Household-specific housing capital gains and consumption: Evidence from Canadian microdata

Kul Bhatia [a],[*], Chris Mitchell [b],[1]

[a] Department of Economics, University of Western Ontario, Social Science Centre Room 4071, London, ON N6A 5C2, Canada
[b] Financial Stability Department, Bank of Canada, 234 Laurier Ave. W., Ottawa, ON K1A 0G9, Canada

ABSTRACT

This paper analyzes how household consumption decisions respond to accrued capital gains on owner-occupied housing using Canadian microdata. It makes an important methodological contribution by utilizing a hedonic regression approach to estimate household-specific capital gains instead of relying on house-price indices and other proxy measures that have been widely used in earlier literature. The results suggest that household consumption expenditures are increasing in the level of accrued capital gains on housing, and that the sustainability of these gains is important for the magnitude of this relationship. When the level of accrued capital gains is persistent over time, total household consumption increases by approximately 5.4 cents for every dollar of permanent capital gains and non-durable consumption increases by approximately 3.9 cents. Estimates of marginal propensity to consume for households in different age categories and other subgroups are also presented.

1. Introduction

How movements in house prices affect household spending decisions has been a topic of extensive research for a long time. Real estate is the largest single asset-type in household balance sheets in many countries, so a lot of empirical work has been done with national and state-level data from the US, UK, Spain, Australia, and elsewhere. This study presents comparable results for Canada. It is the first analysis of this type to use microdata from the Canadian Survey of Household Spending (SHS), and to the broader literature in this area, it offers a technique for estimating household-specific housing capital gains from information about actual real estate transactions and housing characteristics, instead of relying on price indices and other proxy variables commonly used in earlier research. This enables us to capture house-price movements within a region as well as across a broad geographic area, and this can be done for every home-owning household in the SHS, as well as for many types of homes and dwellings. Armed with such detailed estimates of housing value, it is then postulated that a key decision-making variable is "expected" or "permanent"

gains rather than current-period gains, and statistical results are derived for different specifications of the expected gains variable. A lot of research with microeconomic data has been devoted to the behavior of various subgroups and the credit channels that might be used (Engelhardt, 1996; Bover, 2005; Juster et al., 2005; Aron and Muellbauer, 2006; among others). In addition to estimating marginal propensity to consume (MPC) for the entire population, we also shed light on how different subgroups — young and old, high- and low-income brackets, renters and homeowners — respond to changes in real estate prices.

Our approach is motivated by the observation that housing markets tend to be local or regional in character, house prices do not move in unison everywhere, and there are significant differences in how these markets operate across countries and even across jurisdictions within a given country. It is not surprising, therefore, to find that estimates of housing wealth effects – typically, consumer responses to changing house prices or MPC – range from zero (Elliott, 1980) to 0.09 (Benjamin et al., 2004), and several points in between (0.03, Dvornak and Kohler, 2007; about 0.04, Case et al., 2013; and 0.05 to 0.08, Calomiris et al., 2012). This research by and large relies on aggregate data, annual or quarterly, at the national level or for individual states. When microeconomic data is used, differences in wealth distribution, wealth composition, income, age, and other demographic variables are incorporated into the analysis, and they yield a wide range of MPC

* Corresponding author. Tel.: +1 519 661 2111x85323.
E-mail addresses: kbhatia@uwo.ca (K. Bhatia), mith@bankofcanada.ca (C. Mitchell).
[1] Tel.: +1 613 782 7973.

estimates — anywhere from 0.14 (Engelhardt, 1996), to 0.02 (Bover, 2005), to zero (Juster et al., 2005).

A range of specifications, statistical techniques, and data (time series, ad hoc and sequential surveys, and panel data) are used in these studies, so it is difficult to pin down the sources of discrepancy among the conflicting estimates, much less try to reconcile them. It is important to recognize, however, that an important element in this literature is the many different ways in which capital gains and losses, home values, and movements in home prices are estimated and incorporated into the analysis. The list includes one or more of: perpetual inventory aggregations of construction cost (Elliott, 1980), also employed in US Flow-of-Funds and Census of Housing compilations), homeowners' own estimates of real estate value (the US Panel Study of Income Dynamics (PSID) for instance), and a variety of price indices (Campbell and Cocco, 2007). A principal objective of this paper is to suggest that differences along these lines could well be a significant reason for the wide ranging and often contradictory results in this important area of research. The focus here is on the small but growing literature that uses microeconomic data to study housing and includes topics such as consumption-saving decisions, asset allocation, and housing choices for different age groups.

