How Important are Land Values in House Price Growth? Evidence from Canadian Cities

Kenneth G. Stewart

Associate Professor
Department of Economics
University of Victoria
Victoria, British Columbia
Canada V8W 2Y2

email: kstewart@uvic.ca
phone: (250) 721-8534
fax: (250) 721-6214

April 2021

Abstract

A Cobb-Douglas growth accounting framework is used to study the contributions of structures and land to newly-constructed home prices across major Canadian cities. The data set is unusual in that land prices are directly observed rather than having to be imputed, and quality change is carefully controlled for in the measurement of structures. These data permit testing of constant returns to scale, which in conventional applications must be adopted as a maintained hypothesis, as well as the introduction of dynamic effects. Whereas standard analyses find land costs to be the dominant contributor to the growth in housing costs, I find that this varies greatly by city. Yet, despite this novel empirical result, the evidence supports other recent work that endorses the constant-returns Cobb-Douglas methodological framework.

Keywords: Land; Land value; Land prices; House prices; Residential construction costs

JEL classifications: R31; R32

Declarations of interest: none

Disclosure

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

1 Introduction

A major issue in contemporary housing economics is the relative contributions of land and construction costs to the escalation of house prices that has taken place in recent decades. Several prominent articles have used national-level data to conclude that, as Davis and Heathcote (2007, p. 2618) find for the United States, "...house price fluctuations are primarily driven by changes in the price of land ..." More recently Knoll, Schularick, and Steger (2017; henceforth KSS) construct time series for fourteen countries and conclude (p. 349) that "...the late twentieth century surge in house prices was due to sharply rising land prices. About 80 percent of the increase in real house prices in advanced economies in the second half of the twentieth century can be explained by higher land values." Nor is this finding limited to national-level data: Davis, Oliner, Pinto, and Bokka (2017) use property-level data for Washington DC to estimate the price of land and conclude that changes in house prices are dominated by land values.

A recognized deficiency of this literature is a lack of good data on land prices: "Unfortunately, market data on land values have been notoriously piecemeal and subject to numerous measurement challenges." (Albouy, Ehrlich, and Shin, 2018, p. 454.) Instead land values must typically be inferred as a residual from data on the sale prices of houses and construction costs. As Davis and Heathcote (2007, p. 2596) remark, "...we do not directly measure the price of land, but rather infer it from data on house prices and structures costs. With the exception of land sales at the undeveloped fringes of metro areas—where land is relatively cheap—there are very few direct observations of land prices from vacant lot sales, because most desirable residential locations have already been built on." Davis, Oliner, Pinto, and Bokka (2017) similarly impute land costs for their property-level data.

In Davis and Heathcote, KSS, and much of the related literature, these residual land costs are inferred as a literal "Solow residual"-type quantity using growth accounting. And, just as with the Solow residual of macroeconomic growth accounting, this imputed residual is to some extent a measure of our ignorance: it attributes any difference between sales prices and construction costs to land values. Yet, conceivably, structure construction may not fully measure all non-land inputs, so that some non-land costs are left unobserved.

This paper examines the legitimacy of this methodological framework using a time series data set for newly-constructed houses in Canadian cities, a data set that appears to be unique among those available from national statistical agencies and is previously unexploited for this purpose. Because the data are for new house construction, land costs are observed. A second advantage is that changes in housing quality over time are carefully controlled for.

These data yield contributions that are both methodological and economic; in some respects the findings are confirmatory, in others contradictory. With respect to methodology, I argue that the evidence supports a Cobb-Douglas framework generally and the assumption of constant returns to scale specifically. Support for the former is consistent with Ahlfeldt and McMillen (2014) and Baum-Snow and Han (2019). Support for constant-returns Cobb-Douglas is consistent with the theory of Epple, Gordon, and Sieg (2010) and corroborates the recent empirical work of Combes, Duranton, and Gobillon (2017) that uses French data. Furthermore, I go beyond the existing literature and investigate the sensitivity of this modelling framework to seasonality and dynamics. It stands up well to these generalizations.

Yet, despite this support for the received methodology, there is an important contradictory finding with respect to the economics of house price growth. In contrast to the conventional conclusion, city-level data indicate that—at least for new house construction—there is considerable heterogeneity across cities in the role of land costs. In some cities the growth in house prices is indeed mainly, or even entirely, attributable to land costs; but in others, construction costs dominate.

2 Modeling Framework

For the purpose of studying these issues I follow the usual practice (e.g. KSS; Davis and Heathcote 2005; Hornstein 2009; Diewert, Nishimura, Shimizu, and Watanabe 2020, Chap. 4, sec. 4.2.2) of using a Cobb-Douglas technology to describe the production of housing. The quantity of housing Q is produced from land L and structure construction S according to, for general returns-to-scale possibilities,

$$Q = \gamma L^{\alpha} S^{\beta}. \tag{1}$$

Minimization of costs $C = p^{\mathrm{L}}L + p^{\mathrm{S}}S$ subject to this production technology yields the Cobb-Douglas cost function

$$C = \psi(p^{\mathrm{L}})^{\alpha/\nu}(p^{\mathrm{S}})^{\beta/\nu}Q^{1/\nu}, \tag{2}$$

where $\nu \equiv \alpha + \beta$ denotes the degree of returns to scale, and $\psi \equiv \nu(\gamma\alpha^{\alpha}\beta^{\beta})^{-1/\nu}$. As usual with Cobb-Douglas production, the exponents α and β have interpretations both as production elasticities and, under competitive markets, as factor shares. We therefore expect their numerical values to be positive fractions.

Under constant returns to scale (CRS), $\nu = 1$, average costs are

$$\frac{C}{Q} = \psi(p^{\mathrm{L}})^{\alpha}(p^{\mathrm{S}})^{\beta}. \tag{3}$$

With the average cost of houses C/Q at time t measured with a house price index p_t^{H}, and land and structure costs measured by indexes p_t^{L} and p_t^{S}, in principle this cost function can be estimated loglinearly as

$$\log p_t^{\mathrm{H}} = \log\psi + \alpha\log p_t^{\mathrm{L}} + \beta\log p_t^{\mathrm{S}}. \tag{4}$$

Notice that, although constant returns $\nu = \alpha + \beta = 1$ is assumed in obtaining the average cost function (3), it is still a testable restriction on this loglinear regression. A rejection of that restriction calls into question the legitimacy of the step from (2) to (3) and, in fact, the modelling structure more generally.

In this framework, growth in total house prices can be decomposed into components attributable to price growth in each of land and structures:

$$\Delta\log p_t^{\mathrm{H}} = \alpha\Delta\log p_t^{\mathrm{L}} + \beta\Delta\log p_t^{\mathrm{S}}. \tag{5}$$

The shares of each factor in total house price growth are therefore:

$$\text{Growth shares:} \qquad \alpha\frac{\Delta\log p_t^{\mathrm{L}}}{\Delta\log p_t^{\mathrm{H}}}, \qquad \beta\frac{\Delta\log p_t^{\mathrm{S}}}{\Delta\log p_t^{\mathrm{H}}}. \tag{6}$$

Equivalently, the growth relationship (5) can be averaged over the sample,

$$\frac{1}{n}\sum_t \Delta\log p_t^{\mathrm{H}} = \alpha\frac{1}{n}\sum_t \Delta\log p_t^{\mathrm{L}} + \beta\frac{1}{n}\sum_t \Delta\log p_t^{\mathrm{S}},$$

and the growth shares (6) stated in terms of these sample averages. That is, the growth shares are naturally evaluated using the sample mean growth rates.[1] In terms of terminology, it is important to distinguish between these *growth shares* and the *factor shares* α and β.

To take this framework to the data, the typical circumstance—described in the introduction—is that, whereas house prices p_t^{H} and construction costs p_t^{S} are observable, the land price p_t^{L} is not. The simplest approach to handling this, used by KSS for example, is to rearrange the growth relationship (5) to solve for land:[2]

$$\Delta \log p_t^{\mathrm{L}} = \frac{1}{\alpha}\Delta \log p_t^{\mathrm{H}} - \frac{\beta}{\alpha}\Delta \log p_t^{\mathrm{S}}.$$

An assumed or data-based value for α, and the assumption of CRS, so $\beta = 1 - \alpha$, therefore yields an imputed series for the growth of the land price, which can in turn be used to obtain the decomposition (6). Notice that, in this approach, the adoption of CRS as a maintained hypothesis is essential. Just as most housing data do not provide a land price, so too do they not provide information on land's factor share α.

For this purpose Davis and Heathcote use a constructed series for the factor share of land, while KSS use an assumed value. In contrast, with data on the land price p_t^{L} in addition to p_t^{H} and p_t^{S}, it becomes possible to estimate the loglinear cost function (4) and use the estimated factor shares to calculate the growth shares (6). An appealing by-product is that CRS becomes testable, shedding light on the legitimacy of the conventional approach that must assume it.

