Passing the Buck: Impacts of Commodity Price Shocks on Local Outcomes

Arthur Grimes and Sean Hyland

Motu Working Paper 13-10
Motu Economic and Public Policy Research

September 2013

Author contact details
Arthur Grimes
Motu Economic and Public Policy Research
arthur.grimes@motu.org.nz

Sean Hyland
Motu Economic and Public Policy Research
sean.hyland@motu.org.nz

Acknowledgements
We thank the Kelliher Trust for a grant that enabled this research to be conducted. We also thank Inessa Love for making her STATA panel-VAR routines openly available, and to Motu colleagues for helpful comments on an earlier presentation of these materials. The authors alone are responsible for the analysis and views expressed.

Motu Economic and Public Policy Research
PO Box 24390
Wellington
New Zealand

Email	info@motu.org.nz
Telephone	+64 4 9394250
Website	www.motu.org.nz

 ISSN 1176-2667 (Print), ISSN 1177-9047 (Online).

Abstract

The extent to which exogenous international agricultural price fluctuations are internalised by rural communities is of major interest for policy-makers concerned with regional economic performance. So too is the link between rural sector performance and urban outcomes, especially in agriculturally-based economies. Through vector autoregressive (VAR) modelling we estimate the causal effect of exogenous commodity price innovations on both rural and urban community outcomes. Our analysis demonstrates that restricting the focus to national effects may lead to incorrect inference. We therefore extend the analysis to a VAR using panel data covering all New Zealand districts over 1991–2011. House prices and housing investment are used as quarterly indicators of regional economic and population outcomes. By exploiting the variation in production bundles across communities we find that an increase in commodity prices leads to a permanent increase in housing investment and house prices across the country. However, we find that rural communities are relatively insulated from commodity price shocks, whereas urban areas are most affected by commodity price shocks. We discuss the reasons why this paradoxical result may arise.

JEL codes

E32; R11; R12; R31

Keywords

Agricultural commodity prices; rural-urban interactions; house prices; housing investment

Contents

1. Introduction

International commodity prices are a major determinant of economic performance for a small open economy with a significant degree of commodity exports (Mendoza, 1995; Kose, 2002; Grimes, 2006). Such prices are set to equalise international supply and demand. As a result, the economic experience of local communities which have a significant degree of commodity production is vulnerable to external commodity price shocks.

The focus of this study is to improve the understanding of how commodity price fluctuations are realised across a commodity-producing economy. That is, to what extent are exogenous commodity price shocks passed from producers to the surrounding rural economy, and how does rural economic performance determine the outcomes of urban communities? This research is important to understand the extent to which volatility in rural (or urban) economies occurs following international commodity price movements, with implications for the role for policy. Characterising the link from commodity producers to rural and urban economies can also provide valuable information for the analysis of other issues. For instance, this research can be used to provide an estimate of the final spatial incidence of climate change policies or removal of agricultural supports which place the initial incidence on the producer. Our results indicate that the final incidence across rural versus urban areas of the effects of such changes in policy or other exogenous variables may be very different from the location of the direct impacts.

Recent research has touched on these questions. Wu and Gopinath (2008) used "agricultural net returns", in place of international commodity prices, as an instrument for outcomes in local activity variables (although that is not their major focus). Hornbeck and Keskin (2012) consider the implication of new technology on agriculturally focused counties and the degree of economic spillovers to other counties, while Hornbeck (2012) examines impacts on rural outcomes of a local environmental disaster. However, our study is novel in examining the effects of a rural (commodity price) shock on urban community outcomes, as well as those of rural areas.

A number of variables can be used to quantify the dynamic adjustment of a local economy to shocks, such as population, wages or housing market outcomes. We draw on the last of these due to data availability as well as theoretical reasons; housing values reflect local productivity and amenity values whilst the rate of new housing investment reflects changing population location preferences and employment opportunities. We employ these two housing measures to summarise the price and population reactions of a local economy to a shock. One might also consider farm sales prices to be an important indicator of rural community performance, and one that is sensitive to commodity price innovations. We estimate farm sales price responses at the national level; however, the analysis cannot be credibly replicated at a subnational level due to the small number of farm sales within many districts.

