large hole in the distribution above the cutoff. The size of the hole is larger than the size of excess bunching, although the difference between the two is not statistically significant from zero. Thirdly, the hole in the distribution spans a £25,000 range

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 $\hat{B} = \sum_{i=\bar{v}^-}^{\bar{v}} (c_i - \hat{c}_i)$. We may also define an estimate of missing mass (the hole) above the notch as $\hat{M} = \sum_{i=\bar{v}}^{\bar{v}^+} (\hat{c}_i - c_i)$, but this statistic is not used in the estimation of house price responses and house price elasticities (see Online Appendix A.6). Standard errors on all estimates are calculated based on a bootstrap procedure as in Chetty *et al.* (2011). As a robustness check we have tried values between 4 and 7 for the order of the polynomial and our results are not significantly altered.

17. To simplify, our estimation of the counterfactual distribution ignores the marginal shift in the distribution above the hole due to intensive responses in the interior of the upper bracket. It is feasible to account for this shift in the distribution when estimating the counterfactual by using an initial estimate of the house price elasticity (based on ignoring the shift in the upper distribution) to obtain an initial estimate of the distribution shift, re-estimating the counterfactual and the house price elasticity to respect the initial estimate of the distribution shift, and continuing the procedure until the estimation converges. However, given the size of the incentive (a marginal tax rate change of 1–2% above the notch) and the house price elasticities that we find, this shift will be extremely small and have no substantive effect on any of our conclusions.

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. Fourthly, the average house price response is £10,000 at both the £250,000 notch and the £500,000 notch, a response that is twice as large as the tax jump. Interpreting the third and fourth results in the context of our model with downpayment constraints in section 2, a response of £25,000 is consistent with the most responsive households being downpayment constrained and having a loan-to-value ratio of 80%. In the U.K., many households have a loan-to-value (LTV) ratio of 80% as this level is associated with a large notch in mortgage interest rates (see Best *et al.* 2015). Similarly the average price response suggests an average LTV ratio amongst bunchers of 50%, consistent with a mixture of downpayment constrained, highly leveraged buyers, and unconstrained buyers.

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 £60,000 until 16 March 2005, at £120,000 between 17 March 2005 and 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 onward. This section takes a static approach by considering bunching responses within each of these 5 periods separately, while the next section investigates dynamic adjustment paths around the reform episodes. Figure 4 shows results for the 5 periods in separate panels, each of which is constructed as in Figure 3. The findings for the lower notch are qualitatively consistent with those for the other 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 £3,500 and £5,000, or about 4–5 times the size of the tax liability jump, implying that responses are proportionally larger at the bottom. Moreover, based on visual inspection of the hole (which is not as sharp here as for the upper notches), the most responsive households reduce their transacted house value by as much as 20 times the size of the notch. These extremely large responses are consistent with the most responsive households being downpayment constrained at an LTV ratio of 95%, not uncommon among first-time buyers in our sample period.¹⁸ These findings are consistent with the observation that a greater share of properties at the lower end of the price distribution are starter homes being bought by highly leveraged first-time buyers.¹⁹

4.2. *Moving price notches*

This section investigates the dynamics of behavioural adjustment to the changes in the position of the lower notch introduced by the stamp duty holiday. This stimulus policy moved the notch from £125,000 to £175,000 on 3 September 2008, and then moved it back from £175,000 to £125,000 on 1 January 2010. Recall that the start of the holiday was unanticipated while the end of the holiday was anticipated. When interpreting the findings, it is useful to keep in mind that there

18. For example, in 2006, 13% of mortgage originations were at LTV ratios above 90% (Financial Conduct Authority, 2014).

19. In 2011 and 2012, the government introduced two new notches affecting very high value properties, one at £1 million on 6 April 2011 and another one at £2 million on 22 March 2012. Even though these notches were introduced close to the end of our data period, we see very clear house price distortions in the data as shown in Online Appendix Figure A.2. This figure is constructed in the same way as the previous ones, except that the counterfactual distribution is obtained differently. We take advantage of the tax reform (notch introduction) by comparing the empirical house price distribution after the introduction of the notch to the empirical distribution in the year leading up to the introduction of the notch. The results are qualitatively very similar to the previous results, with an average house price response of £30,000 at the £1 million notch (3 times the tax liability jump of £10,000) and £100,000 at the £2 million notch (2.5 times the tax liability jump of £40,000).

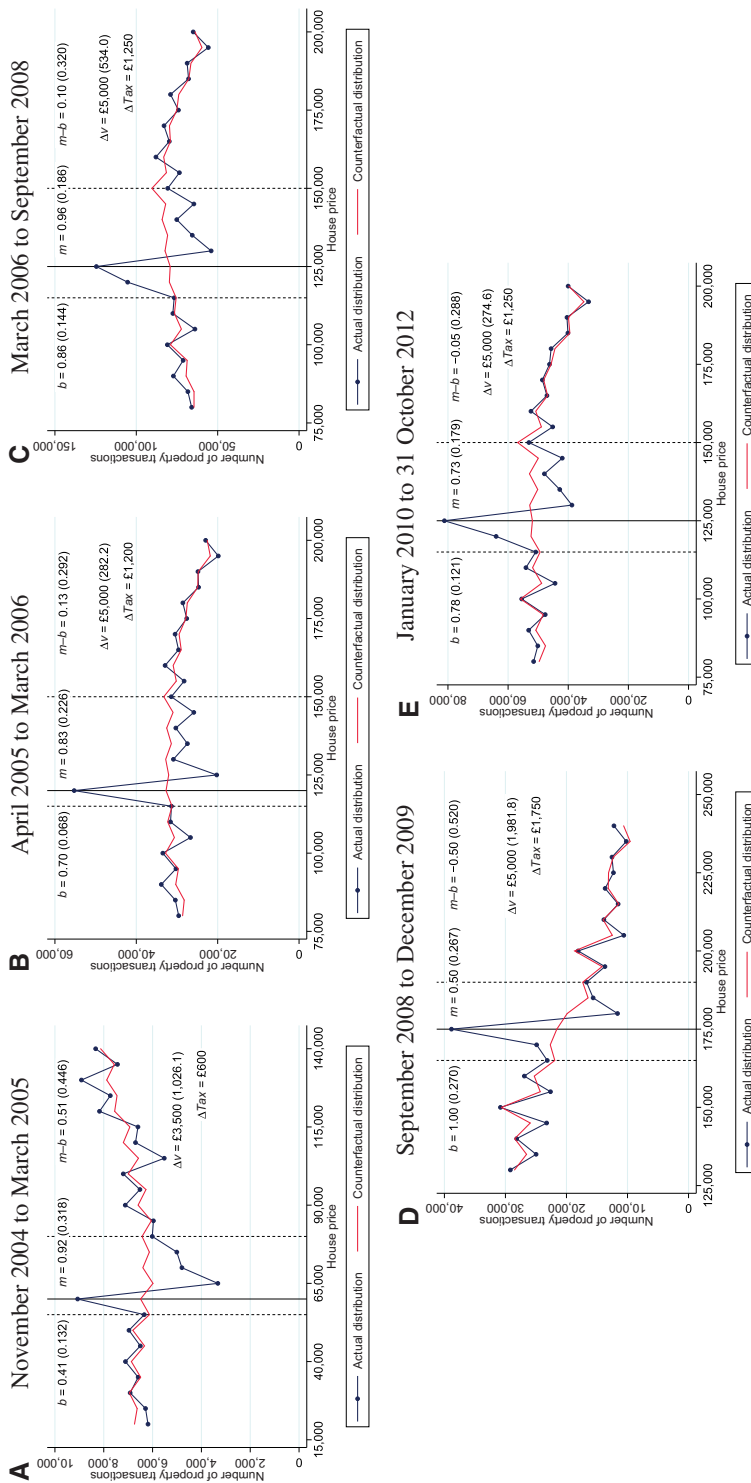


FIGURE 4

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. The notch is at £60,000 in panel A; £120,000 in Panel B; £125,000 in panels C & E, and £175,000 in panel D. The notch was moved on 17 March 2005, 23 March 2006, 3 September 2008, and 1 January 2010. The data used for these estimates excludes transactions that claim relief from the stamp duty land tax (excepting those who claimed first time buyers' relief) as the regular tax schedule does not apply to these transactions. The counterfactual is estimated as in equation (11), using bins £100 wide and a polynomial of order 5 in panels A, C, D, and E and of order 4 in panel B. 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 polynomial 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, and $m-b$ is our estimate of the difference between the missing mass and the bunching mass. The figures also show the average house value change created by the notch $\Delta \bar{v}$, and the tax liability change at the notch ΔTax . All standard errors are obtained by bootstrapping the procedure 200 times.