3 Data and Descriptive Evidence

Given this importance of a data set that provides series for all three of p_t^{H}, p_t^{L}, and p_t^{S} obtained from a unified and coherent methodological framework, let us consider Statistics Canada's New House Price Index (NHPI). To quote directly from the Statistics Canada webpage:[3]

> The New Housing Price Index (NHPI) is a monthly series that measures changes over time in the contractors' selling prices of new residential houses, where detailed specifications pertaining to each house remain the same between two consecutive months. The survey covers the following dwelling types: single homes, semi-detached homes and townhouses. The survey also collects contractors' estimates of the current value (evaluated at market price) of the land. These estimates are independently indexed to provide the published series for land. The current value of the structure is also independently indexed and is presented as the house series.

Notice in particular that the land values are separately collected and independently indexed. Hence, to the extent that CRS may be found to be supported by the data, it is *not* as an artifact of land prices having been imputed from the other series using a Cobb-Douglas model with CRS. (That is, we are not simply finding something in the data because it has been introduced by construction.)

Significantly for our purposes, the NHPI does *not* include two categories of housing: custom homes and condominium apartments. With respect to the former (from the same webpage), "The NHPI measures price change for house models that can be priced repeatedly and whose detailed specifications remain the same over time. Consequently, the observed population for the NHPI excludes custom built houses." The implication is that the NHPI is probably dominated by tract housing on the urban fringe. Land values in the urban core are such that, when a building lot becomes available for development or redevelopment, the value of the property is maximized with a home that is custom designed for that lot. This optimizes what can be created within local zoning restrictions, sight lines, and the existing development of neighbouring properties. To be sure, occasionally larger areas of land become available within the urban core to which these considerations may not apply to the same extent, and so may be captured within the NHPI. But the amount of housing involved is probably fairly modest relative to the newly-developed suburbs of city peripheries. This interpretation of the NHPI is consistent with some of the descriptive evidence, as we are about to discuss.

Another important feature of the NHPI is that quality changes are carefully controlled for. From the same webpage:

> In addition to pricing data and the reasons for price change, detailed information is collected which describes the characteristics of each selected house model (e.g. living area, lot specifications, etc.). This information is used to determine the dollar value of changes in the characteristics or "quality" of the model for which prices have been collected. The reported model prices are adjusted for quality changes so that the NHPI will measure "pure" price change over time.

For some cities the NHPI begins as early as January 1981. In order to study a common sample from major cities in every region of the country, I use the sample May 1984–May 2020.[4] This includes the most important period of escalation in average prices in the past century, the "hockey stick" described by KSS (p. 335): "The long-run trajectory of global house prices displays a hockey-stick pattern: real house prices remained broadly stable from the late nineteenth century to World War II. They trended upward in the post-war decades and have seen a particularly steep incline since the late 1980s."

The NHPI indexes for total house prices, land, and structures are published as nominal values. I convert these to real values using the city-specific consumer price indexes which, unlike alternatives such as the GDP deflator, are available both monthly and by city. For the ten cities I study, these deflated total house, land, and structure price indexes are shown in Figure 1. Table 1 shows the associated average growth rates.

A few preliminary regularities emerge from an inspection of Figure 1. Perhaps the most pervasive is the extent to which total house prices and construction costs tend to track one another, with land being less strongly related. Growth accounting is the tool that will allow us to clarify the links.

A second, but less pervasive, regularity is the speculative boom in real estate prices in the run-up to the 2008–9 financial crisis after a period of fairly stagnant prices in the 1990s. This is most dramatic for the prairie cities of Saskatoon, Edmonton, and Calgary, which experienced house price escalations as dramatic as that of late-1980s Toronto. Often these booms are most strongly manifested in land prices. Quebec City and Winnipeg are striking examples, where land values increased much more than did house prices.

At least as interesting as the boom is the variation in what followed. In some cities (Saskatoon, Edmonton, Calgary) there was a subsequent retrenchment in total house prices (bust is probably too strong a word), in others little (Quebec City, Montreal, Ottawa) or not at all (Toronto, Winnipeg). Why different cities responded differently to the housing bubble is a matter for speculation; growth accounting will not prove enlightening in this regard. However the volatility of land prices relative to construction costs or total house prices that is evident for so many cities is consistent with the findings of Davis, Oliner, Pinto, and Bokka (2017). They studied the Washington, DC metropolitan area over the boom-bust years 2000–2013 and found that (p. 239) "...land prices were more volatile than house prices everywhere, but especially so in the areas with initially inexpensive land." The latter would include the urban periphery that I have argued probably dominates the NHPI.

Table 1 shows how these plots translate into average growth rates over the 36 years. The typical pattern for most cities is real total price growth that ranges between around 0.5 percent (Ottawa and Saskatoon) and 1.6 percent (Calgary). The behaviour of land and structure costs that underlies these price growths varies. In Ottawa, for example, the real land costs of new construction have been essentially constant. We can therefore anticipate that the growth decomposition (6) will find that Ottawa's price growth is entirely due to constructions costs.

The most atypical growth patterns happen to be at opposite ends of the country. In Halifax real house prices and both their land and structure components have seen growth that is just slightly positive, perhaps reflecting the steady modest growth in the regional economy.

To anyone even casually familiar with the Canadian housing market, the biggest surprise of Figure 1 and Table 1 is Vancouver, where nominal house prices are well known to have escalated dramatically in recent decades. The key to understanding the secular declines of Figure 1 and the associated negative average growth of Table 1 is to remember that these data are for newly-built non-custom homes in real terms, quality adjusted. This being the case, recall once again the conjecture that the NHPI is dominated by construction on the urban periphery. To the extent that one accepts this, the Vancouver series suggest that, as the city has expanded, land, structure, and total house prices on the periphery have actually declined modestly in real quality-adjusted terms. Similar considerations would explain Ottawa's decline in real land costs. For greater Vancouver residents willing to move to the periphery and incur the city's notoriously high (at least by Canadian standards) commuting costs, there is a modest reward in the cost of housing: commuting costs are the "price" of lower housing costs. (Or, for those who do not commute, the reward is for not contributing to congestion in the core.)[5]

4 Analysis

Turning from these descriptive remarks to inferential analysis based on the Cobb-Douglas modeling framework of Section 2, the results from OLS estimation of the loglinear cost function (4) are shown in Table 2.

4.1 Static OLS Estimates

Recall that, from the Cobb-Douglas production function (1) that is dual to the cost function, the parameters α and β have interpretations both as factor shares and as the elasticities of output (housing) with respect to the factor inputs land and structures, respectively. In terms of the cost function (3), they are the elasticities of average housing costs with respect to the prices of land and structures.

4.1.1 Factor Shares

The first two columns of Table 2 show that the estimates of these elasticities/factor shares are fairly consistent across the cities. The land elasticity falls in the range 0.20–0.35, while the structure elasticity is in the range 0.65–0.82; the standard errors show them to be well-estimated. City-by-city, these estimates sum to being very close to unity. Perhaps surprisingly, then, F-tests typically reject CRS (the p-values are given in Table 3): the fairly large sample size relative to the small number of parameters being estimated means that the estimates are fairly precise, and so restrictions like CRS tend to be rejected. (As we will discuss in connection with Table 4, the rejections are less pervasive when temporal dependence is treated.) Nevertheless, that these elasticities come so close to summing to one, and consistently so across the cities, tends to be supportive of the conventional methodology that employs CRS as a maintained hypothesis. These findings are fairly consistent with other analyses. For example, in their assessment of the literature Baum-Snow and Han (2019, p. 3) conclude that "…housing production function estimates … indicate [an] approximately Cobb-Douglas form with a land share of 0.2–0.3." In the same spirit, Ahlfeldt and McMillen (2016, p, 27) "…conclude that the housing production function is likely closer to the convenient Cobb-Douglas form than long believed in the literature."

To compare with KSS, recall the interpretation of α and β as factor shares. The analysis of KSS requires an assumed value of α, land's factor share. They cite descriptive evidence in support of a benchmark value of $\alpha = 0.5$, and explore the sensitivity of their results to varying this in the range 0.25–0.75. However they were analyzing the total housing market, which is dominated by resales of existing housing. On the premise that newly-constructed non-custom homes are predominantly on the urban fringe, they should have a lower land share, which is what the estimates of α in Table 2 show.

For this reason a more direct comparator to my work is Combes, Duranton, and Gobillon (2017), who find (p. 1) "…that the production function for housing is reasonably well, though not perfectly, approximated by a Cobb-Douglas function and close to constant returns …" Their study is notable in using a very different analytical framework and data set: observations by building permit in France during the period 2006–2012. The great similarity is that, like me but unlike KSS and so many other studies, their data are for newly-constructed single-family homes. Consequently (p. 1), "We estimate an elasticity of housing production with respect to non-land inputs of about 0.80 . . . ," remarkably similar to the upper end of my range for β.

4.1.2 Growth Shares

Turning from the factor shares α and β, consider the growth shares given by the decomposition (6). Evaluating these using the mean growth rates of Table 1 yields the final three columns of Table 2. As the values indicate, nothing in the nature of the decomposition (6) forces these shares to sum to one by construction; nevertheless it is reassuring that most do, approximately.

