How Do Stamp Duties Affect the Housing Market?*

Andrew Leigh
Economics Program
Research School of Social Sciences
Australian National University
Email: andrew.leigh@anu.edu.au
Website: http://andrewleigh.org

This version: 10 November 2009

Abstract
Land transfer taxes are a substantial portion of the cost of moving house in many developed countries. However, little is known about the effect of such taxes on the housing market. Since stamp duties are endogenous with respect to the house price, I create an instrumental variable that is the stamp duty on a property, given that postcode's starting house price and the national house price trend. In a specification with postcode and year fixed effects, this instrument effectively captures policy changes and nonlinearities in the stamp duty schedule. I find that the impact of an increase in the tax rate is to lower house prices, with the magnitude of the effect rising slightly over the medium run. I also observe impacts of stamp duty on housing turnover. A 10 percent increase in stamp duty lowers turnover by 1-2 percent in the first year, and by 4-5 percent if sustained over a 3 year period.

Keywords: tax incidence, land sales taxation, residential mobility
JEL Codes: H22, H24, H71, R21, R23, R28

* I am grateful to seminar participants at the National Tax Association's 2008 meetings in Philadelphia and RMIT University for valuable comments on earlier drafts.

1. Introduction

A key insight of public economics has been to demonstrate that the economic incidence of a tax can differ from its statutory incidence. Put another way, the person who ends up paying a tax may not be the person upon whom the tax is levied. Studies of tax incidence have variously shown that payroll taxes are mostly borne by workers, that retail sales taxes are mostly borne by consumers, and that corporate taxes are borne by consumers, workers, and investors.

Over recent decades, many developed countries have made increasing use of land transfer taxes, also known as stamp duties. Yet little is known about the impact of these taxes on housing prices. From an administrative standpoint, these taxes are typically levied on the buyer (ie. the statutory incidence of the tax is on the purchaser). But is the economic incidence of stamp duty entirely on the buyer, entirely on the seller, or shared between both parties? And what impact do higher house prices have on housing turnover?

From a theoretical perspective, inelastic factors bear the economic burden of taxes. Thus if buyers are more price-inelastic than sellers, then buyers will bear most of the tax burden (and house prices will not change much in response to a change in house sales taxes). Conversely, if sellers are more price-inelastic than buyers, then sellers will bear most of the tax burden (and house prices will fall by most of the value of a change in house sales taxes). Regardless of incidence, theory also predicts that higher taxes will increase the 'tax wedge' between buyers and sellers, and reduce total sales.

In this paper, I investigate the impact of stamp duties, using data from Australia, a jurisdiction where stamp duty averages around 3 percent of the property value. I exploit a rich dataset containing the full universe of housing sales over a 13-year period, which happens to be one of the periods of most rapid increase in Australian property values.

An advantage of using Australian data is that service delivery is largely homogenous across jurisdictions, due to federal formulas that equate funding across states and territories. This reduces the probability that changes in tax rates are correlated with changes in the quality of service delivery. Another factor that makes this unlikely is that stamp duty represents a relatively small portion of a state's revenue base. While stamp duty bill are substantial from the perspective of those paying them, stamp duty revenue amounts to less than 1 percent of personal income during all the years in my analysis.

To preview my results, I find that stamp duties reduce house prices and turnover rates. The effect of these taxes seems to be larger close to state boundaries, where there is more competition from the neighbouring jurisdiction. The price impacts imply that the incidence of stamp duty is mostly on the seller.

2. Data and Empirical Specification

The data used in this study were purchased from Australian Property Monitors (APM), which is Australia's leading firm that compiles house price data. APM obtains data from state and territory Valuer-General's offices, which is then cleaned by supplementing it with information from real estate agents (via an arrangement that APM has with the Real Estate Institute). The cleaning process is necessary because the data from the Valuer-General's offices sometimes has non-credible sales figures (e.g., sales that are an order of magnitude higher or lower than other recent sales in the same street), or is incomplete in some important detail (e.g., missing a street number). In many cases, errors in the Valuer-General's database can be corrected by reference to data held by real estate agents.[1] Following the cleaning process, APM estimates that their database covers more than 95 percent of all house sales.[2]

Since APM do not sell their full database, this analysis is based upon postcode-level means rather than data for individual sales.[3] Consequently, I am unable to control for

---

[1] For example, if a sale was listed as $35,000 by the Valuer-General and the real estate agent database states that the last asking price was $370,000, APM might assume that the correct sale price was $350,000.

[2] The source of this information is email and telephone conversations with Eva Knight, APM's Head of Research and Analytics during the period when I purchased the data.