To estimate the impacts, we estimate a set of vector autoregressive (VAR) models and simulate the effects of a shock to world commodity prices on house prices and housing investment, or farm prices, within New Zealand. Cespedes and Velasco (2012) show that the impact of commodity price shocks on output and investment dynamics in commodity production-intense economies is decreasing in the flexibility of a country's exchange rate regime. Because New Zealand has maintained a floating exchange rate across our focal period, we include a commodity-weighted local exchange rate to allow for this transmission mechanism through which international price movements are realised by domestic producers.

First, we analyse impacts at a national level. At this level, we may expect to see only muted responses to a commodity price shock, as much of the national economy is unaffected directly by commodity prices, whilst aggregation averages out some of the rich region-specific experiences that result from heterogeneity in local production bundles. Furthermore, analysis at this level involves a relatively small aggregate sample, while macroeconomic conditions may offset the impacts of the external shock. Consistent with such caveats, we find little evidence of an effect from commodity price shocks on national housing outcomes. In fact, the estimation suggests that both national house prices and housing investment fall following a positive commodity price shock. The results for farm prices are in the expected direction, but are insignificant.

Given such analytical caveats, we extend the VAR model to panel data on Territorial Local Authorities (TLAs), which correspond to the 72 mainland local council boundaries of New Zealand. We exploit cross-sectional variation in a local land value-weighted commodity price index to estimate the causal impact of international prices on local outcomes. The analysis indicates that a permanent increase in commodity prices of 1% leads to a 0.09% increase in long run house prices, and a 0.01% increase in the long run housing stock, averaged across all TLAs.

To examine where such effects might be concentrated we use the total share of TLA land value that is attributable to commodity production to categorise each local authority as rural, quasi-rural or urban. The subnational analysis is then conducted on each classification, as well as for the sample of all communities. We find evidence of a significant degree of spatial heterogeneity in the associated commodity price effects. Our analysis suggests that both rural and quasi-rural communities do not internalise commodity price innovations; commodity price changes have little significant effect on either housing outcome within such TLAs. In contrast, urban areas experience strong house price growth and a noticeable increase in housing investment in response to an increase in the returns to commodity production. We estimate that a permanent increase in commodity prices of 1% leads to a 0.45% increase in long run house prices and a 0.08% increase in the long run housing stock in urban areas. Thus we find strong evidence that urban economic performance (in a commodity producing country) reflects the profitability of commodity production, whilst rural communities are relatively insulated from the associated shocks.

One potential concern in the above analysis is that the estimated urban effect is primarily driven by the land use type of the small proportion of land value in commodity production in urban districts, which may be little more than a statistical artefact. As a robustness check, and to clarify the link between rural production and urban community experiences, we generate an alternate price series where the weight on each commodity price reflects the total share of land value attributable to commodity production for each TLA, and the composition of commodity production within the wider Regional Council. Importantly, the results of this analysis do not differ significantly from those discussed above, providing further support for our conclusions.

Our analysis cannot identify the channels by which urban communities internalise commodity price shocks; however there exist a number of potential mechanisms. The difference in regional effects is likely due to the constraints on employment opportunities in rural communities and the indirect income effect experienced by urban areas. This latter effect may be driven by greater rural demand for urban-based professional services and durable goods sold in urban areas due to increased international returns. Alternatively, the effect may be due to an income effect associated with the resulting appreciation of the exchange rate.

The remainder of this paper is structured as follows. Section 2 details our data whilst Section 3 provides a brief descriptive analysis of the national data as further motivation for the study. Section 4 discusses the methodology, in terms of stationarity, modelling, and the optimal number of lags within the models. The results from the national level VARs are presented in Section 5, followed by the results from the TLA panel VARs across various subsamples in Section 6. It is these latter results that form our key focus. A brief set of conclusions, with suggestions for extensions to the analysis, completes the paper.