Whereas the factor shares α and β each fall within fairly modest ranges across cities, the growth shares exhibit considerably more variation. At one extreme is Quebec City, where land has been more than twice as important as structures in contributed to total house price growth. This reflects the dramatic escalation in land relative to construction costs between 2005 and 2010 shown for that city in Figure 1. At the other extreme is Ottawa, where real new house price growth has been entirely attributable to construction costs. Figure 1 suggests that, in contrast to Quebec City, this is because real land costs in Ottawa today are little different from what they were 35 years ago. In fact, land's growth share for Ottawa is actually slightly negative (and, according to its standard error, statistically significantly so) because a slight decline in the real price of land has offset the growth in house prices that would have occurred from construction costs alone.

Beyond these remarks, what is it about Quebec City and Ottawa that lead to these differences, given their obvious similarities of geography, climate, and population (not to mention their status as capital cities, of Quebec and Canada respectively)? Here we run up against the limitations of growth accounting which, although it uses inferential results (the factor share estimates) as inputs, is ultimately an *accounting* exercise in that it does not explain *why.* Growth accounting—in all fields of application, not just this one—is useful in pointing out differences in the sources of growth, but by itself provides no explanation for where those differences come from.

Across all cities, Table 2 shows that the growth shares of land range over the full continuum of the extremes of Ottawa and Quebec City. This contrasts starkly with KSS:

> At a country-by-country level we find that the contribution of land prices in explaining house price growth ranges from 73 percent (United Kingdom) to 96 percent (Finland), while the median is 86 percent. The contribution of land

> prices to national house price growth is 77 percent for Denmark, 81 percent for Belgium, the Netherlands, and Sweden, 83 percent for Switzerland, 89 percent for the United States, 90 percent for Australia, 92 percent for Norway, 93 percent for France, and 95 percent for Canada. (KSS p. 348)

But although analyses such KSS and Davis and Heathcote (2007) demonstrate convincingly that growth in *average* prices for the national housing stock is dominated by land values, my results suggest two qualifications. First, within countries there can be significant heterogeneity between cities, and this heterogeneity does not follow obvious or simplistic patterns. For example, it is not the case that land costs are a greater contributor to price growth in larger cities. On the contrary, in Montreal, Toronto, and Vancouver the growth share of land is at the low end of the range.

A second qualification is that it would be inappropriate to necessarily interpret the "hockey stick" of increasing land values in social welfare or equity terms—to conclude, for example, that aspiring homeowners are increasingly and everywhere trapped into paying exorbitant land prices, inordinately harming, say, young families. For buyers have the option of paying marginal rather than average costs, in this case "marginal" quite literally including the margins of urban areas. The extent to which these marginal increments to the housing stock are dominated by land values varies dramatically by location. In some locations (like Quebec City and Halifax) growth in land values is indeed the dominant contributor to new house price growth. But in many others it is construction costs that dominate, in a few instances even strongly so.

5 Extensions

These are the key conclusions that come out of the results of Table 2. But are they robust to variations on the implementation of the methodology? Several variations naturally suggest themselves: first, the imposition of CRS; second, treating seasonality in the monthly data; and, third, allowing for temporal dependence in the time series.

5.1 Constant Returns to Scale

The results of Table 2 just discussed do not impose constant returns to scale. By contrast, in the standard application of the methodology of Section 2 the maintained hypothesis of constant returns is essential to imputing the land growth share. This constant returns assumption is well-founded in the broader housing literature. For example, Epple, Gordon, and Sieg (2010, Section I.A) provide a unified formulation of "...the canonical model of housing production that underlies almost all theoretical models in modern urban and local public economics." Consistent with the framework of Section 2, this baseline model conceives of housing as being produced from two factors, land and non-land inputs. One of the core assumptions (not predictions) of their model is constant returns, which they justify as follows.

> The constant returns to scale assumption is typically motivated by the observation that the housing construction industry is characterized by firms of varied sizes. This fact is consistent with constant returns to scale since average costs are independent of firm size. In principle, it should not be difficult for an efficiently operated construction firm to build similar houses on similar lots for approximately the same costs. There are few fixed factors to which one could justify decreasing returns to scale. Managerial ability might be the only factor that is fixed in the short run, but clearly not in the long run. Similarly it is difficult to find a compelling reason that would justify an increasing returns to scale scenario. There is no empirical evidence that would suggest the presence of strong agglomeration effects or other spillovers that might be used to justify increasing returns to scale. (Epple, Gordon, and Sieg, 2010, p. 908.)

How does this intuition stand up empirically? Table 3 repeats the analysis of Table 2, imposing CRS. Although, as already noted, F-tests find this restriction to be rejected for most cities, the substantive results of the analysis are little affected by its imposition. The factor share of land α is in the range of 0.19 (Ottawa) through 0.38 (Vancouver), roughly the same as when CRS is not imposed (Table 2).

The translation into growth shares shows, as before, great variation across cities in the extent to which land has contributed to the growth in house prices. In fact the rank-order of the ratio of land to structure shares (final column of Table 3) is unchanged from Table 2, with growth in housing costs in Quebec City being dominated by land and in Ottawa by structures.

Finally notice that, as in Table 2, the growth decomposition (6) does not by construction yield values that sum to one identically, even when CRS is imposed; nevertheless many do, approximately.

5.2 Seasonality

It is natural to believe that there is some seasonality to the residential real estate market. Everyone knows that the market is most active in the spring as families prepare to move between school years, less active in the fall and winter. Often the effect on prices of these seasonal demand and supply shifts roughly balances. But sometimes prices follow these volume fluctuations to a degree that is evident even descriptively in the data. For example, following a long plateau of prices, Winnipeg saw price increases in the market season of June-July-August 2007 that are most evident in an almost-discrete jump in land values even more dramatic than in other prairie cities.

Less obvious is whether seasonality is systematically present, not in nominal prices, but in real quality-adjusted new house prices in relation to their land and structure components. The simplest way of getting a sense of this is to introduce monthly dummy variables into the static OLS regressions of Table 2. The results turn out to be remarkable in their consistency: for every city, monthly dummies are individually and jointly insignificant, by a wide margin. (Across the cities, the *lowest* p-value for an F-test of the joint significance of monthly dummies in the loglinear cost function (4) is 0.995. Consequently Table 2 does not devote space to reporting them.)

Of course, seasonality can be manifested in ways other than deterministic shifts in the regression intercept. It can also be manifested stochastically, as one source of temporal dependence. Let us consider this more generally.

5.3 Dynamics

A simple and natural way of examining the sensitivity of these findings to temporal dependence in the data is to regard the loglinear cost function (4) as a long run equilibrium relationship within an error correction model (ECM). The availability of observed data on land prices makes this a more credible exercise than if they were imputed.

In doing so, the standard issues of specifying such models arise, in particular determining appropriate lag lengths. On one hand, enough lags should be included to adequately treat the autocorrelation in the data. But on the other, the model should not be over-parameterized to the extent that estimates have little precision, weakening the power of hypothesis tests to reject restrictions such as CRS.

Given that our principal interest is in gauging the sensitivity of the results to dynamics, as opposed to arriving at any one "best" model specification, it is useful to consider the results yielded by alternative information criteria. The two most widely used model selection criteria are the Akaike information criterion (AIC) and Schwarz's Bayesian information criterion (BIC). It is well known that the BIC has the desirable property of yielding correct lag lengths asymptotically, but in finite samples tends to under-select. The AIC, by contrast, yields longer lag lengths.

These differences are manifested in Tables 4 and 5, which report the models selected by the two criteria. City by city, the AIC-selected models of Table 4 have longer lag lengths than the BIC-selected models of Table 5. In consequence the AIC-selected models yield portmanteau statistics that do not reject the null of absence of residual autocorrelation at conventional significance levels, whereas that hypothesis is often rejected in the BIC-selected models. In this respect the AIC-selected models are probably the more believable.

But despite these differences, the big take-away from Tables 4 and 5 is how little a treatment of dynamics alters the substantive conclusions of the static OLS regressions of Table 2. For any given city, the land and structure factor shares α and β are about the same as in the static regressions; across cities they are estimated to be in the ranges of 0.15–0.33 and 0.57–0.81, respectively. Given the expanded number of coefficients in the dynamic models, the standard errors are larger than in the static regressions and CRS is less frequently rejected; even so, they indicate that the factor shares continue to be well-estimated.

Given these estimates of the factor share coefficients α and β from the equilibrium error of the ECM's, the growth shares are obtained as previously by evaluating the formulas (6) using the sample mean growth rates of Table 1. Here too the substantive economic findings mirror those of the static OLS analysis. City by city the land and structure growth shares of Tables 4 and 5 are similar to the static estimates of Table 2. Remarkably, even the rank-order of the relative growth shares (final columns of Tables 4 and 5) are unchanged from previously, with land costs dominating price growth in Quebec City but being negligible (or even slightly offsetting construction costs) in Ottawa.

6 Conclusions

This paper has used data on Canadian cities to assess the conventional growth accounting methodology for studying the degree to which house price growth is due to land costs. The data are for newly-constructed homes, and so land costs are directly measured rather than having to be imputed; as well, changes in structure quality are carefully controlled for.

Across cities, static OLS estimation of Cobb-Douglas cost functions yields land and structure factor shares in the range of 0.20–0.35 and 0.65–0.82. These are consistent with the broader literature, including the recent work of Baum-Snow and Han (2019) and Combes, Duranton, and Gobillon (2017).