[3] A small number of postcodes overlap state borders. In these cases, I have separate data for sales on either side of the border, and I treat them as separate units of

changes in quality from one year to the next. To help address this concern, I exclude sales of units (condominiums) from the dataset, since units are likely to be more heterogeneous (within a given neighbourhood) than houses. Units comprise only about one-fifth of the sales in the sample frame (about 1 million of the 5 million sales that underlie the dataset).

The coverage of the APM dataset varies across Australia's eight states and territories, being 1993-2005 for the Australian Capital Territory (ACT), New South Wales (NSW), Queensland (Qld), South Australia (SA), and Western Australia (WA); 1995-2005 for Victoria (Vic); 1998-2005 for the Northern Territory (NT); and 2003-2005 for Tasmania (Tas). For all states, the dataset covers house sales that had been registered by the end of 2006, which allows for up to a 12-month lag in the official registration.[4]

The key price variable is the mean of the log house prices in a postcode (also known as the log of the geometric mean). This measure is preferable to the arithmetic mean, which is sensitive to changes in the prices of the most expensive houses. It is also preferable to the median house price, which is unaffected by changes that only impact the tails of the distribution. A simple way to think about the geometric mean is that if the cheapest house in a postcode increases in value by 10 percent, this has approximately the same impact on the geometric mean as if the most expensive house in a postcode increases in value by 10 percent.

The other main measure of the housing market that I use is the log of the number of sales in a postcode in a given year. This excludes postcodes in which no houses were sold, which comprise 4 percent of the postcodes in the sample.

---

observation. Formally, the analysis is based upon postcode×state observations, but for expositional simplicity I refer to these as postcodes.

[4] I drop two postcode-year observations for the Northern Territory, which appear to be dominated by the sale of extremely large cattle stations (postcode 872 in 2002, where a single property sold for $5 million; and postcode 862 in 2004, where two properties sold with a geometric mean of $29 million). The next-highest set of prices are for central Sydney and Melbourne, with much larger numbers of sales per postcode.

For some of the specifications, I exploit the distance to the state boundary. This was calculated using a dataset purchased from FindMap Pty Ltd, which contains the distance between the centroids of all postcodes in Australia. For each postcode, I calculate the shortest distance to a postcode in another state, and assign this as the distance from the state border.

Data on tax rates were obtained from legal archives. Where the tax schedule changes part-way through the year, I pro-rata the two rates. For example, if rates change at the end of April, I assign a tax rate to that year which is 1/3rd the rate prevailing from January to April, and 2/3^{rds} the rate prevailing from May to December.

Figure 1 plots aggregate revenue from turnover taxes as a share of state personal income. In most states, this figure increased from about 0.3 percent of personal income in the early-1990s to about 0.6 percent of personal income in the early-2000s (with the rate peaking at nearly 1 percent of personal income in both NSW and Qld in 2003).