2. Data

The empirical analysis detailed in this study draws upon official data. All series used are of quarterly frequency (in some cases having been converted to quarterly frequency where quarterly data were not available). The focal period extends from 1991Q2 to 2011Q2, yielding 81 quarterly observations. The beginning of this period reflects the formation of the Territorial Local Authorities (TLAs), which represent our cross-sectional unit of observation in the subnational panel data analysis. We use the official TLA boundaries directly preceding the 2010 amalgamation of Auckland's seven TLAs; thus we bundle Christchurch City and Banks Peninsula as a single TLA. All series used in the analysis are discussed briefly in the paragraphs below, and are shown in Table 1 along with their notation and the source of the series. For a more in-depth discussion of both the publicly available and derived series, see the Appendix.

To consider movements in world commodity prices, we draw upon the ANZ Commodity Price Index (ANZCPI) series, which reports the monthly prices of an overall bundle of commodities as well the indices for the component categories: meat, skins and wool (MSW); dairy; horticulture; and forestry.[1]. The analysis draws only on the component series, as we wish to consider movements in the bundle of prices that are relevant to each district. Importantly, the component series are available in both New Zealand dollars and in international prices (which draws upon different currencies depending on the component index), which allows us to consider international price movements and define an implied exchange rate. Each series is indexed to 100 as at January 1986.

Figure 1 depicts the evolution of each component commodity price index over time. This graph highlights the relatively high and low volatility of the dairy and forestry indices respectively, but also captures the motivation to consider the analysis at a subnational level. To see this, consider the performance of each price series immediately following 1992Q3, indicated by the vertical line. Over the next 5 quarters the international dairy price decreases whilst the MSW index increases. To the extent that local outcomes respond positively to commodity prices, communities devoted to dairy production may be expected to experience a negative impact on outcomes whilst MSW intensive areas could be expected to enjoy positive effects. However, if the two districts are aggregated we may observe little change in both the aggregate commodity price index and outcomes, and as a result fail to estimate the true relationship between commodity prices and economic outcomes.

[1] There also exist component indices for seafood and aluminium, but the expression for our derived commodity price index places little weight on such series and therefore we combine such weights into a catch-all 'other' category.

As implied above, communities will care about individual commodity prices to differing degrees. To generalise the impact of commodity prices on outcomes we construct a real commodity price index ($PCom$) for each TLA (and New Zealand as a whole) which reflects the value-weighted composition of commodity production within each locality. Weights are calculated according to the value of land devoted to each commodity; land that is not used in commodity production is also assigned a weight (so that the weights sum to one) and the corresponding nominal price attributable to production is the Consumer Price Index (CPI). The deflator for the combined commodity index is also the CPI so that $PCom$ tends towards a constant as the weight for the value of land attributed to non-commodity production tends toward one. A locality-specific exchange rate (XE) is derived in a similar manner, as a weighted average of the component indices expressed in world prices relative to that in NZ dollars. Importantly, the international commodity price components of the ANZCPI are defined by different bundles of exchange rates, reflecting different international markets; therefore, we obtain a series exhibiting regional variation in effective exchange rates.

The national and local outcomes we consider are real house prices (PH) and the housing investment rate (HC/H_{-1}), where the latter approximates the growth rate of the housing stock. We also consider real farm sales prices (PF) at the national level as an important outcome by which commodity price fluctuations may be realised. The house prices measure draws upon the quality-adjusted nominal sales price index reported by Quotable Value New Zealand (QVNZ), which is available for New Zealand as well as for all TLAs, and is constructed using the Sales Price Appraisal Ratio method of Bourassa et al (2006). The housing investment rate is defined as the number of new housing building consents (HC) approved by a TLA in a quarter, relative to the estimated size of the housing stock (H) in the previous quarter. The latter series is constructed using the five-yearly censal dwelling count, adjusted by consents for new construction between censal observations. The national farm price series is constructed as a weighted average of the average sales price per hectare across all farm sales for a given commodity category in a quarter. Both farm and house prices are deflated by the CPI to consider inflation-adjusted movements. We do not examine the impact on farm prices across district subsamples, as the small number of sales for many districts induces considerable noise in the resulting series.