In contrast to the most common findings, however, the growth shares that these imply vary greatly across cities. In some cities growth in land costs is the dominant contributor to house price growth, consistent with the national-level findings of widely-cited research such as Davis and Heathcote (2007) and Knoll, Schularick, and Steger (2017). But in others land costs contribute only modestly, or even offset price increases that would arise from construction costs alone. This suggests that, within the national-level statistics, there is significant heterogeneity across cities—the most important economic contribution of this paper.

These conclusions are robust to several variations on the analysis. First, they are unaffected by the relaxation of constant returns. Second, these data are monthly time series, and monthly dummy variables are found to be individually and jointly insignificant. This suggests that, whereas seasonality may be important for other aspects of housing markets, it is not important for the growth accounting methodology. Finally, the substantive economic findings are remarkably insensitive to the incorporation of dynamics.

Constant returns to scale in housing production is widely assumed throughout urban economics, for theoretical reasons well-articulated by Epple, Gordon, and Sieg (2010). Conventional applications of the growth accounting framework therefore impute land costs by assuming constant returns. Although I find constant returns to be formally rejected, especially when temporal dependence in the data is not treated, it is nevertheless the case that estimated factor shares come remarkably close to summing to unity. In turn, the substantive economic results regarding the growth shares are robust to whether constant returns is imposed.

This endorsement of the constant-returns Cobb-Douglas framework—even in its most basic static form—serves to corroborate the emerging evidence from very different sources such as Ahlfeldt and McMillen (2014) and Combes, Duranton, and Gobillon (2017), and is the most important methodological contribution of this paper.

Notes

[1]This is as opposed to, say, attempting to average the ratios $\Delta \log p_t^{\mathrm{L}} / \Delta \log p_t^{\mathrm{H}}$ over the sample. This ratio is undefined when the denominator is zero, as it sometimes is in these data—house prices occasionally do not change between months.

[2]More complicated alternatives include the explicitly hedonic approach of Diewert and collaborators, in which land and structures each consist of observable characteristics for which hedonic prices can be estimated. See Diewert (2013) and Diewert et al. (2020).

[3]Googling "new housing price index" will take you to

https://www23.statcan.gc.ca/imdb/p2SV.pl?Function=getSurvey&SDDS=2310

[4]This also avoids anomalies in the data associated with the recession of the early 1980s—at the time, and for some parts of the country still, the most severe recession since the great depression. In particular, Vancouver experienced a historically unique (for the region) speculative bubble in real estate that peaked and burst in 1981. Even today, inferences drawn from the Vancouver NHPI are influenced by that event when it is included in the sample.

[5]Incidentally, the decline in the Vancouver real NHPI would not be explained by an increase in the mix of condominium apartment construction relative to other housing types. As noted in the opening paragraphs of this section, the NHPI measures single and semi-detached homes and townhouses; it does not include condominium apartments.

References

Ahlfeldt, G., and D. McMillen (2014) "New estimates of the elasticity of substitution of land for capital," *European Regional Science Association Conference Papers.*

Albouy, D., Ehrlich, G., and M. Shin (2018) "Metropolitan land values," *Review of Economics and Statistics* 100, 454–466.

Baum-Snow, N., and L. Han (2019) "The microgeography of housing supply," Working Paper.

Combes, P.-P., Duranton, G., and L. Gobillon (2017) "The production function for housing: Evidence from France," IEB Working Paper 2017/07.

Davis, M.A., and J. Heathcote (2005) "Housing and the business cycle," *International Economic Review* 46, 751–784.

Davis, M.A., and J. Heathcote (2007) "The price and quantity of residential land in the United States," *Journal of Monetary Economics* 54, 2595–2620.

Davis, M.A., Oliner, S.D., Pinto, E.J., and S. Bokka (2017) Residential land values in the Washington, DC meto area: New insights from big data," *Regional Science and Urban Economics* 66, 224–246.

Diewert, W.E. (2013) "Decomposing an RPPI into land and structures components," in *Handbook on Residential Property Price Indices (RPPIs),* ed. by J. de Haan and W.E. Diewert (Luxemburg: Publications Office of the European Union).

Diewert, W.E., Nishimura, K.G., Shimizu, C., and T. Watanabe (2020) *Property Price Index—Theory and Practice* (Springer).

Epple, D., Gordon, B., and H. Sieg (2010) "A new approach to estimating the production function for housing," *American Economic Review* 100, 905–924.

Hornstein, A. (2009) "Problems for a fundamental theory of house prices," Federal Reserve Bank of Richmond *Economic Quarterly* 95, 1–24.

Knoll, K., Schularick, M., and T. Steger (2017) "No price like home: Global house prices, 1870–2012," *American Economic Review* 107, 331–353.

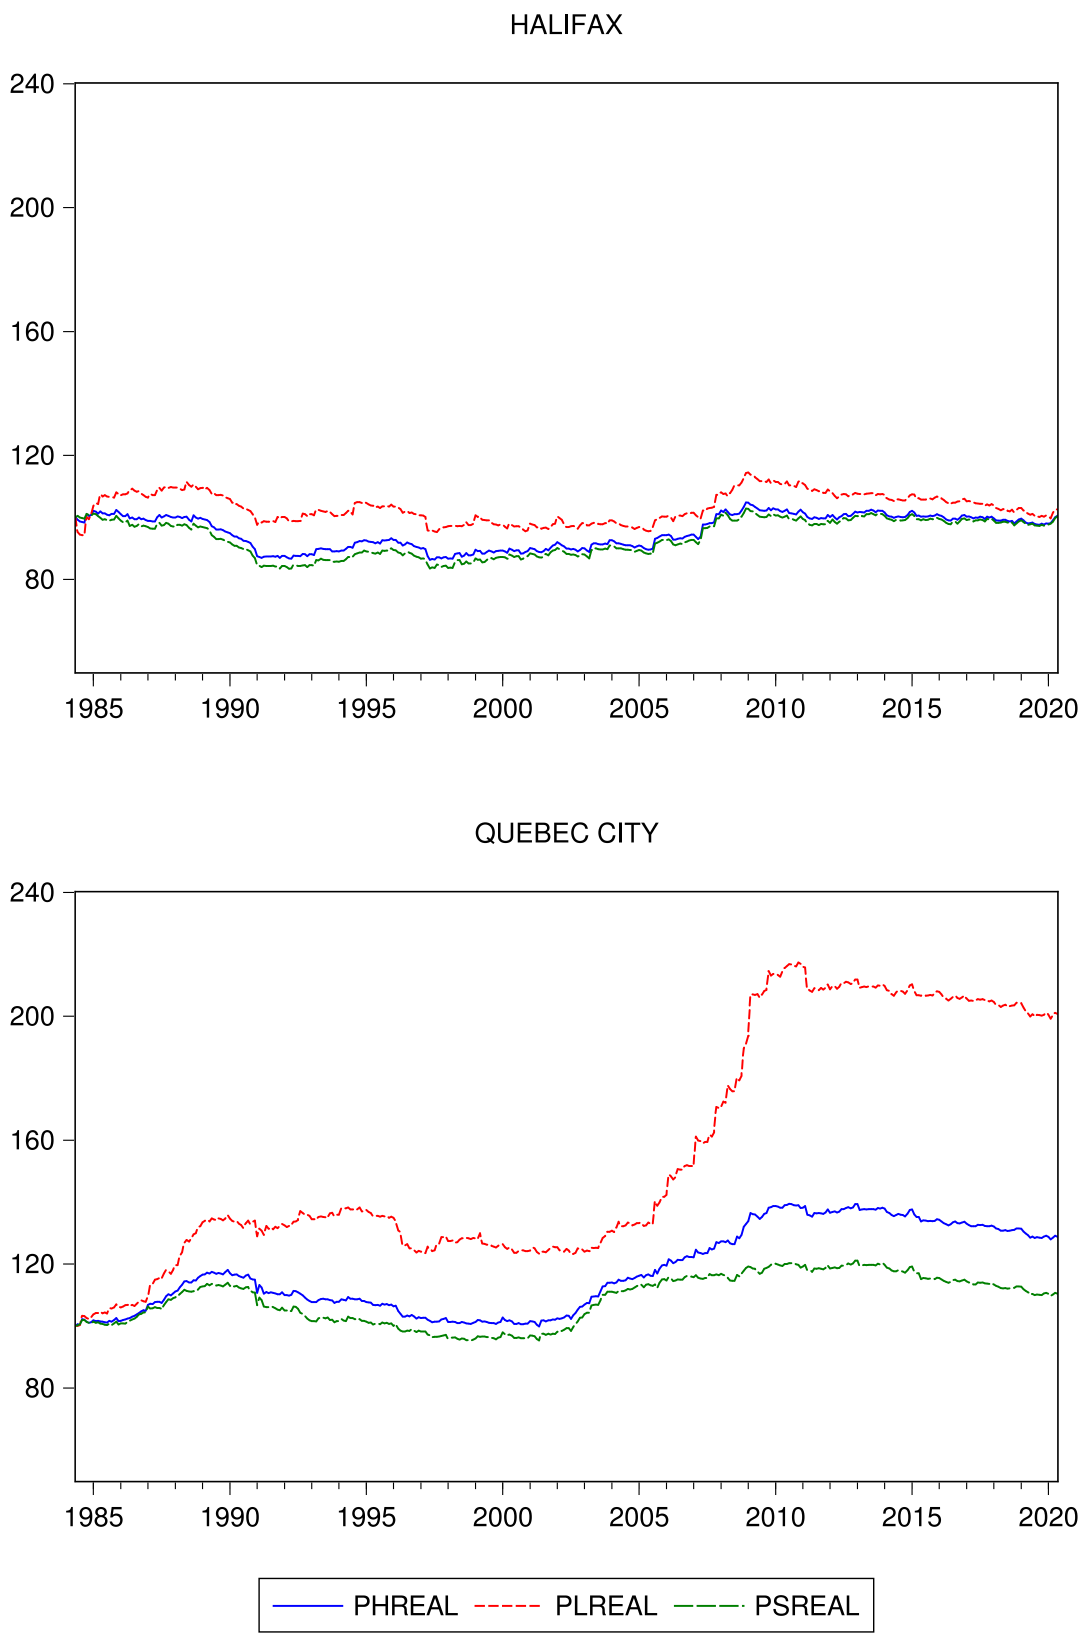

Figure 1(a): Real house, land, and structure prices, May 1984–May 2020 (indexes, May 1984=100)