FIGURE 5
Dynamics of bunching over time

Notes: The figure shows our estimates of $b(\bar{v})$, the bunching mass just below \bar{v} scaled by the counterfactual frequency at \bar{v} , by month from January 2007 to August 2011 and for two values of \bar{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 typical lag of up to 90 days to conclude transactions, leading to inertia in bunching responses. The second vertical line is at December 2009 when the stamp duty holiday came to an end as anticipated, moving the notch from £175,000 back to £125,000.

is a lag between agreeing on a purchase price and completing the housing contract. In the U.K. housing market, this lag is under 90 days for most transactions and about 60 days on average (Besley *et al.*, 2014). 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 after an unanticipated tax change is bounded from below by about 3 months.

To analyse dynamics we estimate bunching separately for each month between January 2007 and August 2011 at both £125,000 and £175,000. Figure 5 summarizes the results by showing the monthly bunching estimate b at the £125,000 threshold (blue dots) and the £175,000 threshold (orange crosses) with 95% confidence intervals around each series. The solid vertical lines in Figure 5 demarcate the implementation dates of the reforms, while the dashed vertical line demarcates the *de facto* time at which the first, unanticipated reform took full effect given the lag between agreed and completed house purchases. The individual monthly bunching figures underlying the estimates if b in Figure 5 are presented in the Online Appendix Figures A.3 and A.4.²⁰

The following main findings emerge from the figure. First, at the beginning of the holiday it takes 3–4 months for bunching at £125,000 to disappear, and about 3 months for bunching at the new £175,000 threshold to build up to a steady state. These response times correspond roughly to the lower bound for unanticipated tax changes created by the contract completion lag described above. Secondly, steady-state bunching at £175,000 ($b \approx 0.9$) is larger than at £125,000 ($b \approx 0.6$), consistent with the former being a larger notch. Third, at the pre-announced end of the holiday, response times are even faster: Bunching at £175,000 disappears almost immediately

20. Animated versions of the monthly bunching figures that show the dynamics vividly can be found on our websites, for example, at http://www.henrikkleven.com/uploads/3/7/3/1/37310663/best-kleven_landnotches_april2013_videos.pdf

and the re-emergence of bunching at £125,000 happens within about 2 months. The immediate disappearance of bunching at £175,000 shows that buyers and sellers did indeed anticipate the end of the holiday and made sure to complete their housing contracts before the end of December 2009. The next section investigates such short-term timing behavior in greater detail.

The figure highlights how sharply house prices react to tax notches even at the monthly level. 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. Compared to recent bunching evidence from labour markets (*e.g.* Saez 2010; Chetty *et al.* 2011; Kleven and Waseem 2013; Gelber *et al.* 2015), the remarkable sharpness of our evidence suggests that behavioural 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. TIMING AND EXTENSIVE RESPONSES TO TRANSACTION TAXES: STIMULUS

We saw in the previous section that house prices respond sharply and quickly to transaction taxes. In this section, we investigate the effect of transaction taxes on transaction volumes, again studying the impact of the stamp duty holiday that temporarily eliminated the stamp duty on house purchases in the £125,000–£175,000 price range. The stamp duty holiday was an unanticipated stimulus programme with a fixed and fully anticipated end date. As shown in section 2, such a policy has two potential effects on transaction volumes. First, there will be a *timing effect* as some agents who would have transacted a house in the future bring that transaction forward to the current period. Second, there will be an *extensive margin effect* as some agents engage in additional house transactions over their lifetime.

To evaluate fiscal stimulus programmes of this kind, it is crucial to obtain estimates not just of the total stimulus effect during the programme (timing and extensive margin effects), but also of the degree to which it is driven by timing (all of which will be reversed after programme withdrawal) and the length of the horizon over which there is re-timing (which determines the speed of reversal). This section provides compelling evidence on all three questions. We also combine our stimulus estimates with survey data on moving-related household spending in order to estimate the effect on real economic activity. Our quasi-experimental research design implies that these are partial equilibrium effects, not including Keynesian multiplier effects or other general equilibrium aspects.

We begin the analysis by studying short-term timing responses around the anticipated end date of the stamp duty holiday, and then turn to the analysis of medium-term timing and extensive margin responses.

5.1. *Short-term timing responses to anticipated tax changes*

The fact that the tax increase at the end of the stamp duty holiday was fully anticipated implies a *time notch* on 1 January 2010 for houses between £125,000 and £175,000. This time notch creates an incentive for individuals to conclude their transactions before New Year, and bunching in the timing of transactions allows us to estimate this short-term timing response.

Before discussing the empirical results, we make two remarks. First, the housing market almost shuts down between Christmas and New Year, so the notch is effectively a notch just before Christmas. 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. Secondly, the existence of the Christmas holiday (with or without a tax notch) may in itself lead to a piling up of house

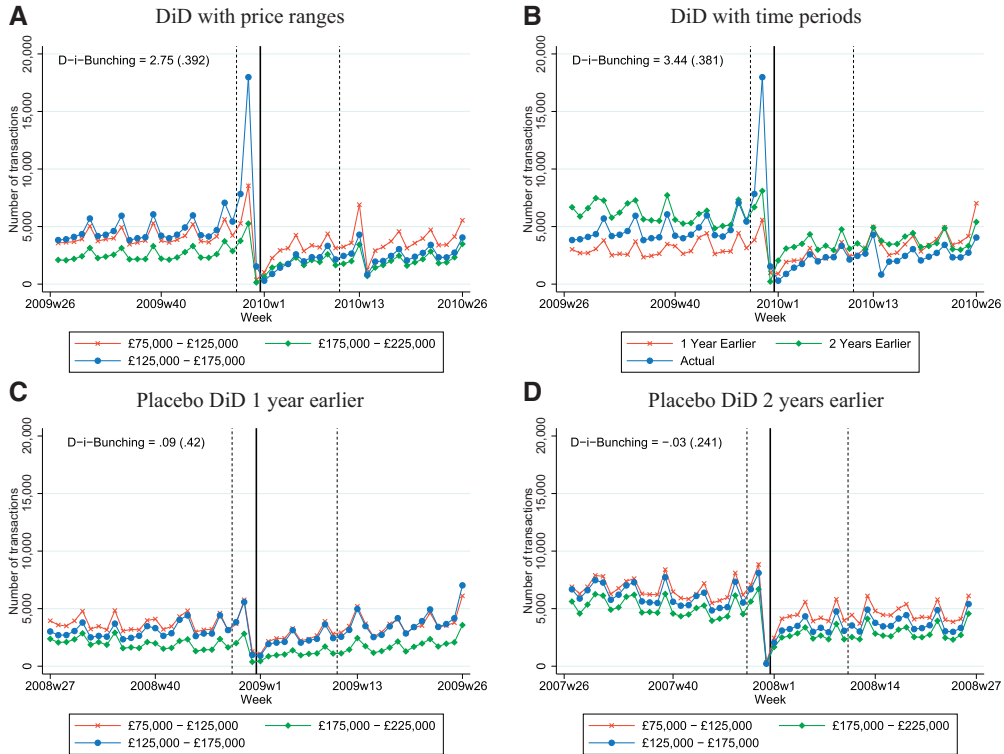


FIGURE 6
Time notch

Notes: The figures show the weekly number of transactions 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 in the treated price range £125,000 – £175,000 (blue circles) alongside the number of transactions in the price ranges £75,000 – £125,000 (orange crosses) and £175,000 – £225,000 (green diamonds). Panel B shows the number of transactions taking place in the treated price range (£125,000 – £175,000) 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 same price ranges as in panel A, but using data from 1 year earlier. Similarly, panel D shows 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. The counterfactual is estimated according to equation (12):

$$c_w = \sum_{j=0}^7 \beta_j (z_w)^j + \eta I\{w \in \text{end of month}\} + \sum_{k=\bar{w}^-}^{\bar{w}^+} \gamma_k I\{w=k\} + \mu_w$$

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 a fixed effect for weeks at the end of the month (which feature heavier trading in every month), while the third term excludes weeks in the excluded range (\bar{w}^- , \bar{w}^+). Each picture shows 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 (normalized) bunching in the treatment group and the average bunching in the two control groups.

transactions in the third week of December. This means that we cannot analyse the time notch using a plain 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 6 shows the weekly number of transactions around New Year in different price ranges and different years. Panel A compares the treated price range £125,000–£175,000 in the treated

period 2009/10 to surrounding price ranges in the 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 3 weeks of the year using a bunching approach analogous to our approach for the price notches.²¹ 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 in the treated period to the same price range in other periods (1 or 2 years earlier). The results are very similar, with estimated excess mass before the notch being somewhat larger and still strongly significant. The placebo tests in the bottom panels repeat the strategy in Panel A (comparing different price ranges), but 1 or 2 years earlier. In each case, the timing effect is close to 0 and statistically insignificant.

Overall, this provides very compelling evidence of short-term timing responses to anticipated tax changes, consistent with the sharpness of price responses discussed above. These findings contribute to the previous literature on the timing of the realisation of taxable income (Auerbach, 1988; Burman and Randolph, 1994; Goolsbee, 2000) and medical expenditures (Einav *et al.*, 2015).