Other local outcome measures could also be of interest, but our choice is limited by data availability at the TLA level. One particular additional variable of interest is population. In the long run, however, housing consents should reflect population changes. Individuals moving into an area will require somewhere to live, whilst if residential growth is constrained it is likely that local employment and internal migration will also be constrained. This is highlighted in Figure 2, where we plot the 15-year change in population and total consents issued over the same period for each TLA in New Zealand. The graph highlights the strong correlation between the two variables, suggesting an additional consent is associated with an additional 2.67 individuals in the district, with an R-squared for the fitted line of 0.94. Thus, to the extent that we can identify housing investment impacts from a shock, we can also infer long-run population changes.

3. Descriptive Analysis

Before we consider the results of any dynamic parametric relationship, it is useful to examine national comovements between the international price series and those of our focal outcomes, to explore whether a prima facie relationship exists.

First, we consider the exchange rate. The relevant exchange rate for our analysis is a weighted commodity price implied exchange rate, (XE), as described in the previous section. Alternatively, we could consider the simple unweighted ratio of the aggregated ANZ World Commodity price index, expressed in international prices, relative to the same index expressed in New Zealand dollars. Our derived (national) exchange rate and its unweighted alternative, together with the USD/NZD exchange rate, are depicted in Figure 3. As might be expected, there is a high degree of correlation between the three exchange rate measures. This supports the use of our construct as well as our understanding of its movements, which is preferred for its cross-sectional variation.

Secondly, consider the comovements between the natural logarithm of the derived national commodity price index in world prices and the analogous exchange rate index. We consider the natural log on the prior that percentage changes in commodity prices have a relationship with percentage changes in the exchange rate. The comovements are detailed in Figure 4. The two series are strongly positively correlated, consistent with the conclusion of Cespedes and Velasco (2012) that the exchange rate is a material component of the transmission mechanism from commodity price shocks to local outcomes.

Now consider how each focal outcome moves with world commodity prices at the national level. First, consider the relationship with national house prices, depicted in Figure 5, where the upper graph plots the natural log of the commodity price and national house price indices, whilst the lower graph depicts the change in the log series. There is evidence of a strong positive correlation between the levels of the national commodity price and house price indices. Both series are trending upwards from the beginning of the sample period until the late 1990s, flat around the year 2000, and increasing during the mid 2000s. The lower panel also suggests that a dynamic relationship exists. We see that strong growth in one series is generally coupled with strong growth in the other. Of course this may simply pick up a correlation of price appreciation in a number of asset markets, as observed during the 2000s; thus, care will be required for a causal interpretation in the analysis that follows.

The comovement between the log national commodity price index and housing stock is depicted in Figure 6. It is difficult to visualise a relationship in levels, given the scale, and slow adjustment, of the housing stock as displayed in the upper panel of Figure 6. However, there is a clear relationship in changes; both series are increasing during the early to mid 1990s, and both series display similar dynamics for most of the 2000s, including the drop associated with the global recession.

Of course, there exists an interdependency between house prices and the housing stock, but to include both outcomes simultaneously in our analysis requires knowledge of their statistical ordering, akin to Granger causality. To consider the reduced form relationship, we plot the comovements in the log national house price index and housing stock in Figure 7. The series in changes highlights a strong relationship between the two series where house price growth appears to lead housing stock changes, which we approximate by the housing investment series. Given these results, together with the fact that housing consents are reported quarterly only, it is reasonable to assume that house prices affect housing investment, without feedback, in the contemporaneous period. This is consistent with the q-theory of housing construction in Grimes and Hyland (2013).