Figure 1: Land turnover tax revenue as a share of personal income

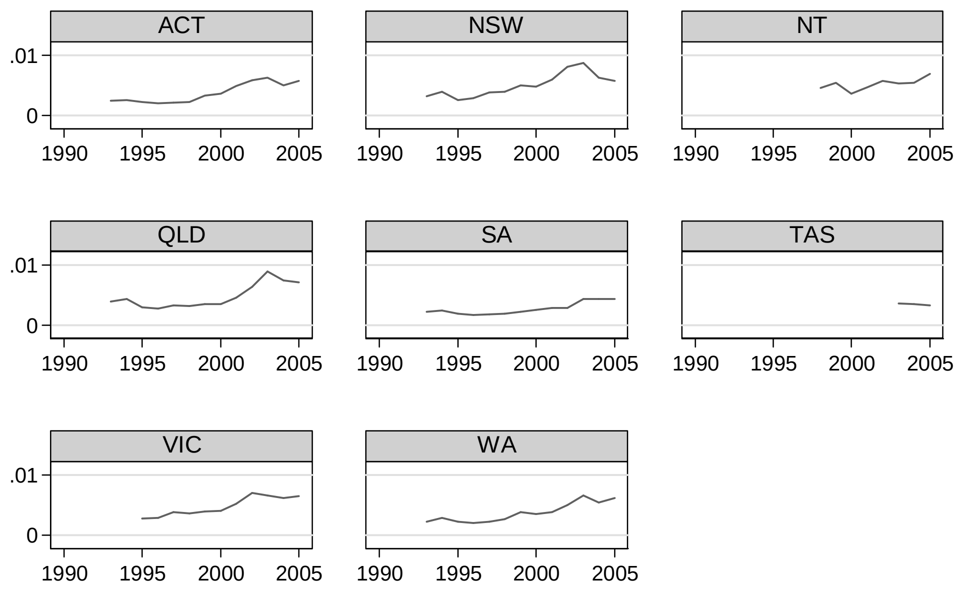

A key empirical challenge in estimating the relationship between taxes and prices is that there is a mechanical relationship between the stamp duty paid on a property and the sale price. Therefore, if one were to simply regress the sale price on the tax payable on that property, the coefficient would capture both the *mechanical* fact that the tax amount is a function of the price, as well as any *behavioural* impact of taxes on prices. (Similar issues arise in estimating the impact of income taxes on wages: see e.g. Feldstein and Wrobel 1998; Leigh 2008.)

To address this problem, I form an instrumental variable that is the predicted tax bill in a given postcode, assuming that property prices in that postcode followed the national trend. In a regression with postcode and year fixed effects, the instrument variable is effectively identified from policy changes and the non-linear nature of the stamp duty schedule.

To provide some intuition for this approach, Figure 2 plots data for the nine postcodes with the highest turnover rates in the year 2000. The solid line shows actual house prices (the geometric mean), while the dashed line shows predicted house prices, assuming that prices had followed the same national trend. By construction, the two series are equal in 1993, the first year of the data. Note that the dashed line has the same slope across all nine postcodes, since it reflects the rate at which the average national price increased over the period 1993-2005. For eight of the nine postcodes, the actual prices rose more rapidly than the national average (the exception being Queensland postcode 4350). This is hardly surprising, given that high-turnover postcodes are potentially also ‘hot’ neighbourhoods, which therefore experienced house price growth that outpaced the national average.

Figure 2: Actual and predicted house prices for the 9 zipcodes with the highest turnover

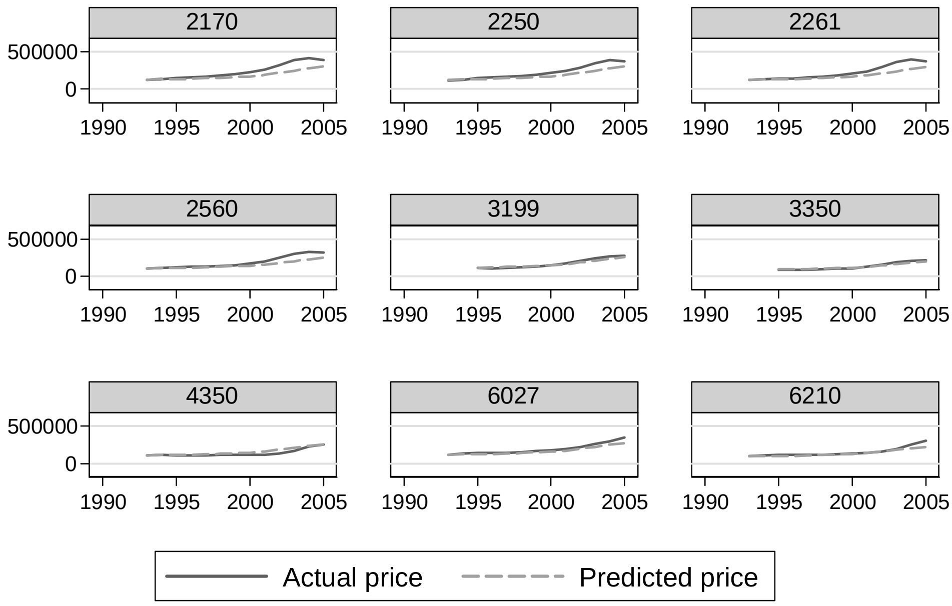

Using the actual and predicted geometric mean sales prices for these postcodes, I then estimate the actual and predicted stamp duty rates. These are plotted in Figure 3. Given that the predicted prices are close to the predicted tax rates, it is not surprising that the predicted tax rates are close to the actual tax rates. Across the nine postcodes, the tax rate rises from an average of about 2 percent to an average of just under 4 percent.