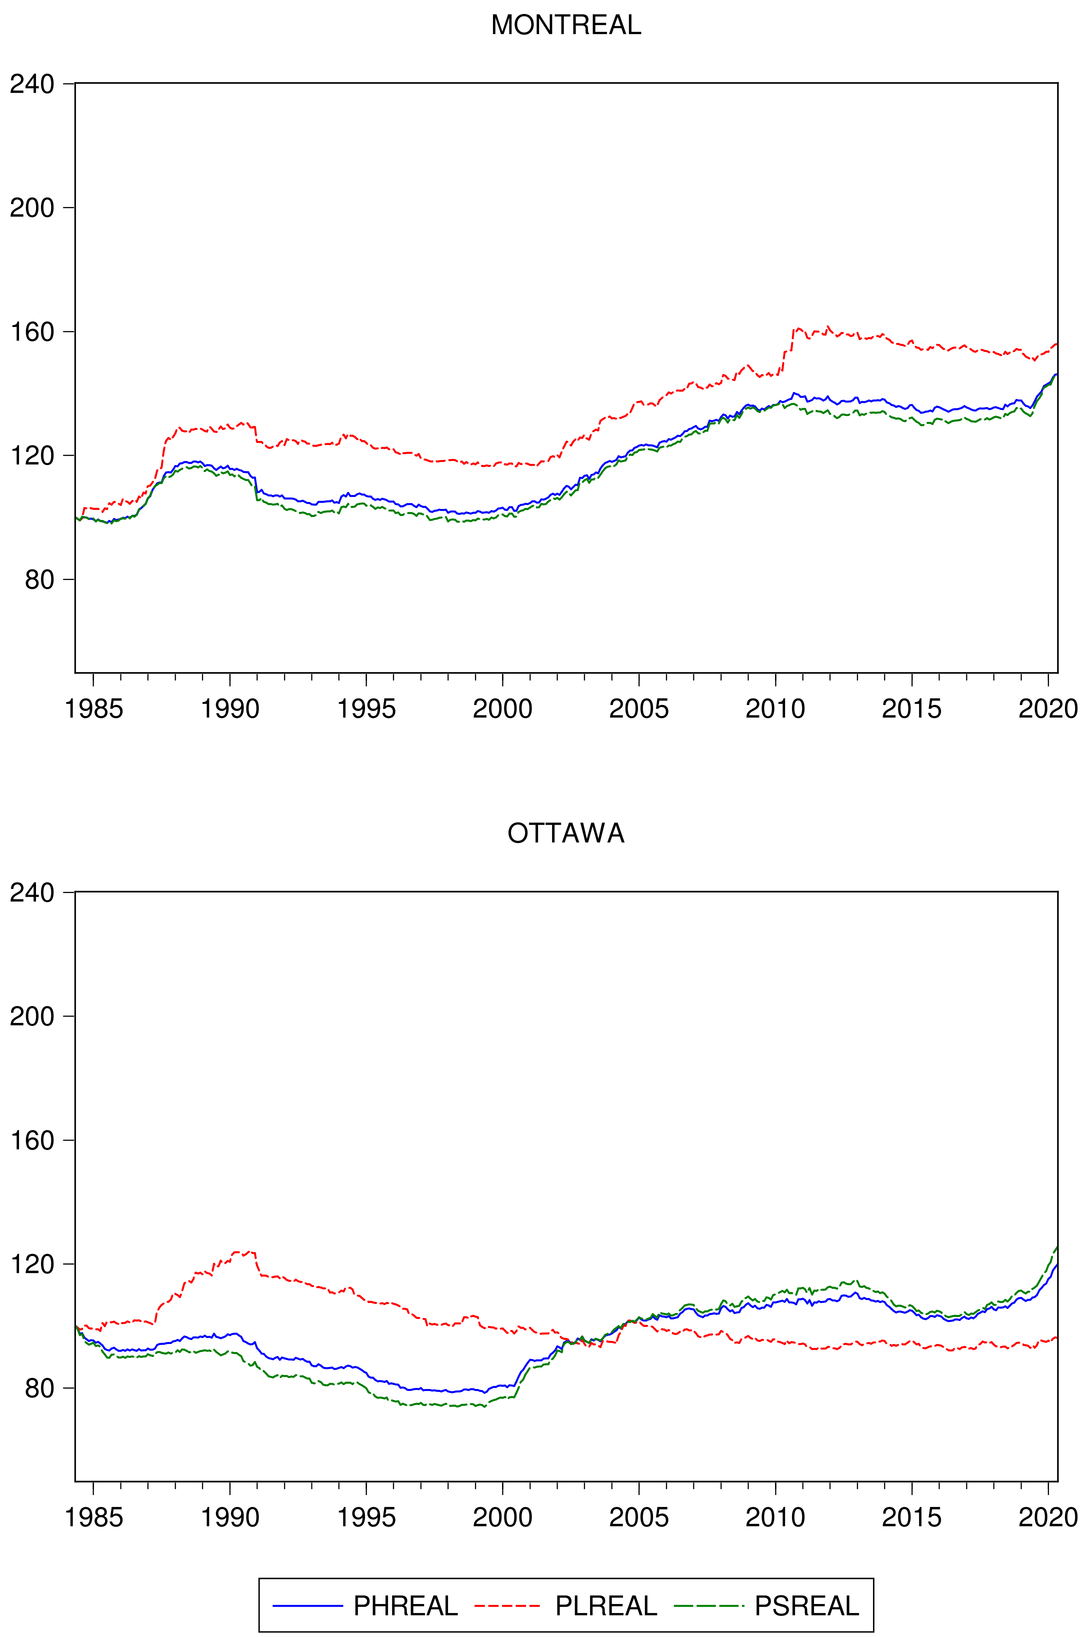

Figure 1(b): Real house, land, and structure prices, May 1984–May 2020 (indexes, May 1984=100)