5.2. Medium-term timing and extensive margin responses to stimulus

The stamp duty holiday temporarily eliminated the transaction tax in the price range £125,000—£175,000 without changing the tax in neighbouring price ranges, presenting us with an opportunity to pursue a difference-in-differences approach to estimate medium-term timing and extensive margin responses to the stimulus. A naïve first cut at this (that we refine shortly) is to compare the transaction volumes over time in the treated range £125,000–£175,000 and in a nearby control range. This is done in Figure 7, which shows the log number of transactions in the treated range £125,000–£175,000 (blue dots) and a control range defined as £175,000–£225,000 (orange crosses) in each month from September 2006 to October 2012. We have normalized the log number of transactions in each series by subtracting the average log number of transactions in the pre-treatment period (the 2 years leading up to the holiday) in order to make visual comparison of the two series easier. The solid vertical lines mark the beginning (3 September 2008) and the end (31 December 2009) of the stamp duty holiday.²²

21. To quantify bunching, we estimate a counterfactual number of weekly transactions based on the following regression

$$c_w = \sum_{j=0}^7 \beta_j (z_w)^j + \eta I\{w \in \text{end of month}\} + \sum_{k=\bar{w}^-}^{\bar{w}^+} \gamma_k I\{w=k\} + \mu_w, \quad (12)$$

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 a fixed effect for weeks at the end of the month (which feature heavier trading in every month), while the third term excludes weeks in a range (\bar{w}^-, \bar{w}^+) , which we set to include the last 3 weeks of 2009 and the first 10 weeks of 2010.

22. As described in section 3.1, a stamp duty relief scheme was implemented for first-time buyers in the price range £125K–£250K between 25 March 2010 and 24 March 2012 (*after* the end of the stamp duty holiday). Since we are also interested in estimating reversal after the stamp duty holiday, it is important to make sure that the first-time buyers' relief scheme is not a confounding factor during the reversal period. This motivates using a control range (£175K–£225K) just




FIGURE 7

Effects of the stamp duty holiday stimulus: naive diff in diff

Notes: The figure shows how the level of housing market activity changed over time in the price range affected by the stamp duty holiday (£125,000–£175,000) and the neighbouring price range £175,000–£225,000. The figure shows the normalized log monthly number of transactions 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 24 months leading up to the start date of the Stamp Duty Holiday (September 2006 to August 2008).

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 first months of the holiday and features a sharp spike in the last month as people rushed to take advantage of the stimulus before it expired. After the holiday, there is a sharp dip in the treated series during the first month, but only slight additional reversal thereafter as the treated group is marginally below the control group for about a year and then converges with the control group in the later part of the sample. Taken at face value, this graph implies that housing stimulus gave a large boost to housing market activity during the policy with very weak reversal after the policy (apart from the short-term timing effect shown by the spike and dip right around the stimulus end date analysed in section 5.1 above).²³ However, this overstates the positive impact of the stimulus policy and understates the slump after the end of the policy.

The issue with interpreting the evidence in Figure 7 is that treatment assignment (whether a transaction takes place in the £125,000–£175,000 price range) 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 range, the holiday creates a new notch at £175,000 that induces agents to move from a region above the cutoff to a point just below the cutoff (bunching). We have seen in section 4 that bunching responses at £175,000 do indeed occur,

above the treatment range (£125K–£175K), ensuring that both groups fall within the range eligible for first-time buyers' relief and therefore face the same incentive from this scheme. There could still be a concern that the treatment and control range respond differently to the first-time buyer incentive, which would be a confounding factor in the reversal estimates. To alleviate this concern, we drop all transactions claiming first-time buyers' relief throughout the analysis in this section. Including those observations only strengthens our findings below of incomplete reversal after the end of the stamp duty holiday.

23. Note that the control group also features a (much smaller) spike and dip around the end of the stamp duty holiday driven by the Christmas/New Year effect as discussed in section 5.1 above.

and this increases activity in the treated range compared to the control range. 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 seen that the disappearance of bunching at £125,000 also occurs, and this further increases activity in the treated range compared to the control range. Hence, the positive effect of housing stimulus in Figure 7 combines the true effect on overall activity levels with endogenous price responses resulting from the change in the location of the notch.

We consider two ways of dealing with this endogeneity issue. The first and simplest way is to widen the treatment range on each side (below £125,000 and above £175,000) in order to ensure that any price manipulation around notches occurs *within* the treatment range and so does not affect measured activity levels in this range. By including transactions outside the tax holiday area in the treatment group, this strategy captures an intent-to-treat effect and therefore understates the impact on the actually treated. We consider this intent-to-treat strategy in Figure 8 in which we expand the treatment range to £115,000–£195,000 and compare with the range £195,000–£235,000. Based on the bunching evidence presented earlier, this expanded treatment range fully encompasses the bunching interval below £125,000 and the hole interval above £175,000.

Figure 8 presents our findings in four panels. Panel A shows the normalized log number of transactions in each month in the treatment and control ranges (exactly as in Figure 7). Panel B shows the differences in these log counts in each month, thus highlighting the stimulus effects in a more striking way by differencing out the seasonal and cyclical aspects of the two time series. These two panels show that the intent-to-treat strategy produces stimulus effects that are qualitatively similar to those in the previous figure, but that the boost in activity during the holiday is smaller and the lull in activity after the holiday is larger. Panel B shows that the lull in activity lasts for approximately 12 months, after which the two series are parallel again. In order to illustrate the total extent of reversal after the holiday, Panel C depicts the cumulated sums of the flows from both Panel A (levels) and Panel B (differences). This graph shows clearly that reversal is far from complete: the cumulated differences fall by less than half of their initial rise during the 2.5-year post-stimulus period in the figure.

To put numbers on these effects, we run the following difference-in-differences regression on a panel of monthly activity levels in price bins of £5,000:

$$n_{it} = \alpha_0 Pre_t + \alpha_H Hol_t + \alpha_R Rev_t + \alpha_P Post_t + \alpha_T Treated_i + \beta_H Hol_t \times Treated_i + \beta_R Rev_t \times Treated_i + \beta_P Post_t \times Treated_i + v_{it}, \quad (13)$$

where n_{it} is the log number of transactions in price bin i and month t , Pre_t is a dummy for the pre-period September 2006 to August 2008, Hol_t is a dummy for the stamp duty holiday period September 2008 to December 2009, Rev_t is a dummy for the post-holiday reversal period January to December 2010, $Post_t$ is a dummy for the later months January 2011 to October 2012, $Treated_i$ is a dummy for the treated price range £115,000–£195,000, and finally v_{it} is an error term that we cluster by price bin to account for serial correlation (Bertrand *et al.*, 2004). The coefficients we are interested in are β_H (positive effect during stimulus) and β_R (negative effect after stimulus due to re-timing).

Panel A of Figure 8 shows our estimates of the coefficients β_H , β_R , and β_P . The coefficient $\hat{\beta}_H = 0.17$ (0.083) implies that average monthly activity was approximately 17% higher during the holiday than it would have been in the absence of stimulus. The coefficient $\hat{\beta}_R = -0.10$ (0.021) implies that average monthly activity was about 10% lower in the first year after the stimulus than it otherwise would have been. Together, these estimates imply that $-(12\hat{\beta}_R)/(16\hat{\beta}_H) = 42\%$ of the additional activity created by the stimulus programme was a timing response by people

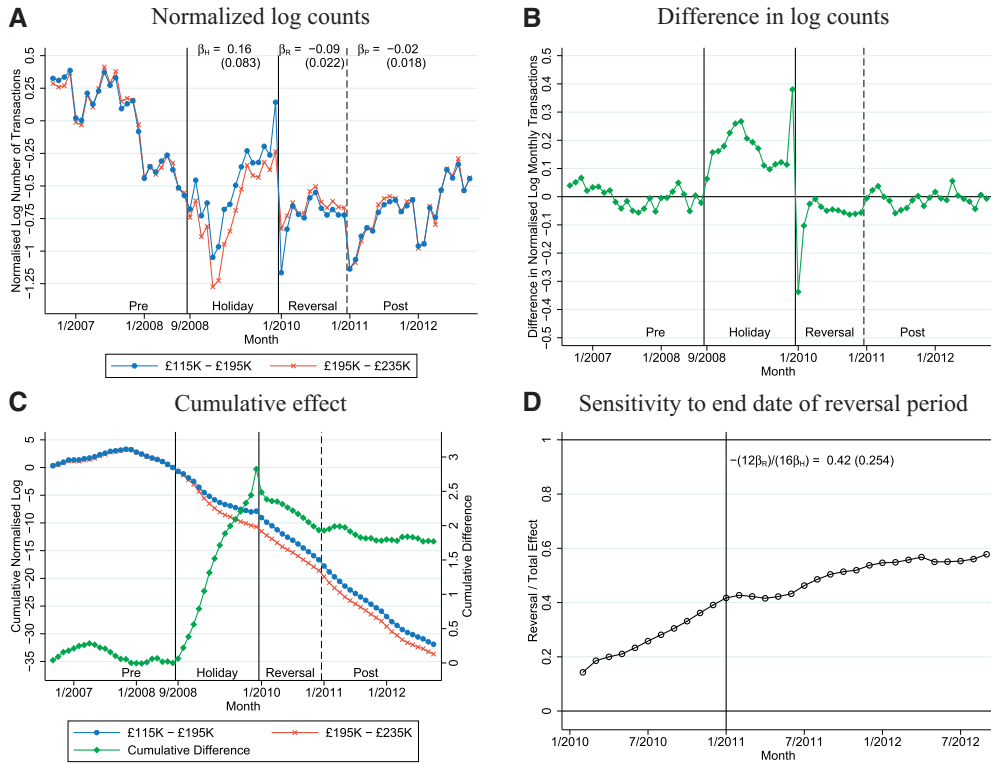


FIGURE 8

Effects of the stamp duty holiday stimulus: diff in diff with wider treatment range