Figure 3: Actual and predicted tax rates for the 9 zipcodes with the highest turnover

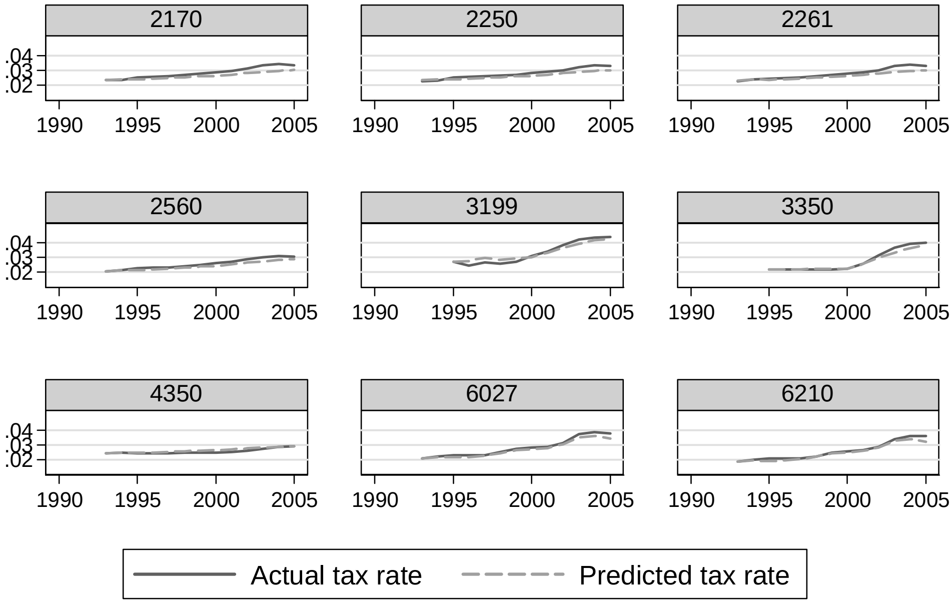

In Figure 4, I calculate the weighted mean tax rate for each state and territory (using both actual and predicted prices). The two series track each other quite closely. All states increase their average tax rates over the period for which data are available, with the largest increases being in the ACT, Vic and WA.

Figure 4: Actual and predicted land turnover taxes by state

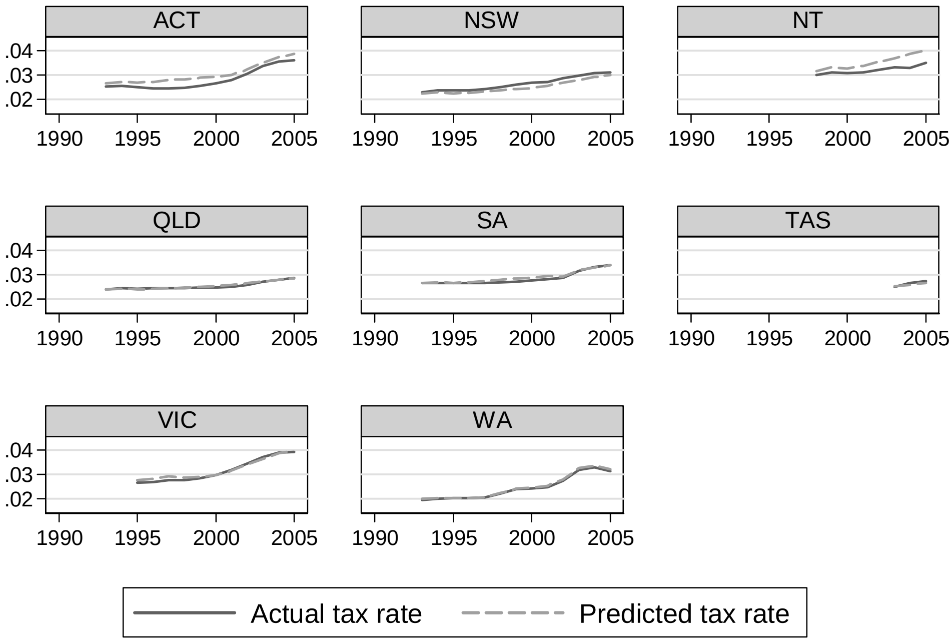

Formally, using data on geometric mean sale prices and turnover in postcode *i* in year *t*, I calculate two stamp duty amounts. The first is τ, which is the actual tax bill based on the geometric mean sale price. The second amount is T, which is the predicted tax bill, based on the assumption that house prices rose at the same pace as the national trend.

In the first stage, I regress the actual tax bill on the predicted tax bill, with postcode and year fixed effects. In the second stage, I use the fitted values to test the impact of tax changes on *Y*, which is either the log of the geometric mean house price, or the log of the number of houses sold. φ and β are parameters.

$$\ln(\tau)_{it} = \varphi \ln(T)_{it} + I_i^{Zipcodes} + I_t^{Years} + \mu_{it} \quad (1)$$

$$\ln(Y)_{it} = \beta \ln(\hat{\tau})_{it} + I_i^{Zipcodes} + I_t^{Years} + \varepsilon_{it} \quad (2)$$

Standard errors are clustered at the postcode level, to account for possible serial correlation within postcodes over time (Bertrand, Duflo and Mullainathan 2004).[5] In specifications where the dependent variable is the log of the house price, observations are weighted by the number of sales.[6] Where the dependent variable is the log of the number of sales, the regressions are unweighted.