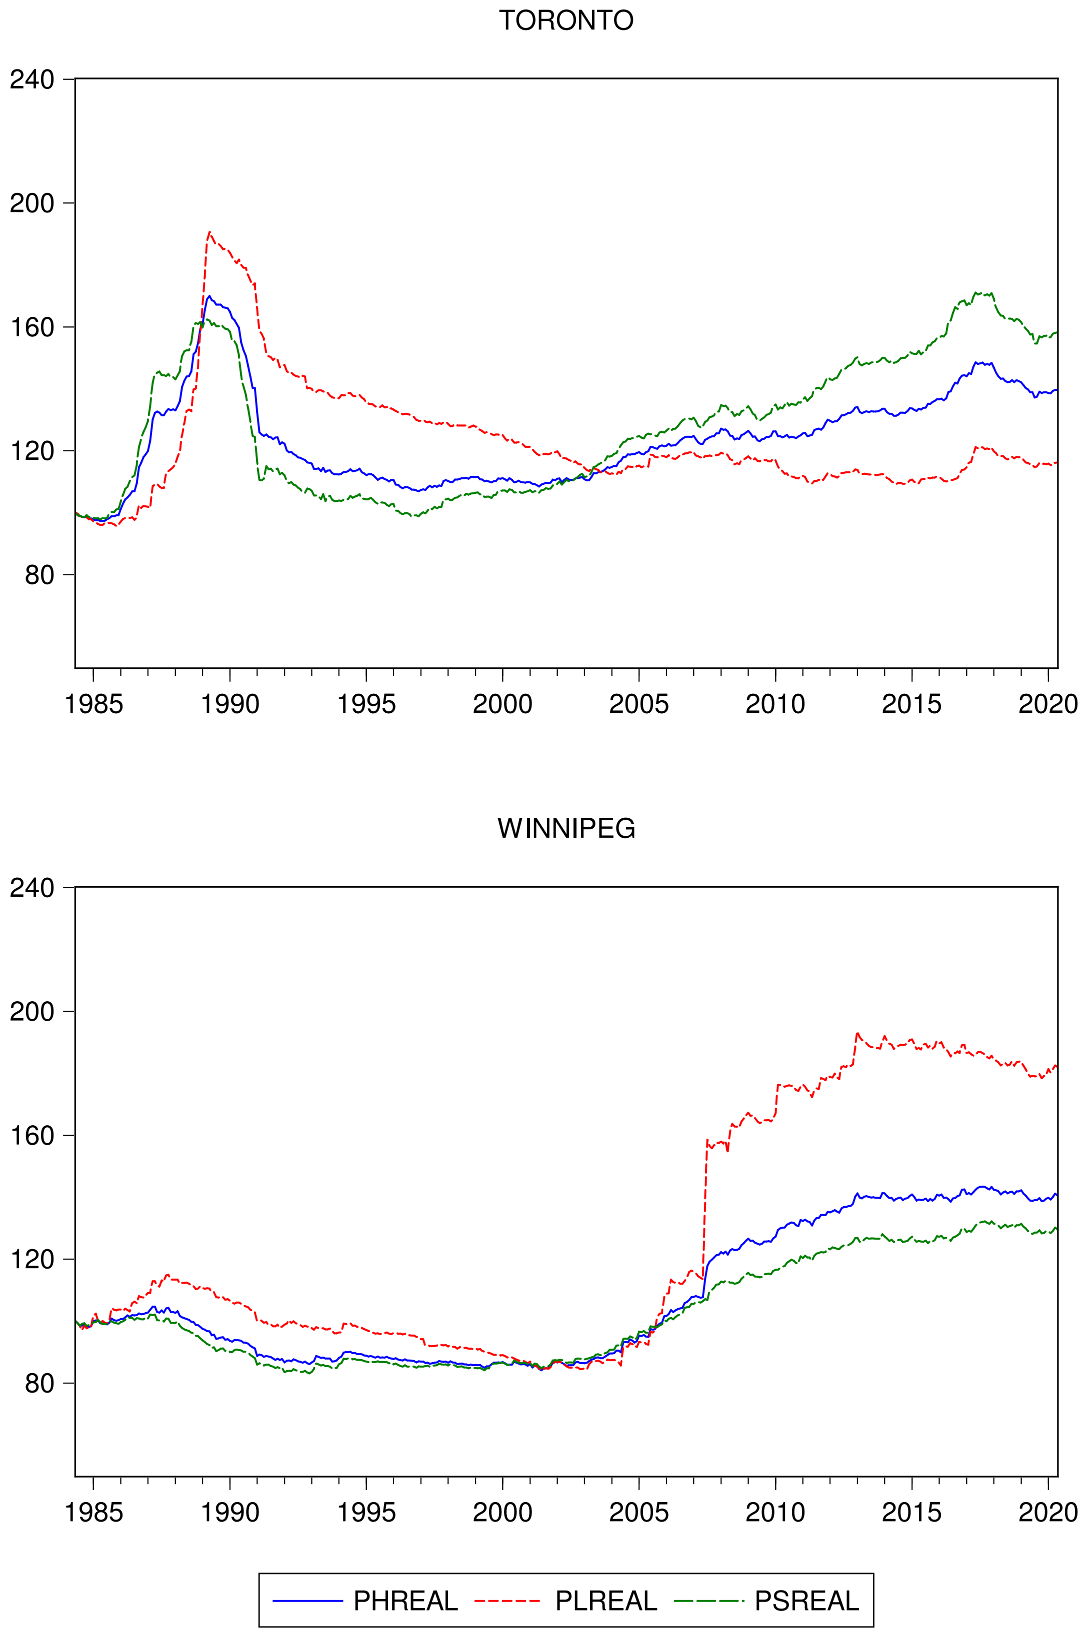

Figure 1(c): Real house, land, and structure prices, May 1984–May 2020 (indexes, May 1984=100)

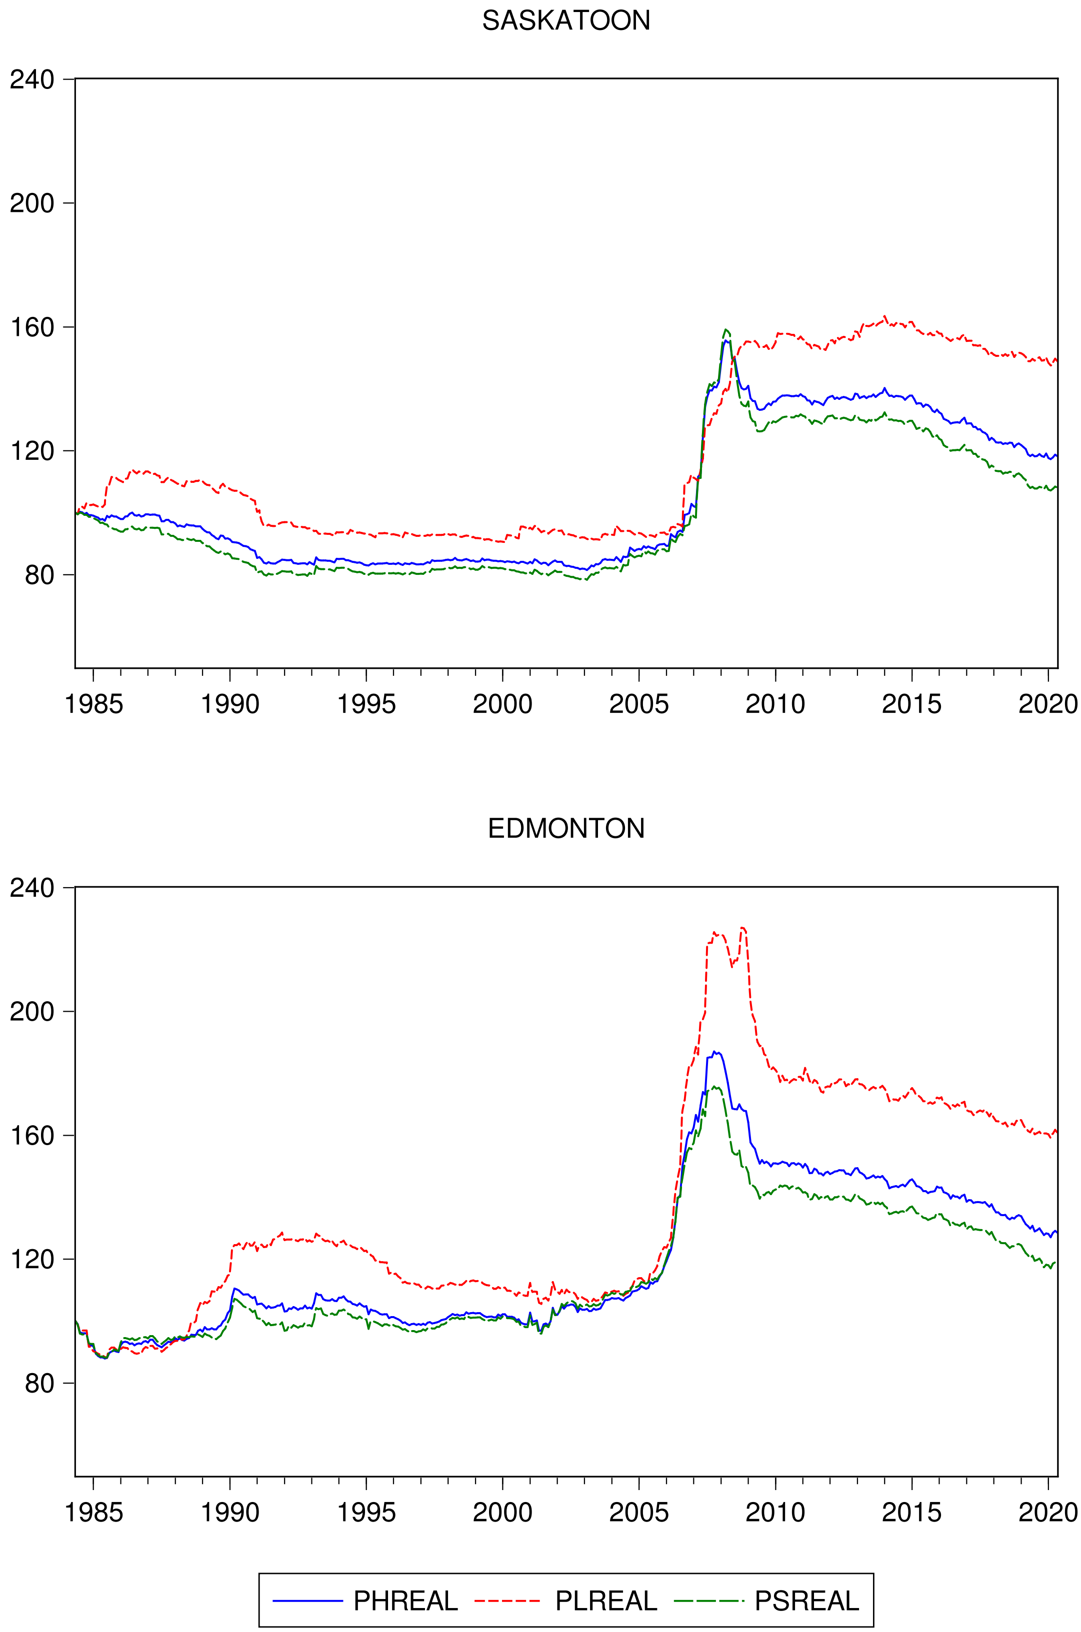

Figure 1(d): Real house, land, and structure prices, May 1984–May 2020 (indexes, May 1984=100)