Notes: The figure shows the effect of the stamp duty holiday stimulus on market activity using £115,000–£195,000 as the treated price range (blue circles) and £195,000–£235,000 as the control range (orange crosses). Panel A shows the log monthly number of transactions normalized by subtracting the average log number of transactions in the 24 months leading up to the stamp duty holiday (September 2006 to August 2008). We also show estimates of β_H , β_R and β_P from equation (13) with standard errors clustered by £5,000 bin to account for serial correlation (Bertrand *et al.*, 2004). Panel B shows the differences between the monthly counts in panel A (blue dots and orange crosses) and the cumulative sum of the normalized log counts in panel A (blue dots and orange crosses) and the cumulative sum of the differences in panel B (green diamonds). Panel D shows how the proportion of the effect of the stamp duty holiday that is undone by reversal changes as we vary the end date of the reversal period, that is, it shows $(\beta_R \sum_t Rev_t) / (16\beta_H)$ For different end dates of the period defining Rev_t . The vertical line is our preferred choice for the first month of *Post*, January 2011, giving an estimate of the proportion of the effect undone by reversal of 0.42 (0.234).

bringing forward their purchases in order to benefit from the tax cut, while the remaining 58% was a lasting extensive margin effect. These large and lasting effects are consistent with our theoretical framework in section 2 in which even small changes in the transaction tax are leveraged up into large changes in the house that downpayment-constrained households can afford, and hence can have large impacts on their decision to participate in the housing market.

As the end date of the reversal period (December 2010) was chosen visually, there might be a concern that our estimate of total reversal is sensitive to the chosen end date. In order to address this, Panel D of Figure 8 shows total reversal estimates for all possible end dates, with the vertical line corresponding to the baseline assumption of 12-months reversal.²⁴ The graph shows that the

24. The point estimates are obtained by performing the regression (13) using different reversal period cutoffs and calculating $-(\sum_t Rev_t \times \hat{\beta}_R) / (16\hat{\beta}_H)$, where $\sum_t Rev_t$ denotes the length of the reversal period in the particular regression.

reversal estimate is below 60% in all specifications and hence the finding of incomplete reversal is robust.

We now turn to the second and more sophisticated way of dealing with endogenous movements across notches. This strategy exploits the fact that we have monthly bunching estimates of price responses to notches and can therefore directly control for them. That is, we may consider the number of transactions in different price brackets adjusted for the effect of bunching behaviour in each month. To be precise, in every month, the estimated bunching mass just below £125,000 is reallocated to the treatment range £125,000–£175,000 while the estimated bunching mass just below £175,000 is reallocated to the control range £175,000–£225,000. By using these bunching-adjusted counts in our difference-in-differences strategy, we avoid bias from selection into treatment.

The results of this bunching-adjusted strategy are presented in Figure 9, which is constructed exactly as the previous figure. We see that the effects are qualitatively similar, but slightly stronger than those from the intent-to-treat strategy: the holiday boost is about 20% (as opposed to 17% before) and the reversal is about 8% for 12 months (as opposed to 10% before).²⁵ These effects reduce our estimate of total reversal as a share of total stimulus, and the estimate is even more robust to the chosen end date, never rising above 40%. These estimates of the boost and reversal effects can be converted to extensive margin elasticities in the short run (boost effect) and long run (boost effect net of reversal). The short-run elasticity equals $\hat{\eta} = \frac{\hat{\beta}_H}{\Delta\tau/(1+\tau)} = 20.62$ (2.18), while the long-run elasticity equals $\hat{\eta} = \frac{\hat{\beta}_H + (12/16)\hat{\beta}_R}{\Delta\tau/(1+\tau)} = 14.3$ (3.26).

To be clear, the strength of our difference-in-differences evidence lies in the following features of the data: the completely parallel trends of the treatment and control groups before *and* after the stimulus policy, along with the sharp changes occurring in the exact months that the policy is first introduced and subsequently withdrawn. It is very unusual for difference-in-differences studies to have such double verification of parallel trends before and after, along with sharp treatment effects at the monthly (rather than annual) level. For our estimates to be driven by confounding shocks, there would have to be both positive and negative shocks that affect neighbouring price bands differently and according to a very specific monthly pattern.

To eliminate any lingering concerns about confounders, Figure 10 shows a series of placebo estimates of the holiday effect β_H . These estimates are based on a difference-in-differences regression like (13), but instead of having only two groups (treated and untreated) in a relatively narrow price range, we allow for each 10K price bin to have its own holiday effect and consider a wider price range. To be precise, we estimate the holiday effect β_H for each 10K bin between £115,000–£375,000 using the range £375,000–£425,000 as the comparison group. If the underlying trends were non-parallel across the house price distribution, this would result in non-zero estimates of β_H even as move up through the untreated part of the distribution (*i.e.* above £195,000 when accounting for bunching responses just above the treated range). The graph lends strong support to our identification strategy: the series of placebo estimates are consistently close to zero and are almost always statistically insignificant.

25. Since these estimates require us to run the difference-in-differences regression (13) using bunching-adjusted activity levels in £5K bins, we have to reallocate bunching mass below the two cutoffs to specific £5K bins above the cutoffs. We reallocate bunching mass below a cutoff to the five bins above the cutoff in proportion to the amount of missing mass (difference between the estimated counterfactual mass and the observed mass) in each bin. Furthermore, since activity levels are adjusted using *estimated* bunching at the thresholds, we are introducing measurement error to our dependent variable coming from potential misspecification of the counterfactual when calculating the amount of bunching at £125K and £175K. However, since this measurement error is effectively noisy in the dependent variable, it does not cause bias in our estimates, but simply increases our standard errors.