I also carry out two robustness checks.[7] One is to estimate the impact of taxes in the region close to the state border. This allows for the possibility that the behavioural effect of taxes might be larger for individuals who can readily purchase a house in

[5] Results are estimated using Stata's *xtivreg2* command (Schaffer 2007), which allows clustering, and does not require weights to be constant within panels.
[6] The coefficient is approximately twice as large in unweighted specifications (-0.29 in the IV specification using all postcodes, and -0.23 in a reduced-form specification using all postcodes).
[7] A further robustness check is to add state×year fixed effects into both equations (1) and (2). Since tax policies only vary at the state-year level, the results in this specification are identified only from non-linearities in the tax schedule (it also allows for the possibility that other time-varying state policies are correlated with changes in tax rates). The results from this robustness check are generally similar to those in the primary specification.

another jurisdiction. Another is to explore the impact of lagged tax rates, which accounts for the possibility that the housing market may take some time to adjust to a change in tax rates.

3. Results

Table 1 shows the relationship between stamp duty and house prices. In the first column, I present results using an IV specification (instrumenting $\ln(\tau)$ with $\ln(T)$), while in column 2, I estimate a reduced-form regression (using $\ln(T)$ directly). In the IV specification, the first stage is very strong, with an F-statistic on the excluded instrument of 222, and a partial R-squared of 0.09. The elasticity of house prices with respect to stamp duty is -0.16, suggesting that a 10 percent rise in stamp duty leads to a 2 percent fall in house prices.

In the reduced-form specification, the coefficient is only slightly smaller (-0.13). The difference between the IV and reduced-form specifications is a measure of the size of the coefficient on $\ln(T)$ in the first-stage of the IV regression. If that coefficient is 1, then the reduced form and IV specifications will produce the same elasticity. If the coefficient on the excluded instrument in the first-stage IV regression is greater than 1, then it will act to 'scale down' the IV elasticity, relative to the reduced-form specification. Conversely, if the coefficient on the excluded instrument in the first-stage IV regression is smaller than 1, then it will act to 'scale up' the IV elasticity, relative to the reduced-form specification. In this case, the coefficient on $\ln(T)$ in the first-stage of the IV regression is 0.85 (SE=0.05), which effectively 'scales up' the reduced-form coefficient from -0.13 to -0.16. The more closely that the values of the instrumented stamp duty and the actual stamp duty track one another, the closer the IV and reduced-form specifications will be.

In column 3, I re-estimate the regressions, but this time restricting the sample to postcodes less than 50 kilometres (31 miles) from the nearest state border. The notion underlying this cut-off is that such a distance represents a plausible commuting zone, potentially allowing individuals to move to a different state without changing their job. I find that the elasticity of house prices with respect to stamp duty rates is substantially higher in these regions. In this sub-sample, the elasticity is around -0.81, implying that a 10 percent rise in stamp duty lowers house prices by 8 percent.

However, it is possible that the IV results are higher towards the state border because the first-stage is less precisely estimated. Measurement error in either ln(τ) or ln(T) will affect the coefficient on the excluded instrument in the first-stage regression. Indeed, in the full sample, the coefficient on the predicted stamp duty rate in the first stage is 0.85 (SE=0.05); while in the bordering postcodes sample, the coefficient on the predicted stamp duty rate in the first stage is 0.49 (SE=0.13).

Nonetheless, the larger elasticity in bordering postcodes does not appear to be solely due to differences in the first-stage regression. Even in a reduced-form specification, the elasticity in the bordering postcodes sample (column 4) is 0.39: three times as large as the elasticity in the reduced-form specification with all postcodes (column 2).

Table 1: Stamp Duty and House Prices
Dependent variable is the mean log house price

	[1] Full sample (IV)	[2] Full sample (reduced form)	[3] <50km from border (IV)	[4] <50km from border (reduced form)
Log (stamp duty)	-0.155**	-0.133***	-0.808*	-0.393***
	[0.062]	[0.044]	[0.415]	[0.099]
Postcode fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	27645	27645	3696	3696
Postcodes	2508	2508	332	332
R-squared	0.83	0.88	0.6	0.91

Note: ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively. Standard errors, clustered at the postcode level, in parentheses. In columns 1 and 3, the log of the actual stamp duty (ln(τ)), is instrumented using the stamp duty on an average property in that postcode, assuming prices rose with the national trend (ln(T)). In columns 2 and 4, I directly use ln(T) as the stamp duty measure. All specifications are weighted by the number of house sales in that postcode-year cell.