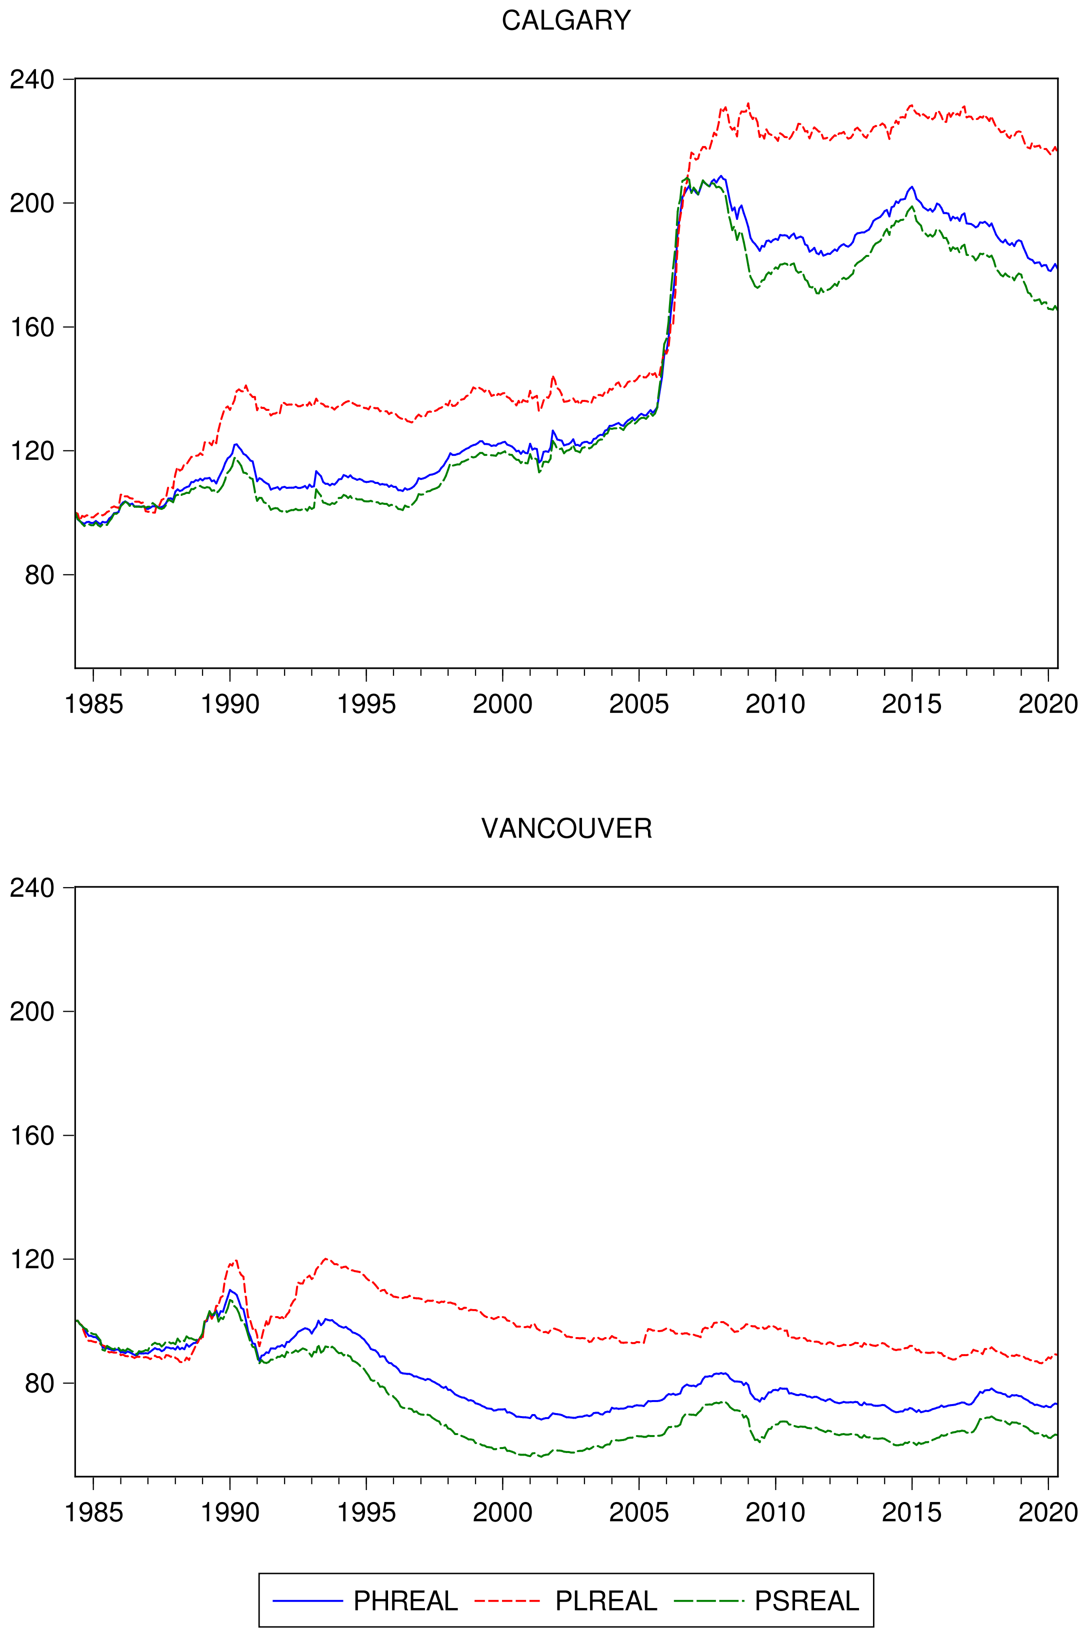

Figure 1(e): Real house, land, and structure prices, May 1984–May 2020 (indexes, May 1984=100)

-

Table 1: Growth in Real Prices, May 1984–May 2020 (annualized percent)

	Total house: $\Delta \ln p^{\mathrm{H}}$		Land: $\Delta \ln p^{\mathrm{L}}$		Structure: $\Delta \ln p^{\mathrm{S}}$	
City	mean	standard deviation	mean	standard deviation	mean	standard deviation
Halifax	0.0170	8.0149	0.0702	9.9607	0.0114	8.2087
Quebec City	0.6955	7.0926	1.9176	11.5719	0.2740	6.8974
Montreal	1.0600	6.5770	1.2493	8.7194	1.0454	6.8542
Ottawa	0.5123	7.2840	−0.0727	8.2396	0.6297	8.2189
Toronto	0.9317	10.5260	0.4110	12.8517	1.2876	12.2057
Winnipeg	0.9524	8.0989	1.6603	18.0181	0.7267	7.1221
Saskatoon	0.4869	13.2234	1.0936	13.4530	0.2333	15.2581
Edmonton	0.6403	13.2413	1.3011	16.8116	0.4056	13.1054
Calgary	1.6103	12.0327	2.1374	13.0962	1.4062	13.2634
Vancouver	−0.8731	10.0798	−0.3290	11.8515	−1.2830	11.6772

Note: $n = 433$ observations.

Table 2: Shares of Land and Structures in House Price Growth: Static OLS

City	Factor shares: α	Factor shares: β	Growth Shares: Land share $\alpha\frac{\Delta \ln p^{\mathrm{L}}}{\Delta \ln p^{\mathrm{H}}}$	Growth Shares: Structure share $\beta\frac{\Delta \ln p^{\mathrm{S}}}{\Delta \ln p^{\mathrm{H}}}$	Growth Shares: Ratio of land to structure shares $\frac{\alpha\Delta \ln p^{\mathrm{L}}}{\beta\Delta \ln p^{\mathrm{S}}}$
Halifax	0.2418	0.7488	0.9978	0.5018	1.9885
	(0.0030)	(0.0022)	(0.0124)	(0.0015)	(0.0294)
Quebec City	0.2651	0.7796	0.7310	0.3072	2.3795
	(0.0018)	(0.0057)	(0.0049)	(0.0022)	(0.0312)
Montreal	0.2203	0.7817	0.2596	0.7709	0.3368
	(0.0027)	(0.0029)	(0.0032)	(0.0028)	(0.0054)
Ottawa	0.1974	0.8154	−0.0280	1.0022	−0.0280
	(0.0021)	(0.0012)	(0.0003)	(0.0014)	(0.0003)
Toronto	0.3453	0.6483	0.1523	0.8960	0.1700
	(0.0023)	(0.0018)	(0.0010)	(0.0025)	(0.0012)
Winnipeg	0.2511	0.7585	0.4378	0.5788	0.7564
	(0.0012)	(0.0022)	(0.0020)	(0.0016)	(0.0056)
Saskatoon	0.2560	0.7532	0.5750	0.3610	1.5930
	(0.0013)	(0.0014)	(0.0029)	(0.0007)	(0.0108)
Edmonton	0.2854	0.7230	0.5800	0.4579	1.2666
	(0.0030)	(0.0043)	(0.0060)	(0.0028)	(0.0205)
Calgary	0.2847	0.7170	0.3779	0.6261	0.6036
	(0.0030)	(0.0033)	(0.0040)	(0.0029)	(0.0091)
Vancouver	0.3056	0.6186	0.1152	0.9090	0.1267
	(0.0046)	(0.0021)	(0.0017)	(0.0031)	(0.0021)

Note: Standard errors are in parentheses.