Finally, consider the analogous comovement between the national commodity price index and the national farm sales price per hectare series, displayed in Figure 8. Whilst the correlation in levels appears negative before 1995Q4, thereafter we see strong evidence of a positive relationship; both series are declining following 1995Q4, then relatively flat around 2000, growing during the mid 2000s period and then having similar dynamics in the face of the global financial crisis. It is harder to see such a relationship as we consider the series in changes. Noise, aggregation and measurement difficulties in the farm price data may explain the negative relationship in levels early in the sample period, as well as the weakly correlated changes. Another possibility is that farm prices are driven as much by other factors (e.g. credit availability) as they are by commodity returns, which cannot be stripped out in such graphical analysis, although the same may be true of house prices.

Overall, a prima facie relationship appears to exist between changes in commodity prices, the exchange rate and housing market outcomes, as well as between the two housing market outcomes. This motivates the structural restrictions applied to our analytical model, which looks at the impact of commodity price innovations on the change in each price series and the level of the housing investment rate. We proceed to parametrically test the effect on farm prices at the national level also, however caution is required in the interpretation given the descriptive evidence of this section.

4. Methodology

Impulse response functions (IRFs), based on vector autoregressive models (VARs), are the key analytical tool used in this study. Estimating the parameters of a VAR requires that the series are covariance stationary, thus we examine individual stationarity in Subsection 4.1, given the relationships identified in Section 3. However, even with this requirement met, it is still possible for a system of individually stationary series to be unstable. That is, shocks to the system can cause the IRFs to explode over time. In such cases, interpretation of IRFs is invalid. The number of lags present in the model, as discussed in Subsection 4.4, is chosen to be that which yields stable national and subnational systems and thereafter minimises a model selection criterion. Subsections 4.2 and 4.3 describe the structure of our national and subnational VAR models.

4.1. Stationarity Testing

Given the relationships displayed in Section 3, we test the stationarity of the change in each price series as well as the level of the housing investment rate (which approximates the percentage change in the housing stock).

At the national level, we test for individual stationarity using an augmented Dickey-Fuller (ADF) test, where the null hypothesis is a series contains a unit root. The number of lags utilised in each test is chosen to minimise the Schwarz Bayesian Information Criterion (SBIC). The ADF test statistics are presented in Table 2. Based on these results we strongly reject the null hypothesis of a unit root for all price series in changes at conventional significance levels.

However, we fail to reject the null hypothesis for the level of the housing investment rate. This is surprising as Figure 7 suggested that the change in house prices and the housing investment rate have a very close relationship. The divergence in series occurs in the period following the onset of the Global Financial Crisis; house price changes reverted towards their mean quickly, whilst housing investment has remained well below the historical average. Consequently, we could test for a unit-root in the national housing investment rate series up to 2007Q3; the p-value from the corresponding ADF test (with 1 lag) is 0.0077, implying the series is stationary subject to a structural break at the time of the GFC.

Evidence of a stationary national aggregate series does not imply that the series at a local level are stationary, which is required if we are to extend the VAR model to a subnational level. We test the stationarity after cross-sectionally demeaning and Helmert-differencing the series to remove time and area fixed effects, respectively, as this is the form in which they appear in the model. We hereafter refer to this form of each series as the transformed series, and denote this by a tilde above the variable name. Further discussion of this transformation, and the motivation for it, appears in Subsection 4.3.

To test for stationarity in each subnational transformed series we use the Levin-Lin-Chu (LLC) and Im-Pesaran-Shin (IPS) panel unit root tests, each of which has the null hypothesis of a unit root, where the number of lags is chosen to minimise the SBIC. The associated test statistics and p-values, across all districts, are shown in Table 3. The LLC and IPS tests reject the null hypothesis of a unit root for all series at all conventional significance levels. This evidence allows for subnational analysis on price changes and the housing investment rate in levels, a framework consistent with the national level analysis.