How do changes in stamp duty affect the number of houses sold in a postcode? In Table 2, I estimate models using the log of the number of sales as the dependent variable. In the full sample, instrumenting the actual stamp duty with the predicted stamp duty (column 1), I find an elasticity of -0.21, which implies that a 10 percent increase in stamp duty lowers housing turnover by 2 percent. In the reduced-form specification (column 2), the elasticity is only -0.12. Both coefficients are only significant at the 10 percent level.

There is not clear evidence on whether turnover is more responsive to stamp duty rates as one comes closer to the state border. In the IV specification, the elasticity is larger (compare column 3 with column 1), but in the reduced-form specification, the elasticity is smaller (compare column 4 with column 2). The coefficients in columns 3 and 4 of Table 2 are not statistically significant.

Table 2: Stamp Duty and House Sales
Dependent variable is the log of the number of house sales

	[1] Full sample (IV)	[2] Full sample (reduced form)	[3] <50km from border (IV)	[4] <50km from border (reduced form)
Log (stamp duty)	-0.208*	-0.120*	-0.269	-0.074
	[0.122]	[0.069]	[0.731]	[0.195]
Postcode fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	27645	27645	3696	3696
Postcodes	2508	2508	332	332
R-squared	0.1	0.11	0.1	0.13

Note: ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively. Standard errors, clustered at the postcode level, in parentheses. In columns 1 and 3, the log of the actual stamp duty (ln(τ)), is instrumented using the stamp duty on an average property in that postcode, assuming prices rose with the national trend (ln(T)). In columns 2 and 4, I directly use ln(T) as the stamp duty measure.

Until this point, I have assumed that the effect of stamp duty on house values and turnover occurs in the same year. But it is possible that the effects may take time to manifest themselves. This could occur due to information lags (if buyers and sellers do not immediately realise that stamp duty rates have risen), or because housing transactions take a long time, and parties could be reluctant to adapt their bargaining strategies when the tax rate changes.

In Table 3, I re-estimate the house price models, but with additional lags (one additional year in Panel A, two additional years in Panel B). For each regression, I show the coefficients for each year, plus the sum of the three coefficients, which denotes the impact on prices of a stamp duty rise that persists over 2 years (Panel A) or over 3 years (Panel B).

In the full sample, the summed coefficients are slightly larger than the corresponding estimates from Table 1. The elasticity over a 2 year period (columns 1 and 2 of Panel A) is -0.19 in the IV specification and -0.21 in the reduced form specification. Over a 3 year period, the elasticity in the full sample (columns 1 and 2 of Panel B) is -0.18 in the IV specification and -0.31 in the reduced form specification.

The specification that only includes postcodes near state borders is presented in columns 3 and 4 of Table 3. In the IV specification (column 3), neither elasticity is significant. In the reduced form specification (column 4), the 2-year elasticity is -0.52, while the 3-year elasticity is -0.63. These are both larger than the 1-year elasticity, which is -0.39 (Table 1, column 4). Again, these results suggest that the impact of stamp duty on prices is slightly larger when sustained over a 2 or 3 year period.

Table 3: Stamp Duty and House Prices (Medium Run)
Dependent variable is the mean log house price

	[1] Full sample (IV)	[2] Full sample (reduced form)	[3] <50km from border (IV)	[4] <50km from border (reduced form)
	Panel A: Two-year impacts			
Log (stamp duty)$_t$	0.18	-0.112	2.073	-0.133
	[0.341]	[0.075]	[3.205]	[0.192]
Log (stamp duty)$_{t-1}$	-0.366	-0.093	-2.62	-0.385
	[0.283]	[0.068]	[2.640]	[0.237]
Postcode fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	24296	24296	3224	3224
Postcodes	2411	2411	318	318
R-squared	0.844	0.892	0.224	0.913
Sum of stamp duty coefficients	-0.186**	-0.205***	-0.548	-0.517***
	[0.088]	[0.042]	[0.78]	[0.108]
	Panel B: Three-year impacts			
Log (stamp duty)$_t$	0.414	-0.177**	3.499	0.041
	[0.257]	[0.069]	[2.669]	[0.163]
Log (stamp duty)$_{t-1}$	-0.353	0.103*	-4.303	-0.699***
	[0.235]	[0.062]	[3.573]	[0.195]
Log (stamp duty)$_{t-2}$	-0.243***	-0.239***	0.406	0.028
	[0.059]	[0.042]	[1.716]	[0.217]
Postcode fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	21365	21365	2827	2827
Postcodes	2246	2246	306	306
R-squared	0.872	0.905	-0.746	0.921
Sum of stamp duty coefficients	-0.182**	-0.313***	-0.398	-0.631***
	[0.078]	[0.045]	[0.937]	[0.141]