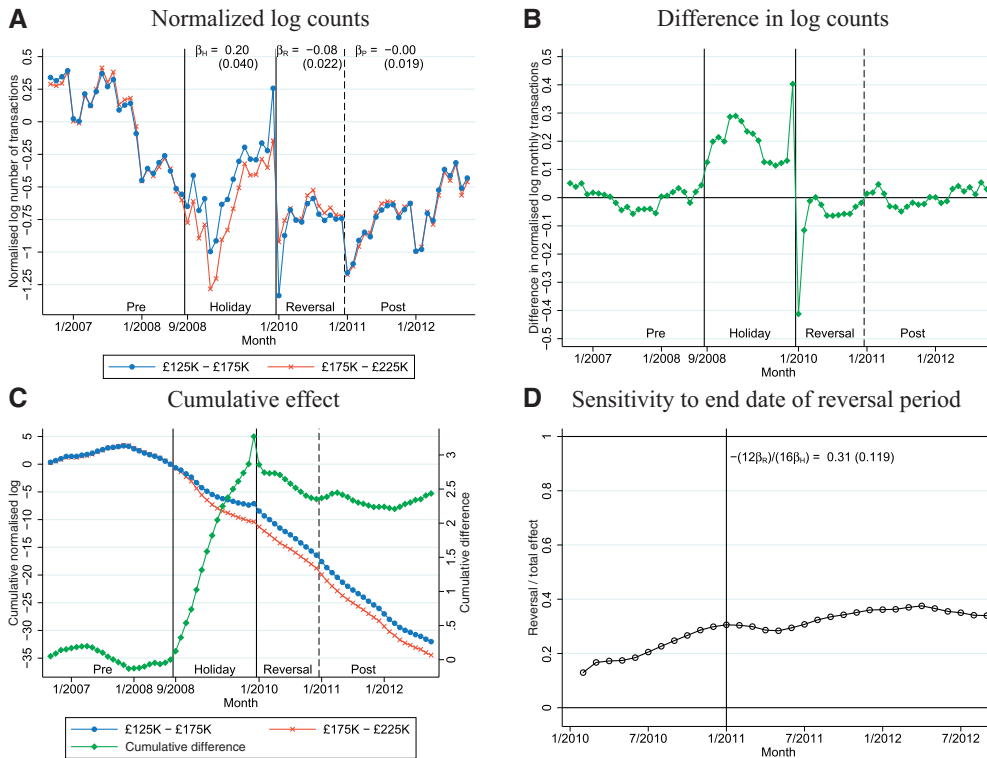


FIGURE 9

Effects of the stamp duty holiday stimulus: adjusting for bunching

Notes: The figure shows the effect of the stamp duty holiday stimulus on housing market activity using £125,000–£175,000 as the treated price range and £175,000–£225,000 as the control price range, but adjusting all counts for price manipulation using bunching estimates by moving excess transactions at £125,000 to prices between £125,000 and £150,000 and moving excess transactions at £175,000 to prices between £175,000 and £200,000. Panel A shows the log monthly number of transactions normalized by subtracting the average log number of transactions in the 24 months leading up to the stamp duty holiday (September 2006 to August 2008). We also show estimates of β_H , β_R and β_P from equation (13) with standard errors clustered by £5,000 bin to account for serial correlation (Bertrand *et al.*, 2004). Panel B shows the differences between the monthly counts in panel A. Panel C shows the cumulative sum of the normalized log counts in panel A (blue dots and orange crosses) and the cumulative sum of the differences in panel B (green diamonds). Panel D shows how the proportion of the effect of the stamp duty holiday that is undone by reversal changes as we vary the end date of the reversal period, that is, it shows $(\beta_R \Sigma_t Rev_t) / (16\beta_H)$ For different end dates of the period defining Rev_t . The vertical line is our preferred choice for the first month of $Post_t$, January 2011, giving an estimate of the proportion of the effect undone by reversal of 0.31 (0.119).

In addition to providing placebo tests, Figure 10 sheds light on the potential for spillovers between the treated price range and neighboring price ranges. Spillovers between treatments and controls may arise from the presence of *real estate chains*, that is, linked house transactions whereby households buying new houses are simultaneously selling their existing houses (exactly as modelled in section 2). Such chains introduce the possibility that the stimulus-induced transactions in the treatment range trigger additional transactions outside that range. Such spillovers would make us underestimate the total effect of the tax holiday. However, Figure 10 suggests that chain spillovers outside the holiday region are not significant: because chain spillovers should be larger close to the treated region than farther away, they would create a declining series of placebo estimates above the treated range. The fact that we do not see such a declining pattern suggests that the stimulus effects are local to the directly treated range. It should be noted, however, that we cannot rule out chain effects in the *immediately* surrounding bins in which bunching responses depress transaction volumes during the holiday and thus hides potential




FIGURE 10

Placebo effects of the stamp duty holiday stimulus on other price ranges

Notes: The figure shows a series of placebo estimates of the holiday effect β_H in an extended version of the difference in difference regression (13) where we allow each 10K price bin between £115,000–£375,000 to be a separate treatment group and use the £375,000–£425,000 range as the control group. While the estimates in the treated range £125,000–£175,000 are positive and strongly significant, the placebo estimates in other price ranges are consistently close to zero and almost always statistically insignificant, supporting the identification strategy.

spillovers. The possibility of such effects provides an argument for choosing the intent-to-treat strategy over the bunching-adjusted strategy, because the former would include any spillovers.²⁶

Our stimulus findings stand in contrast to Mian and Sufi (2012), who find complete and swift reversal following a short (1 month) stimulus programme offering car transaction subsidies in the U.S. Apart from the fact that the housing market may work differently than the car market (for example with respect to downpayments and leverage), the contrast between our findings and those of Mian and Sufi (2012) suggests that stimulus policies that are of extremely short duration do not give households sufficient time to respond along the extensive margin and therefore have only short-term timing effects. Hence, our findings highlight the importance of the length of the stimulus programme.²⁷ Of course, while the strength of reversal is important for evaluating stimulus, it does not by itself indict or validate such policies as their key rationale is to create

26. Note that, while the log effect β_H is smaller under the intent-to-treat strategy, the baseline number of transactions is larger due to the wider treatment range. Hence the stimulus effect in terms of additional transactions created is actually somewhat larger under the intent-to-treat strategy than under the bunching-adjusted strategy, consistent with moderate local spillovers.

27. Consistent with this, Online Appendix Figure A.5 uses the same strategy to analyse a permanent reform early on in our data period. On 16 March 2005, the bottom notch was permanently moved from £60,000 to £120,000. The reform cut the tax from 1% to 0% over the price range £60,000 to £120,000 while leaving the tax unchanged in neighbouring price ranges, allowing us to use our bunching-adjusted difference in differences strategy. Panel A shows the normalized log counts of monthly transaction volumes in the treatment range £60,000–£120,000 (blue circles) and the control range £120,000–£180,000 (orange crosses) together with the estimated treatment effect from a regression analogous to equation (13), while panel B shows the cumulative sums of the normalized log counts in the treatment and control ranges. The estimated coefficient $\hat{\beta}_P = 0.23$ (0.018) implies that this permanent reform increased monthly transaction volumes by approximately 23% on average. This extensive response is even larger than the extensive response to the stamp duty holiday, consistent with the idea that these effects are increasing in the length of the time horizon of the policy as discussed

more economic activity when the economy is slack (even if this comes at the expense of less economic activity when the economy is tight). The next section provides a rough estimation of the immediate increase in real economic activity created by the U.K. housing stimulus programme.

5.3. *Spending effects of stimulus*

While we have established that the stamp duty holiday had a large effect on transaction volume in the housing market (and therefore on household mobility), part of the motivation for the policy was also to stimulate household spending through the potential complementarities between moving house and spending. In particular, moving house is associated with substantial spending on items such as repairs and improvements, removals, furniture, appliances, and commissions. Investigating the spending effect of the U.K. housing stimulus programme also allows for a comparison between our findings and previous work on the spending effect of fiscal stimulus such as income tax rebates (*e.g.* Shapiro and Slemrod 2003a,b; Johnson *et al.* 2006; Agarwal *et al.* 2007; Parker *et al.* 2013).

To estimate the effect of housing stimulus on moving-related spending, we use the U.K. Living Costs and Food Survey (LCFS). This data set includes detailed consumption information and is a repeated cross-section covering 2001–13.²⁸ While in principle one could directly estimate the effect of the stamp duty holiday on spending in a difference-in-differences design, in practice such a strategy is not feasible due to limitations of the survey data such as small sample size, annual rather than monthly information, and unobserved house prices. We therefore adopt an alternative strategy that combines our difference-in-differences estimate of the effect on residential transaction volume (moving) with an “event-study” estimate of the effect of moving on spending using the survey data. This strategy exploits the power and precision of the administrative data as far as it goes, and then leverages the survey data only to estimate the last step from moving to spending.

As a first descriptive exercise, Table 2 compares spending in moving-related categories in the year of moving (year 0), in the year after moving (year 1), and in subsequent years (year 2+). Columns (1)–(3) show the *level* of spending in these different years, while columns (4)–(6) show the *additional* spending in year 0 and 1 (as compared to year 2+). Panel A of the table shows that households spend more on all categories in the year of moving, and that they continue to spend more on repairs and improvements (and only slightly more on other categories) in the year after moving. These numbers are consistent with the idea that households incur substantial moving-related expenditures in the year that they move, and that renovations extend into the following year. The total amount of additional spending in years 0 and 1 equals £6,428. Panel B converts average spending into percentages of the average house price (£230,000), adds commissions to estate agents (an average of 1.98%) and other service providers such as solicitors (an average of 1.24%), thus arriving at an initial estimate of the impact of moving on spending.²⁹ We find that moving is associated with extra spending of about 5% of the house value in the first year and an additional 1% in the next year.

in section 5.2. It is also worth noting that this reform was implemented during the height of the housing market boom, in sharp contrast to the stamp duty holiday implemented at the bottom of the recession.

28. We exclude the 2008 survey round, because in this year the housing tenure variable (which we use in the empirical strategy described below) is corrupted.

29. A 2011 survey by *Which? Magazine* estimates that estate agents’ fees average 1.8% of the house price before VAT, or 1.98% with VAT (see <http://www.which.co.uk/news/2011/03/estate-agents-fees-exposed-248666/>). ReallyMoving (2012) estimates that other commissions and fees total £1,880 on average, and do not vary much with house value, so we scale this by the average value of houses bought in the range affected by the policy (£152,000).