4.2. National Level Modelling

Initially, we employ a VAR model to estimate the effect of commodity price movements on our focal outcomes at the national level. We do so since this is the level of aggregation that is often adopted in applied macroeconomic analysis. As a result of the national stationarity tests, we consider the effect of a one-off shock to the change in world commodity prices, interpreted as a permanent change in the level of commodity prices, on the change in national house prices and on the national housing investment rate. We also consider the effect of a permanent commodity price shock on the change in national farm prices.

We posit that commodity prices can influence each focal outcome contemporaneously and with a lag. We also consider the role of the exchange rate as an intermediate mechanism; the greater the appreciation of the exchange rate that results from an increase in world commodity prices, the lower the expected impact on national outcomes.

We allow for a deterministic time trend to affect at least one series in levels; thus, a constant is included in the national VAR as we analyse variables in changes. Accordingly, we adopt the following compact form of the stacked unrestricted national structural model, using the notation of Fry and Pagan (2011):

$$\beta z_t = B_0 + \sum_{s=1}^{P} B_s z_{t-s} + \varepsilon_t \tag{1}$$

where: $z_t = \begin{bmatrix} \Delta lnPCom_t \\ \Delta lnXE_t \\ \Delta lnPH_t \\ HC_t/H_{t-1} \end{bmatrix}$ or alternatively: $z_t = \begin{bmatrix} \Delta lnPCom_t \\ \Delta lnXE_t \\ \Delta lnPF_t \end{bmatrix}$

where z_t is a stacked column vector of national observations across series, β is a block matrix of contemporaneous parameters, B_0 is column vector of unique constants for each dependent variable, B_s is a block matrix of parameters on lagged variables (at lag length s), and ε_t is a stacked column vector of shocks that have zero mean, a constant covariance matrix and no serial correlation. We obtain 81 quarterly observations for each series at the national level; however, P observations are lost from each series for estimation, which denotes the maximum lag length in the system.

Given the potential contemporaneous correlations between dependent variables, we cannot uniquely identify the coefficients in equation (1) as it stands. Instead we consider the reduced form specification by pre-multiplying the system by β^{-1}, leading to a system where there are no contemporaneous variables on the right hand side. The structural model includes $n(n-1)/2$ more parameters than are contained within the reduced form specification, where n is the number of dependent variables in the system. Thus we require 6 restrictions for the housing outcomes model, or 3 restrictions for that of farm price outcomes, to identify the structural form parameters respectively. A common solution to identifying the structural parameters from reduced form estimation is to use the type of recursive model proposed in Sims (1980), involving the imposition of certain theory-based restrictions that we apply to the β matrix.

There exist several restrictions to employ from economic theory. First, we allow world commodity price shocks (assumed to be determined by factors outside New Zealand) to impact contemporaneously on the demand for New Zealand dollars and hence on New Zealand's exchange rate. Woods and Coleman (2012) show that movements in the value of the New Zealand dollar have no reverse impact on international commodity prices, which are set to equate global supply and demand. Thus we impose the restriction that New Zealand's exchange rate does not impact contemporaneously on world commodity prices.

Woods and Coleman (2012) also consider the extent to which New Zealand as a whole influences international commodity prices and find that there is little evidence of any influence. A natural restriction, therefore, is to restrict the contemporaneous effect of changes to the national outcome series on international commodity prices to be zero.

Furthermore, we assume that domestic housing outcomes should exhibit no contemporaneous influence on the exchange rate, which represents another price that is set on international markets to equate international supply and demand. However, fluctuations in the exchange rate lead to adjustments in the purchasing power of the domestic currency, which can have real effects on domestic income and consumption and thus national outcomes.

These three restrictions mentioned are all that is required to identify the structural form parameters in the system containing farm prices, whilst an additional restriction is required for the alternate system.