Note: ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively. Standard errors, clustered at the postcode level, in parentheses. In columns 1 and 3, the log of the actual stamp duty ($\ln(\tau)_t$, $\ln(\tau)_{t-1}$, and $\ln(\tau)_{t-2}$, where applicable), are instrumented using the stamp duty on an average property in that postcode, assuming prices rose with the national trend ($\ln(T)_t$, $\ln(T)_{t-1}$, and $\ln(T)_{t-2}$, where applicable). In columns 2 and 4, I directly use $\ln(T)_t$, $\ln(T)_{t-1}$, and $\ln(T)_{t-2}$, where applicable, as the stamp duty measures. All specifications are weighted by the number of house sales in that postcode-year cell.

In Table 4, I look at the effect of lagged stamp duty on housing turnover. In the full sample, the summed coefficients are larger than in the corresponding 1 year specification. Instrumenting the stamp duty rate (column 1), the elasticity of turnover with respect to the tax rate rises to -0.33 over 2 years and -0.47 over 3 years. In a reduced form specification (column 2), the elasticity of turnover rises to -0.23 over 2 years and -0.41 over 3 years. For the subsample of postcodes that are close to state borders (columns 3 and 4), the turnover estimates have large standard errors, and are never statistically significant.

Table 4: Stamp Duty and House Sales (Medium Run)
Dependent variable is the log of the number of house sales

	[1] Full sample (IV)	[2] Full sample (reduced form)	[3] <50km from border (IV)	[4] <50km from border (reduced form)
Panel A: Two-year impacts				
Log (stamp duty)$_t$	-0.008	-0.11	-8.308	-0.474
	[1.206]	[0.119]	[69.560]	[0.423]
Log (stamp duty)$_{t-1}$	-0.323	-0.117	12.505	0.531
	[1.162]	[0.122]	[101.216]	[0.458]
Postcode fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	24296	24296	3224	3224
Postcodes	2411	2411	318	318
R-squared	0.122	0.123	-62.869	0.124
Sum of stamp duty coefficients	-0.331***	-0.227***	4.197	0.056
	[0.115]	[0.072]	[31.819]	[0.224]
Panel B: Three-year impacts				
Log (stamp duty)$_t$	-0.38	-0.333***	3.732	-0.702*
	[0.758]	[0.117]	[57.869]	[0.422]
Log (stamp duty)$_{t-1}$	0.684	0.352**	-2.126	1.023
	[1.251]	[0.167]	[20.872]	[0.626]
Log (stamp duty)$_{t-2}$	-0.775	-0.414***	-1.124	-0.218
	[0.658]	[0.137]	[42.108]	[0.455]
Postcode fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	21365	21365	2827	2827
Postcodes	2246	2246	306	306
R-squared	-0.128	0.149	-5.185	0.136
Sum of stamp duty coefficients	-0.471***	-0.394***	0.482	0.104
	[0.141]	[0.082]	[5.15]	[0.269]

Note: ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively. Standard errors, clustered at the postcode level, in parentheses. In columns 1 and 3, the log of the actual stamp duty ($\ln(\tau)_t$, $\ln(\tau)_{t-1}$, and $\ln(\tau)_{t-2}$, where applicable), are instrumented using the stamp duty on an average property in that postcode, assuming prices rose with the national trend ($\ln(T)_t$, $\ln(T)_{t-1}$, and $\ln(T)_{t-2}$, where applicable). In columns 2 and 4, I directly use $\ln(T)_t$, $\ln(T)_{t-1}$, and $\ln(T)_{t-2}$, where applicable, as the stamp duty measures.

4. Conclusion

Using exogenous variation in stamp duty rates, this paper has estimated the impact of changes in stamp duty on house prices and housing turnover. I find statistically significant and economically meaningful impacts of changes in stamp duty on both outcomes. Across all neighbourhoods, the short-term impact of a 10 percent increase in the tax rate is to lower house prices by 1-2 percent. This impact rises over time: for example, a stamp duty increase that is sustained over 3 years lowers house prices by 2-3 percent. It is also larger for homes located near state borders, where a 10 percent increase in the tax rate lowers house prices by 4-8 percent.