For estimating wealth effects, real estate prices enter this framework in many ways: First of all is the home-price variable itself, which could be the sale price of a respondent's home, the average price of homes actually sold in a given locality, or an appropriate index of real estate prices. Reports by local real estate boards about market conditions might also fit the bill sometimes, especially if cohort or other grouped data are used. Second, but likely the most important for a consumer's decision making, is the household's own perception of the value of its home, however derived, for that determines if the household feels more or less wealthy after a price change, and by how much. First differences in such estimates will measure accrued housing gains, to be distinguished from a change in a price index and realized gains (the latter requires a house to be sold, involving transaction costs and possibly a tax liability). This is the most commonly reported house-value variable in large data sets such as the US-PSID. Researchers can then apply concepts such as "expected," "permanent," and "predictable" price movements, or their opposites – "unexpected," "transitory," and "unpredicted" price changes – to analyze consumer decisions. Third, credit constraints figure prominently in this literature, and it has been argued that the observed correlation between consumption and home-price increases might be due to a loosening of credit constraints rather than a direct wealth effect (Buiter, 2010). Since home equity is a widely-used source of collateral for many household loans, how a household's borrowing capacity changes in the wake of a home-price movement is not easy to predict. Among other things, it might depend on the availability of other credit sources, restrictions on mortgage lending, and other supply-side considerations beyond the household's control, one of which could be the price variable a potential lender chooses to consider. Jiang et al. (2013) point out that a well-established technique in mortgage finance research is to rely on local price indices. For estimating wealth effects, they therefore combine the two points made above and use price indices for Metropolitan Statistical Areas (MSAs) to adjust self-reported home values in their data set. Still, they are unable to measure household-specific house price changes, but that can be done with the methodology being used in this study, applied to the SHS data.

The SHS contains comprehensive data on consumer expenditures, income, and household composition, much like the UK Family Expenditure Survey (FES) or some Michigan Research Centre surveys in the US. It also has price data for homes bought in a given year as well as information on characteristics of all housing units from which household-specific estimates of accrued capital gains on housing are computed (more on that in a moment). Our results suggest an MPC estimate of 0.054 for all consumption expenditures and 0.039 if durables purchases are excluded (pooled sample, 2004–2006). There are also new MPC

estimates for homeowners in different age groups, and also for renters. There is quite a range of results for these subgroups in the literature.

The mainstay of this research is a rather innovative application of the hedonic technique which deconstructs the actual purchase price of a house to permit estimation of the implicit prices for individual housing characteristics. By repeating this exercise for successive years, we get a measure of accrued capital gains (CG) that, arguably, reflects price movements more accurately than other methods commonly used in this literature. It is a measure of pure price change, and with microdata it takes into account the specific housing characteristics of each household. It also facilitates computing "expected" or "permanent" capital gains as in earlier analyses with aggregate time-series data. We thus have information for each household in the sample, not just for a group or a cohort of them, which has many advantages. For instance, households with different levels of estimated wealth effects can be precisely identified, and we can consider other household-specific issues. To cite one result, the MPC estimate out of expected gains for older homeowners is 0.06, significantly higher than .031 for younger homeowners. By contrast, Campbell and Cocco (2007, p.616) report that the effect of "predictable house prices" on the consumption of these two groups is not statistically different from each other. Jiang et al. (2013) arrive at yet another MPC estimate for older households (0.03). The model specification, the data, and the estimation techniques in these studies are quite different from ours, so the estimates are not strictly comparable. We mention them to underscore the point that household-level market-based estimates of house price movements are worth taking seriously.

The remainder of the paper is organized as follows: Section 2 discusses other relevant literature and highlights some theoretical and empirical issues. Section 3 describes the data and provides summary statistics on Canadian housing wealth. Section 4 deals with the estimation of accrued capital gains on owner-occupied housing. Everything depends on the quality of these estimates, so their statistical properties are carefully examined, and these are also compared with the only other alternative — the Teranet index. Section 5 describes the consumption functions, the estimation results are presented in Section 6, and the conclusions are summarized in Section 7.

2. Other literature

The econometric literature based on aggregate data has gone through several stages — straightforward regressions of consumption on income, wealth, and other variables; Permanent Income Hypothesis (PIH)-type specifications; error correction models; models incorporating household characteristics and wealth distributions, etc. Housing wealth effects by and large have been found to be positive, and stronger results have emerged when market prices of real estate have been incorporated into the estimation process. An early example, which provided some motivation for the present research, is Elliott (1980). Relying on construction cost data in the US Flow-of-Funds accounts, it found no significant wealth effect, whereas Bhatia (1987), in a similar PIH-type specification, but using numbers that incorporated market prices, found a statistically significant result (Calomiris et al. (2012) provide a useful summary of the more recent literature). From microdata also, a comparison of two studies, Skinner (1993) and Hoynes and McFadden (1994), made by Engelhardt (1996) is quite illuminating, for they report marginal propensity to save estimates with opposite signs from essentially the same US-PSID data. One point of contrast between the two is that Skinner uses self-reported values whereas Hoynes and McFadden rely on MSA level data on market price "… obtained from quarterly surveys of homebuilders, appraisers, mortgage bankers, and savings and loan officers."

The advent of large bodies of survey data during the last few decades has facilitated estimation of household-level consumption functions, for individual households or cohorts of them (Bostic et al. (2009) survey the recent literature). With the SHS data we can incorporate