Table 3: Shares of Land and Structures in House Price Growth: CRS Imposed

City	α	CRS[a] (p-value)	Growth Shares: Land share $\alpha\frac{\Delta \ln p^{\mathrm{L}}}{\Delta \ln p^{\mathrm{H}}}$	Structure share $(1-\alpha)\frac{\Delta \ln p^{\mathrm{S}}}{\Delta \ln p^{\mathrm{H}}}$	Ratio of land to structure shares $\frac{\alpha\Delta \ln p^{\mathrm{L}}}{(1-\alpha)\Delta \ln p^{\mathrm{S}}}$
Halifax	0.2517	0.0000	1.0385	0.5015	2.0709
	(0.0022)		(0.0092)	(0.0015)	(0.0244)
Quebec City	0.2755	0.0000	0.7598	0.2855	2.6614
	(0.0016)		(0.0044)	(0.0006)	(0.0211)
Montreal	0.2201	0.0496	0.2593	0.7692	0.3372
	(0.0028)		(0.0032)	(0.0027)	(0.0054)
Ottawa	0.1888	0.0000	−0.0268	0.9971	−0.0243
	(0.0007)		(0.0000)	(0.0008)	(0.0001)
Toronto	0.3492	0.0301	0.1540	0.8995	0.1712
	(0.0014)		(0.0006)	(0.0020)	(0.0011)
Winnipeg	0.2596	0.0000	0.4525	0.5650	0.8009
	(0.0007)		(0.0012)	(0.0005)	(0.0028)
Saskatoon	0.2585	0.0000	0.5807	0.3553	1.6341
	(0.0017)		(0.0038)	(0.0008)	(0.0145)
Edmonton	0.2939	0.0000	0.5974	0.4472	1.3359
	(0.0024)		(0.0048)	(0.0015)	(0.0152)
Calgary	0.2851	0.0972	0.3785	0.6243	0.6062
	(0.0030)		(0.0040)	(0.0027)	(0.0091)
Vancouver	0.3766	0.0000	0.1419	0.9160	0.1550
	(0.0027)		(0.0010)	(0.0040)	(0.0018)

Note: Standard errors are in parentheses.

[a] F-test of constant returns to scale ($\alpha + \beta = 1$) in the Cobb-Douglas cost function.

Table 4: Land and Structure Shares: ECM with AIC-selected lag length

City	Lags[a]	$Q(12)$[b] (p-value)	Factor shares α	Factor shares β	CRS[c] (p-value)	Growth Shares: Land share $\alpha\frac{\Delta \ln p^{L}}{\Delta \ln p^{H}}$	Growth Shares: Structure share $\beta\frac{\Delta \ln p^{S}}{\Delta \ln p^{H}}$	Growth Shares: Ratio of land to structure shares $\frac{\alpha\Delta \ln p^{L}}{\beta\Delta \ln p^{S}}$
Halifax	7,6,7	0.974	0.2487	0.7436	0.617	1.0262	0.4983	2.0591
			(0.0252)	(0.0178)		(0.1040)	(0.0190)	(0.2494)
Quebec City	4,4,5	0.182	0.2967	0.7529	0.101	0.8181	0.2967	2.7576
			(0.0157)	(0.0404)		(0.0432)	(0.0159)	(0.2753)
Montreal	12,12,5	0.780	0.2430	0.7664	0.025	0.2864	0.7558	0.3789
			(0.0140)	(0.0142)		(0.0165)	(0.0140)	(0.0286)
Ottawa	6,7,6	0.873	0.1681	0.8090	0.565	−0.0239	0.9944	−0.0240
			(0.0295)	(0.0137)		(0.0042)	(0.0168)	(0.0040)
Toronto	6,5,4	0.994	0.3261	0.6392	0.311	0.1439	0.8834	0.1628
			(0.0267)	(0.0176)		(0.0118)	(0.0244)	(0.0134)
Winnipeg	12,7,7	0.823	0.2472	0.7626	0.346	0.4309	0.5819	0.7404
			(0.0105)	(0.0200)		(0.0184)	(0.0153)	(0.0505)
Saskatoon	8,4,4	0.985	0.2747	0.7326	0.001	0.6170	0.3511	1.7575
			(0.0076)	(0.0086)		(0.0171)	(0.0041)	(0.0689)
Edmonton	9,12,12	1.000	0.2868	0.7460	0.020	0.5828	0.4725	1.2334
			(0.0218)	(0.0310)		(0.0443)	(0.0196)	(0.1423)
Calgary	5,5,12	0.541	0.2577	0.7629	0.032	0.3421	0.6662	0.5135
			(0.0243)	(0.0282)		(0.0323)	(0.0246)	(0.0666)
Vancouver	4,3,2	0.156	0.2325	0.5757	0.068	0.0876	0.8459	0.1036
			(0.0857)	(0.0431)		(0.0323)	(0.0633)	(0.0370)

Note: Standard errors are in parentheses.

[a] Lags selected by the Akaike information criterion for an autoregressive distributed lag model in the variables $\ln p^{H}$, $\ln p^{L}$, and $\ln p^{S}$.

[b] Box-Pierce-Ljung portmanteau statistic for residual autocorrelation up to 12 lags.

[c] Wald test of constant returns to scale ($\alpha + \beta = 1$) in the Cobb-Douglas cost function.

Table 5: Land and Structure Shares: ECM with BIC-selected lag length

						Growth Shares		
		$Q(12)$[b]	Factor shares		CRS[c]	Land share	Structure share	Ratio of land to structure shares
City	Lags[a]	(p-value)	α	β	(p-value)	$\alpha \frac{\Delta \ln p^{\mathrm{L}}}{\Delta \ln p^{\mathrm{H}}}$	$\beta \frac{\Delta \ln p^{\mathrm{S}}}{\Delta \ln p^{\mathrm{H}}}$	$\frac{\alpha \Delta \ln p^{\mathrm{L}}}{\beta \Delta \ln p^{\mathrm{S}}}$
Halifax	1,1,1	0.016	0.2396	0.7503	0.295	0.9886	0.5028	1.9661
			(0.0147)	(0.0106)		(0.0606)	(0.0071)	(0.1426)
Quebec City	4,4,4	0.143	0.2934	0.7607	0.069	0.8091	0.3000	2.6991
			(0.0148)	(0.0390)		(0.0407)	(0.0154)	(0.2561)
Montreal	2,2,2	0.099	0.2312	0.7756	0.201	0.2725	0.7649	0.3562
			(0.0152)	(0.0158)		(0.0179)	(0.0155)	(0.0303)
Ottawa	3,3,3	0.005	0.1553	0.8041	0.371	−0.0220	0.9883	−0.0223
			(0.0340)	(0.0148)		(0.0048)	(0.0182)	(0.0046)
Toronto	4,2,2	0.763	0.2995	0.6362	0.081	0.1321	0.8792	0.1503
			(0.0296)	(0.0168)		(0.0131)	(0.0232)	(0.0146)
Winnipeg	2,2,2	0.002	0.2489	0.7621	0.067	0.4340	0.5816	0.7462
			(0.0063)	(0.0119)		(0.0110)	(0.0091)	(0.0303)
Saskatoon	1,1,1	0.635	0.2607	0.7490	0.000	0.5856	0.3589	1.6315
			(0.0055)	(0.0062)		(0.0124)	(0.0029)	(0.0475)
Edmonton	2,2,2	0.000	0.2839	0.7424	0.038	0.5768	0.4702	1.2265
			(0.0193)	(0.0286)		(0.0393)	(0.0181)	(0.1286)
Calgary	3,3,3	0.032	0.2810	0.7479	0.128	0.3730	0.6531	0.5712
			(0.0376)	(0.0426)		(0.0499)	(0.0372)	(0.1066)
Vancouver	1,1,1	0.021	0.2633	0.5845	0.094	0.0992	0.8588	0.1155
			(0.0799)	(0.0421)		(0.0301)	(0.0619)	(0.0359)

Note: Standard errors are in parentheses.

[a] Lags selected by Schwarz's Bayesian information criterion for an autoregressive distributed lag model in the variables $\ln p^{\mathrm{H}}$, $\ln p^{\mathrm{L}}$, and $\ln p^{\mathrm{S}}$.

[b] Box-Pierce-Ljung portmanteau statistic for residual autocorrelation up to 12 lags.

[c] Wald test of constant returns to scale ($\alpha + \beta = 1$) in the Cobb-Douglas cost function.