TABLE 2
Moving-related spending 2001–13

	Spending levels			Movers' additional spending		
	Year 0 (1)	Year 1 (2)	Year 2+ (3)	Year 0 (4)	Year 1 (5)	Year 0 + 1 (6)
Panel A: spending categories						
Repairs and Improvements	2,790	3,383	1,667	1,123	1,716	2,839
Removals and Storage	204	31	2	202	29	231
Furnishings	3,482	1,454	934	2,548	520	3,068
Appliances	226	163	130	96	33	128
Other durables	577	543	479	98	64	162
Total spending	7,279	5,574	3,213	4,067	2,361	6,428
Panel B: total moving-related spending (percentage of house value)						
Total spending				1.77	1.03	2.79
Estate agent commissions				1.98	0	1.98
Other commissions				1.24	0	1.24
Impact of move on spending				4.99	1.03	6.01

Notes: The table shows the moving-related spending by movers in the 2001–13 rounds of the LCFS (excluding the 2008 round in which the housing tenure variable is corrupted). Panel A shows spending in five moving-related categories and total spending on all five categories. Column (1) shows spending by households who moved within the last year, column (2) shows households who moved between 1 and 2 years ago, and column (3) shows households who moved over 2 years ago. Column (4) shows the difference between spending by households who moved in the last year and those who moved over 2 years ago (subtracting column (3) from column (1)). Column (5) shows the difference between spending by households who moved between 1 and 2 years ago and those who moved over 2 years ago (subtracting column (3) from column (2)). Column (6) adds columns (5) and (6). Panel B shows moving-related spending as a fraction of the value of the house. The “Total spending” row shows the last row of panel A as a fraction of the average house price (£230,000). Estate agents’ fees average 1.8% of the house price before VAT, or 1.98% with VAT (see <http://www.which.co.uk/news/2011/03/estate-agents-fees-exposed-248666/>). ReallyMoving (2012) estimates that other commissions and fees total £1,880 on average, and do not vary much with house value, so we scale this by the average value of houses bought in the range affected by the policy (£152,000).

There are two main issues with interpreting these numbers that we now address. One issue is that moving-related spending may be cyclical, and in particular that movers spend less during recession periods such as the one we study. To investigate this, Figure 11 summarizes the results of performing the calculation in Table 2 separately for each year. The series in grey triangles show movers’ additional spending across all categories in year 0 (as in column (4) of the table), while the series in black dots show movers’ additional spending across all categories in year 0 and 1 combined (as in column 6 of the table). We see that moving-related spending is indeed cyclical, with more spending when housing market activity is high (2003–06 and 2013) than when housing market activity is low (2001–03 and 2007–12). In light of this, in what follows we present estimates separately for low-activity years and for all years pooled.

Another issue is that moving may be endogenous to factors that impact directly on spending such as income shocks or childbirth. We can alleviate such concerns in two ways: by controlling for a number of observables available in our survey, and by performing a pseudo event study analysis based on sharp changes close to the time of moving. Denoting the number of years since last move by s (“event time”), we consider the following specification

$$\log c_{is} = \sum_{s=0}^{21} \alpha_s \text{Move}_{is} + X_{is} \beta + \mu_q + v_{is}, \quad (14)$$




FIGURE 11
Cyclicity of moving related household spending

Notes: The figure shows the differences in spending on moving-related categories (repairs and improvements, removals, furniture, appliances, and commissions) between recent movers and those who moved more than 2 years ago separately for each round of the LCFS from 2001/02 to 2013. The grey triangles show the difference between households who moved within the last year and those who moved more than 2 years ago. The black circles show the spending by those who move within the last year plus those who move between 1 and 2 years ago, minus the spending of those who moved more than 2 years ago.

where c_{is} is spending in moving-related categories by household i who moved between s and $s+1$ years ago, $Move_{is}$ is an indicator for household i moving between s and $s+1$ years ago, X_{is} is a vector of controls (log income, baby in the household, family size, age of the household head, and marital status), μ_q are quarter-of-interview fixed effects, and ν_{is} is an error term clustered by quarter of interview.

Figure 12 shows the α_s -coefficients with 95% confidence intervals estimated from equation (14). Panel A pools all years, while panel B restricts the sample to years with low housing market activity (2001–03 and 2007–12). The two panels show a large and sharp effect on spending in year 0, a more modest effect on spending in year 1, and very small and mostly insignificant effects in all other years. The moving effects are smaller during recessions, but not by very much. While we cannot control for within-individual shocks that affect spending (absent panel data), the sharpness of spending effects in year 0–1 tend to rule out serially correlated confounders (such as job switches or promotions that change income permanently) as they would change spending levels beyond year 0 and 1. Our identification relies on the assumption that any omitted consumption determinants are smooth around the exact year of moving. Our estimates imply that moving triggers a spending increase of at least 140% in the year of moving and 50% in the subsequent year.³⁰

Two robustness checks are worth considering. First, the estimates in Figure 12 represent average effects in the population, but the stamp duty stimulus targeted the lower end of the house price distribution where effects might be different. While the LCFS data does not contain information on house prices, we can look at heterogeneity by income as a proxy for house

30. These numbers are based on the estimates $\hat{\alpha}_0 = 0.86$ and $\hat{\alpha}_1 = 0.40$ from the event study regression restricting the sample to years with low housing market activity.

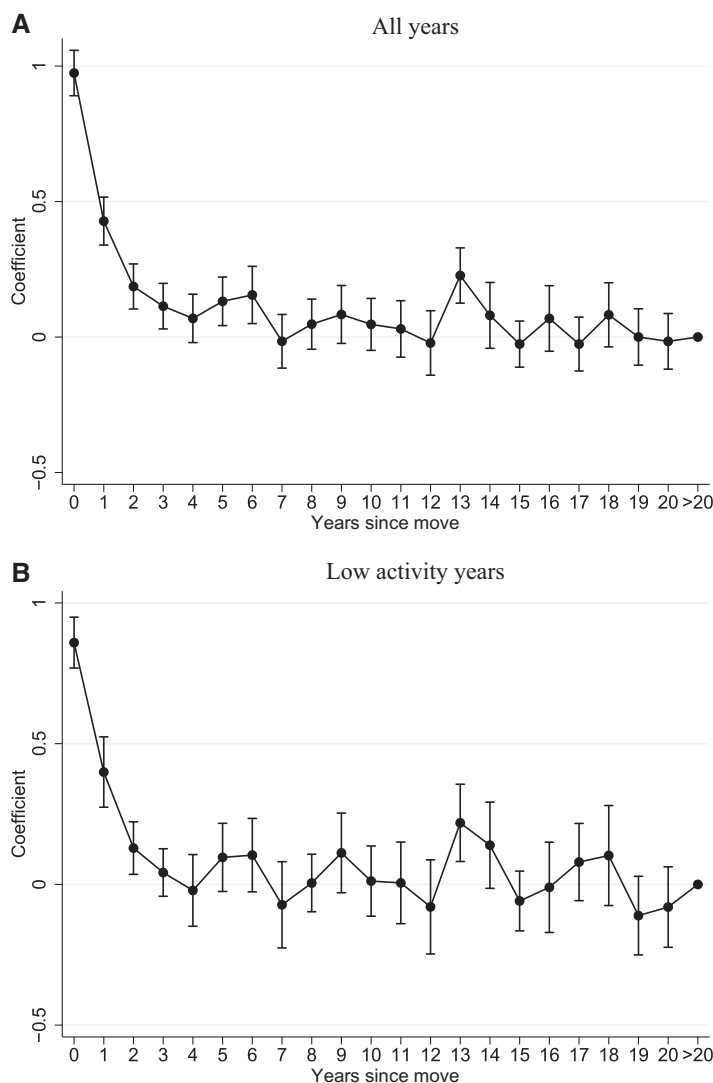


FIGURE 12

Event studies of household spending around moving

Notes: The figure shows the estimated $\hat{\alpha}_s$ coefficients on spending by movers in different years from two estimates of our event study specification (14) along with their 95% confidence intervals. The event study compares log spending on categories related to moving by households with different move dates, controlling for income, household size, age and marital status of the household head, the presence of babies in the household, and quarter of interview fixed effects. Panel A shows the estimates pooling all available years 2001–13 except 2008 (when the housing tenure data is corrupted) and panel B shows the estimates using only years with low housing market activity (2001–03 and 2007–12).

price heterogeneity. Online Appendix Figure A.6 shows that the log effects ($\hat{\alpha}_0$ and $\hat{\alpha}_1$) are larger for households with lower income, consistent with the presence of a fixed cost element in moving-related expenditures. Hence, by using estimates of average moving-related spending (in percentages), we are being conservative in our estimates of the stimulus' impact on spending.

Secondly, a potential concern is that moving-related spending (*e.g.* renovations) merely crowds out other forms of spending (*e.g.* restaurant visits) so that total consumer spending is

not larger when moving.³¹ To address this concern, Figure A.7 in the Online Appendix repeats the event study analysis using total consumption rather than only explicitly moving-related consumption. We still see a very sharp increase in spending in year 0, but no longer in year 1. That is, households who are still incurring moving-related expenditures the year after moving (primarily renovations as discussed above) may be offsetting this with lower consumption in other categories. Hence, a conservative approach to estimating the spending effect of the stamp duty holiday is to only include moving-related spending in year 0.³²

Combining these estimates with our earlier estimates of the impact of the stimulus policy on residential housing transactions (moving), we are able to assess the amount of spending generated by each dollar of foregone tax revenue. Denoting our estimate of moving-related spending as a fraction of the house value by ϕ ,³³ the impact of the stimulus policy on spending is $\Delta C = \phi v^m \Delta n$ where v^m is the average value of houses bought due to the stimulus, and Δn is the number of additional houses bought due to the policy. The foregone tax revenue is given by $\Delta Tax = \tau v^m n$, where $\tau = 0.01$ is the pre-stimulus tax rate and n is the counterfactual number of transactions in the price range affected by the stimulus. From our difference-in-differences analysis in the previous section, we have an estimate of $\beta_H \approx \Delta n/n$, and so we can estimate the effect on spending per dollar of tax cut as $\Delta C / \Delta Tax = \phi \beta_H / \tau$.