Since stamp duty averages only 2-4 percent of the value of the property, these results imply that the economic incidence of the tax is entirely on the seller; that is, prices fall by the full amount of the tax. Indeed, the house price results are in some sense "too large", in that they imply a larger reduction in sale prices than the value of the tax. Assuming that stamp duty amounts to 4 percent of the house value, an elasticity of 0 would suggest that the buyer alone bore the tax; an elasticity between 0 and -0.04 would suggest that the stamp duty was shared between the buyer and seller; and an elasticity of -0.04 would indicate that the seller bore the tax (and thus the net-of-tax sale price would be unaffected by a rise in the stamp duty rate). Elasticities below -0.04 indicate that the seller bears more than the tax (ie. that in dollar terms, the house price drops by more than the size of the increase in the tax bill). However, it is worth noting that for most of the specifications (such as column 1 of Table 1), the 95 percent confidence intervals for the elasticity of house prices with respect to stamp duty include -0.04.

I also observe impacts of stamp duty on housing turnover. In the full sample, a 10 percent increase in stamp duty lowers turnover by 1-2 percent in the first year. However, over a 3-year period, a 10 percent stamp duty increase lowers housing turnover by 4-5 percent. Close to state borders, the effects of stamp duty on housing turnover is imprecisely estimated in most specifications.

Taken together, these results imply that stamp duty can have an economically meaningful impact on housing prices and turnover in Australia. Averaging across the five jurisdictions for which I have house price data in all years, the average stamp duty rate rose from 2.45 percent in 1993 to 3.25 percent in 2005. In percentage terms, this represents a 33 percent increase in stamp duty over this period – relative to what would have occurred if the rate had remained constant. These results suggest that if stamp duty rates had remained at their 1993 levels, Australia would have more housing turnover and higher house prices than it does today.

Naturally, such a conclusion does not automatically imply that stamp duties are welfare-reducing. In a review paper on the pros and cons of housing transaction taxes, van Ommeren (2008) discusses a number of arguments against such taxes. For workers, high transaction taxes may cause individuals to forego better job offers in other regions (thereby reducing productivity), or to commute overly long distances to a new job (thereby increasing road congestion). Housing transaction taxes may also discourage unemployed homeowners from taking a job in another region. In addition, housing transaction taxes may lead to mismatches in the housing stock, by effectively discouraging young families to upsize their housing and by discouraging retiree households from downsizing.

Conversely, housing transaction taxes may have advantages. If land taxes are an efficient way of raising revenue, but are politically unfeasible, then land transfer taxes may be an appropriate second-best policy. In addition, if residential turnover reduces the social capital in a neighbourhood (see eg. Dietz and Haurin 2003; Frijters and Leigh 2008), then stamp duties may internalise the negative externality that movers impose on their neighbourhood. It is also possible that lower stamp duties might create the opportunity for real estate brokerage fees to rise (for example, the US has low stamp duties but high brokerage fees: see Belot and Ederveen 2005). Ensuring that brokerage fees do not rise when stamp duty is reduced should be a priority in any such policy reform.

References

Belot, M. and Ederveen, S. 2005. "Cultural and institutional barriers in migration between OECD countries" CPB Netherlands Bureau for Economic Policy Working Paper.

Bertrand, M., Duflo, E. and Mullainathan, S. 2004. "How Much Should We Trust Differences-in-Differences Estimates?"; Quarterly Journal of Economics, 119(1), 249-75

Dietz, R.D. and Haurin, D.R.. 2003. 'The Social and Private Micro-Level Consequences of Homeownership", *Journal of Urban Economics*, 54: 401-450.

Feldstein, M. and Wrobel, M.V. 1998. "Can State Taxes Redistribute Income?" *Journal of Public Economics* 68(3): 369–96.

Frijters, P. and Leigh, A. 2008. "Materialism on the March: From Conspicuous Leisure to Conspicuous Consumption?" *Journal of Socio-Economics*, 37(5): 1937-1945

Leigh, A. 2008. 'Do Redistributive State Taxes Reduce Inequality?' *National Tax Journal*, 61(1): 81-104

Schaffer, M.E., 2007. xtivreg2: Stata module to perform extended IV/2SLS, GMM and AC/HAC, LIML and k-class regression for panel data models. http://ideas.repec.org/c/boc/bocode/s456501.html

van Ommeren, J. 2008. "Transaction Costs in Housing Markets", Tinbergen Institute Discussion Paper TI2008-099/3. Tinbergen Institute, Amsterdam.