The results are shown in Table 3. Panel A of the table provides a range of estimates of ϕ from the event study analysis, while Panel B provides a range of tax multipliers $\phi \beta_H / \tau$ using different values of ϕ and β_H . The tax multipliers are around 1 or slightly higher. Our preferred multiplier is based on the bunching-adjusted estimate of β_H ($=0.20$) and a conservative spending parameter ϕ based on the year-0 effect alone, a full set of controls, and low-activity years ($\phi = 5.09$). This gives an estimate of the spending effect per dollar of tax cut equal to 1.02. As a robustness check, Online Appendix Table A.1 reports tax multipliers using ϕ estimates based on total consumption. Reassuringly, the estimates are very similar when including only the year-0 effect, with our preferred specification giving a multiplier of 1.10.

These estimates suggest that the stamp duty holiday was not only successful in stimulating housing market activity, but also provided a significant boost to real economic activity through the strong complementarities between moving house and consumer spending. As mentioned earlier, our quasi-experimental research design cannot capture potential general equilibrium effects of stimulus. Previous work using research designs that also identify partial equilibrium effects found significantly smaller effects of tax stimulus on consumer spending (*e.g.* Johnson *et al.* 2006;

31. The effect could also go the other way: moving house may increase spending in categories that are not explicitly moving related as households are induced to re-optimize in various dimensions (say, buy a new car or bicycles for the kids).

32. Besides affecting moving-related spending via the number of movers, the stimulus policy may also affect spending on repairs and improvements conditional on moving. First, sellers who would normally sell their house for a price just above £175,000 may reduce spending on improvements to remain below the £175,000 notch. Secondly, sellers who would normally sell for just below £125,000 may increase spending on improvements as they no longer have to avoid the £125,000 notch. Thirdly, all households buying in the range £125,000–£175,000 pay less tax and may increase spending. Our bunching estimates imply that the first two effects combined much be close to zero, even if we assume that bunching responses consist entirely of spending effects. Bunching at £175,000 is $b = 1.00$ (see Figure 4) corresponding to a spending reduction of £5,000, while bunching at £125,000 is $b = 0.86$ corresponding to a spending increase of £4,300. Defining the fraction of stimulus-affected households at £125,000 as $\gamma_{125} \equiv g_0(125K) / \sum_{125K}^{175K} g_0(v)$ and analogously for £175,000, the total effect is $\gamma_{125} \times 4300 - \gamma_{175} \times 5000 \approx 0$. This leaves only the third, positive effect, and by ignoring this effect our spending estimates are conservative.

33. We calculate ϕ based on the log coefficients from the event studies ($\hat{\alpha}_0$ alone or $\hat{\alpha}_0, \hat{\alpha}_1$ together) evaluated at the average spending of people who moved more than 2 years ago (μ_{2+}) and scaled by the average house value in the period ($v^M = £230,000$). We then add in the estate agents' fees (1.98%) and other fees (1.24%) to arrive at $\phi = (\exp(\hat{\alpha}_0) - 1 + \exp(\hat{\alpha}_1) - 1) \cdot (\mu_{2+} / v^M) + 1.98 + 1.24$ (when including both the year-0 and year-1 effects).

TABLE 3
Spending effects per dollar of tax cut

		Year 0		Year 0 + 1	
		No controls	Controls	No controls	Controls
Panel A: move-related spending: ϕ (% of house value)					
All years		6.29	5.49	7.56	6.23
Low activity years		5.72	5.09	6.93	5.77
Panel B: spending per dollar of tax cut: $\beta_H \cdot \phi / \tau$					
All years	$\beta_H = 0.20$	1.26	1.10	1.51	1.25
	$\beta_H = 0.17$	1.07	0.93	1.29	1.06
	$\beta_H = 0.20$	1.14	1.02	1.39	1.15
Low activity years	$\beta_H = 0.20$	0.97	0.87	1.18	0.98
	$\beta_H = 0.17$				

Notes: The table shows the spending effects of the stamp duty holiday using only the moving-related spending categories listed in Table 2. Panel A shows the impact of moving on spending ϕ (percentage of house value) as implied by our different event-study specifications, while Panel B shows estimates of tax multipliers—spending effect per dollar of tax cut—given by $\phi\beta_H/\tau$. We calculate ϕ based on the log coefficients from the event studies ($\hat{\alpha}_0$ alone or $\hat{\alpha}_0, \hat{\alpha}_1$ together) evaluated at the average spending of people who moved more than 2 years ago (μ_{2+}) and scaled by the average house value in the period ($v^M = £230,000$). We then add in the estate agents' fees (1.98%) and other fees (1.24%) to arrive at $\phi = (\exp(\hat{\alpha}_0) - 1 + \exp(\hat{\alpha}_1) - 1) \cdot (\mu_{2+}/v^M) + 1.98 + 1.24$ (when including both the year-0 and year-1 effects). The columns labelled Year 0 use only $\hat{\alpha}_0$ while the columns labelled Year 0 + 1 use both $\hat{\alpha}_0$ and $\hat{\alpha}_1$. The columns labelled Controls include controls for log income, baby in the household, family size, age of the household head and marital status. The rows labelled “low activity years” use only data from 2001–03 to 2007–12.

Agarwal *et al.* 2007; Parker *et al.* 2013). Overall, our findings suggest that transaction tax cuts (or subsidies) can be very effective at stimulating both housing market activity and real economic activity during downturns.

6. CONCLUSION

This article has studied the impact of housing transaction taxes based on administrative stamp duty records and quasi-experimental variation from notches and stimulus. We have shown that the housing market responds very strongly and quickly to transaction taxes across a range margins, making such taxes very distortionary compared to standard recurrent taxes. We have interpreted our results in the context of a housing model with downpayment constraints (building on Stein, 1995), arguing that the liquidity and leverage channels amplify behavioural responses to transaction taxes and provides a natural explanation for the large behavioral responses we find.

Our most important and novel findings are those related to the 2008–09 stamp duty holiday. These findings contribute to the large macro literature on fiscal stimulus, presenting some of the first evidence on the effectiveness of using temporary tax changes to stimulate the housing market during economic downturns. Our stimulus analysis carries two qualitative lessons. First, stimulating transactions of existing houses may have real effects due to the complementarities between moving house and spending. These complementarities combined with the large effect of transaction taxes on transaction volume create a direct tax multiplier of about 1, not including potential multiplier effects. That is, the direct spending effect of this tax stimulus was about the same as the direct effect one would obtain from traditional spending stimulus. This is important for policy because tax stimulus has political and practical advantages over spending stimulus, because of the aversion to spending increases and the fact that the implementation of spending projects can have considerable time lags. For these reasons, U.S. stimulus over the Great Recession had a

much larger weight on tax stimulus than on spending stimulus, including a first-time homebuyer tax credit similar in spirit to the transaction tax cut analysed here.

Secondly, our findings provide a counter-example to Mian and Sufi (2012), who find complete and swift reversal following the Cash-for-Clunkers stimulus programme in the U.S. Why are the experiences from these two policies so different? Two differences stand out. One is that housing stimulus may work differently than car stimulus due to the amplifying role of downpayment constraints in the housing market as shown by our theoretical model. Of course, financial frictions are also present in car markets, but they may be more severe in housing markets and especially at the time of the Great Recession. The other difference relates to practical implementation. Mian and Sufi (2012) considered a very short (1 month) and unanticipated programme that gave households very little time to respond.³⁴ It seems natural that such a programme only induced short-term intertemporal substitution by households who were already planning to buy a new car in the near future. As shown by our (frictionless) model, a temporary price subsidy *should* create some extensive margin effect in addition to short-term intertemporal substitution. However, when allowing for optimization frictions in responses to short-term price changes, it is possible that the Cash-for-Clunkers programme was over too quickly to trigger an extensive margin effect. Hence, our findings highlight the importance of the length of the stimulus programme. More generally, our article suggests the need for more work studying the details of programme design and implementation that make stimulus effective.

Finally, our study of transaction taxes in the property market contributes to the scant evidence on the effects of transaction taxes in asset markets more broadly, including the financial transaction taxes that have been discussed widely in recent years. An interesting question there turns on the ability of such taxes to affect the emergence of asset price bubbles and the volatility of the economy more generally. Addressing this issue raises some daunting empirical challenges, ideally requiring exogenous variation in transaction taxes across economies, and so is left for future research.

Acknowledgments. We thank Tim Besley, Raj Chetty, Julie Cullen, Michael Devereux, Amy Finkelstein, Gita Gopinath, Roger Gordon, Jonathan Gruber, Daniel Hamermesh, Benjamin Keys, Wojciech Kopczuk, Camille Landais, Attila Lindner, Bruce Meyer, Atif Mian, David Munroe, Benjamin Olken, James Poterba, John Van Reenen, Emmanuel Saez, Jesse Shapiro, Andrei Shleifer, Joel Slemrod, William Wheaton, anonymous referees, and numerous seminar participants for helpful comments and discussions. 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. This work contains statistical data from HMRC which is Crown Copyright. The research data sets used may not exactly reproduce HMRC aggregates. The use of HMRC statistical 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'.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

- AGARWAL, S., LIU, C. and SOULELES, N. S. (2007), "The Reaction of Consumer Spending and Debt to Tax Rebates – Evidence from Consumer Credit Data", *Journal of Political Economy*, **115**, 986–1019.
- ANDREWS, D., SÁNCHEZ, A. C. and JOANSSON, Å (2011), "Housing Market and Structural Policies in Oecd Countries" (Working Papers No. 836, OECD Economics Department).
- AUERBACH, A. J. (1988), "Capital Gains Taxation in the United States: Realizations, Revenue, and Rhetoric", *Brookings Papers on Economic Activity*, **1988**, 595–637.
- BERTRAND, M., DUFLO, E. and MULLAINATHAN, S. (2004), "How Much Should We Trust Difference-in-Differences Estimates?" *Quarterly Journal of Economics*, **119**, 249–275.

34. The programme was signed into law on June 24 of 2009 (with little prior public debate about the programme) and was effective between July 24 to August 24.

- BESLEY, T., MEADS, N. and SURICO, P. (2014), "The Incidence of Transactions Taxes: Evidence from a Stamp Duty Holiday", *Journal of Public Economics*, **119**, 61–70.
- BEST, M. and KLEVEN, H. (2014), "Housing Market Responses to Transaction Taxes: Evidence from Notches and Stimulus in the U.K." (Working Paper, May 2014).
- BEST, M., CLOYNE, J., ILZETZKI, E. and KLEVEN, H. (2015), "Interest Rates, Debt and Intertemporal Allocation: Evidence from Notched Mortgage Contracts in the U.K." (Working Paper).
- BURMAN, L. E. and RANDOLPH, W. C. (1994), "Measuring Permanent Responses to Capital-Gains Tax Changes in Panel Data" *American Economic Review*, **84**, 794–809.
- CAMPBELL, J. Y. and FROOT, K. A. (1994), "International experiences with securities transaction taxes", Chap. 6, in Frankel, J. A., (ed.), *The Internationalization of Equity Markets* (University of Chicago Press) 277–308.
- CHETTY, R. (2012), "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labour Supply", *Econometrica*, **80**, 969–1018.
- CHETTY, R., FRIEDMAN, J., OLSEN, T. and PISTAFERRI, L. (2011), "Adjustment Costs, Firm Responses, and Micro VS. Macro Labour Supply Elasticities: Evidence from Danish Tax Records", *Quarterly Journal of Economics*, **126**, 749–804.
- CUNNINGHAM, C. R. and ENGELHARDT, G. 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, B., DURANTON, G. and TURNER, M. 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.
- EINAV, L., FINKELSTEIN, A. and SCHRIMPF, P. (2015), "The Response of Drug Expenditure to Non-linear Contract Design: Evidence from Medicare Part D", *Quarterly Journal of Economics*, Forthcoming.
- EUROPEAN COMMISSION (2013), "Proposal for a Council Directive Implementing Enhanced Cooperation in the Area of Financial Transaction Tax", http://ec.europa.eu/taxation_customs/taxation/other_taxes/financial_sector/index_en.htm.
- FELDSTEIN, M., SLEMOD, J. and YITZHAKI, S. (1980), "The Effects of Taxation on the Selling of Corporate Stock and the Realization of Capital Gains", *Quarterly Journal of Economics*, **94**, 777–791.
- FINANCIAL CONDUCT AUTHORITY (2014), "Mortgage Product Sales Data", <http://www.fca.org.uk/static/documents/psd/mortgages-2014-data-extended.xlsx>.
- GELBER, A., JONES, D. and SACKS, D. W. (2015), "Earnings Adjustment Frictions: Evidence from the Social Security Earnings Test" (Mimeo: UC Berkeley).
- GOOLSBEE, A. (2000), "What Happens When You Tax the Rich? Evidence from Executive Compensation", *Journal of Political Economy*, **108**, 352–378.
- HOUSE, C. L. and SHAPIRO, M. D. (2008), "Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation", *American Economic Review*, **93**, 737–68.
- JOHNSON, D. S., PARKER, J. A. and SOULELES, N. S. (2006), "Household Expenditure and the Income Tax Rebates of 2001", *American Economic Review*, **96**, 1589–1610.
- KLEVEN, H. (2016), "Bunching", *Annual Review of Economics*, **8**, 435–464.
- KLEVEN, H. and WASEEM, M. (2013), "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan", *Quarterly Journal of Economics*, **128**, 669–723.
- KLEVEN, H., KREINER, C. and SAEZ, E. (2016), "Why Can Modern Governments Tax So Much? An Agency Model of Firms as Fiscal Intermediaries", *Economica*, **83**, 219–246.
- KOPCZUK, W. and MUNROE, D. (2015), "Mansion Tax: The Effect of Transfer Taxes on the Residential Real Estate Market", *American Economic Journal: Economic Policy*, **7**, 214–257.
- KRAINER, J. (2001), "A Theory of Liquidity in Residential Real Estate Markets", *Journal of Urban Economics*, **49**, 32–53.
- LINCOLN INSTITUTE OF LAND POLICY (2014), "Significant Features of the Property Tax", http://www.lincolninstitute.org/subcenters/significant-features-property-tax/Report_Real_Estate_Transfer_Charges.aspx.
- LUNDBORG, P. and SKEDINGER, P. (1999), "Transaction Taxes in a Search Model of the Housing Market", *Journal of Urban Economics*, **45**, 385–399.
- MATHESON, T. (2011), "Taxing Financial Transactions: Issues and Evidence" (Working Paper WP/11/54, IMF).
- MIAN, A. and SUFI, A. (2012), "The Effects of Fiscal Stimulus: Evidence from the 2009 Cash for Clunkers Programme", *Quarterly Journal of Economics*, **127**, 1107–1142.
- NGAI, L. R. and TENREYRO, S. (2014), "Hot and Cold Seasons in the Housing Market", *American Economic Review* forthcoming.
- PARKER, J. A., SOULELES, N. S., JOHNSON, D. S. and McCLELLAND, R. (2013), "Consumer Spending and the Economic Stimulus Payments of 2008", *American Economic Review*, **103**, 2530–2553.
- PIAZZESI, M. and SCHNEIDER, M. (2009), "Momentum Traders in the Housing Market: Survey Evidence and a Search Model", *American Economic Review Papers & Proceedings*, **99**, 406–411.
- POTERBA, J. M. (2002), "Taxation, Risk-taking, and Household Portfolio Behavior", in Chap. 17, Auerbach, A. J. and Feldstein, M., (eds), *Handbook of Public Economics*, vol. 3 (North Holland: Elsevier Science Publishers) 1109–1171.
- REALLYMOVING (2012), "Moving Costs Report September 2012", <http://www.reallymoving.com>.
- SAEZ, E. (2010), "Do Taxpayers Bunch at Kink Points?", *American Economic Journal: Economic Policy*, **2**, 180–212.
- SAEZ, E., SLEMOD, J. and GIERTZ, S. (2012), "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review", *Journal of Economic Literature*, **50**, 3–50.

- SHAN, H. (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.
- SHAPIRO, M. D. and SLEMROD, J. (2003a), "Consumer Responses to Tax Rebates", *American Economic Review*, **93**, 381–396.
- SHAPIRO, M. D. and SLEMROD, J. (2003b), "Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys", *Tax Policy & The Economy*, **17**, 83–109.
- SLEMROD, J., WEBER, C. and SHAN, H. (2014), "The Behavioral Response to Housing Transfer Taxes: Evidence from a Notched Change in D.C. Policy" (Working Paper, University of Michigan).
- STEIN, J. C. (1995), "Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects", *Quarterly Journal of Economics*, **100**, 379–406.
- VAN OMMEREN, J. and VAN LEUVENSTEIJN, M. (2005), "New Evidence of the Effect of Transaction Costs on Residential Mobility", *Journal of Regional Science*, **45**, 681–702.
- WHEATON, W. (1990), "Vacancy, Search, and Prices in a Housing Market Matching Model", *Journal of Political Economy*, **98**, 1270–1292.
- ZWICK, E. and MAHON, J. (2014), "Do Financial Frictions Amplify Fiscal Policy? Evidence from Business Investment Stimulus" (Working Paper).