Consider the interdependency and the required ordering between house prices and housing investment. From a q-theory of investment, as set out in Grimes and Hyland (2013), developers react to disequilibria in house prices and their replacement costs. Thus, there is a strong link from prices to housing consents (housing investment). Due to reporting conventions, where consents data are released only quarterly, it is unlikely that potential house purchasers have knowledge of housing consents in the same period. Further, Figure 7 provides descriptive support for house prices leading housing investment. As a result, we assume that housing consents do not contemporaneously affect house prices.

We premultiply equation (1) by the restricted β^{-1} matrix and obtain a restricted reduced form national VAR model which is estimable for a given autoregressive order, under the assumption that the errors have zero mean, a constant covariance matrix and no serial correlation. The choice of order, as well as the reasonableness of these assumptions, is discussed in Subsection 4.4.

4.3. Subnational Level Modelling

Given the potential for averaging out local effects in national level impacts, we repeat the analysis through a panel VAR, where TLAs comprise the cross-sectional unit of observation. A comparison of dynamic responses across all TLAs, as well as sub-samples of such, can be used to test the responsiveness of rural economies to commodity price innovations, and to ascertain the relationship between rural sector performance and urban economic experiences.

A lack of farm sales in a large number of district-quarter observations results in a considerable degree of noise in the farm sales price data. As such, it is hard to separate the property specific noise from the overall trends in a region, and thus we cannot credibly consider the impact on farm prices at a subnational level.

Given the results of the subnational stationarity tests we seek to estimate the parameters of a VAR using the transformed price series in changes and the transformed level of the housing investment rate. A central requirement of applying VAR analysis to panel data is that each cross-sectional-time observation has the same data generating process; this is unlikely to hold in practice. As such, we can exploit the panel structure of the subnational dataset to allow for heterogeneity along several dimensions. Firstly, districts may experience common period-specific shocks, for example, due to variation in macroeconomic conditions or seasonality. Furthermore, there exists a possibility that commodity prices and local house or land prices are both influenced by movements in world asset prices. To the extent that such factors affect all TLAs equally this can be captured through a time fixed effect. Secondly, we allow districts to have a different average rate of change in each series when all other series are constant. This accommodates systematic effects such as changing lifestyle preferences for certain districts, or systematic differences between districts due to land or construction constraints. In this way, we identify the causal impact of commodity prices on local outcomes, after stripping out area-specific and macroeconomic factors, by exploiting the variation in international prices over time and the variation in the bundle of commodities across districts.

To estimate the model in the presence of period-specific (e.g. macroeconomic) shocks we subtract the period-specific mean across all TLAs from each series, which we refer to as cross-sectional demeaning. This demeaning process has the effect of reducing the volatility of the rural commodity price index whilst inducing a greater degree of volatility in the urban commodity price index than seen in the raw series. The result is an urban commodity price index that has a standard deviation (in changes) that is only 33% less than that of the rural category, and 48% greater than the quasi-rural category. This outcome highlights the strength of our methodology, which uses the *relative* commodity price outcomes across districts to identify their impacts on the focal outcomes.

A common method to remove area fixed effects is to first-difference each district series. However, in the presence of lagged dependent variables present in the model such differencing introduces a correlation between the error term and the lagged dependent regressor. The solution we employ is to consider a Helmert transformation of the data, essentially a forward-mean differencing over each district-specific series, following the discussion set out in Love and Ziccino (2006) and detailed in Hamilton (1994). The cross-sectionally demeaned, Helmert-differenced form of the series is referred to as the transformed series.

Following the presentation of the national structural model, as well as the discussion of the restricted reduced form national VAR, we estimate the reduced form of the restricted panel VAR with the transformed subnational series, under the same set of restrictions, as follows:

$$z_{it} = A_0^C + \sum_{s=1}^{P} A_s^C z_{i,t-s} + e_{it} \qquad (2)$$

$$\text{where } z_{it} = \begin{bmatrix} \Delta \widetilde{lnPCom_{it}} \\ \Delta \widetilde{lnXE_{it}} \\ \Delta \widetilde{lnPH_{it}} \\ \widetilde{HC_{i,t}/H_{i,t-1}} \end{bmatrix}$$

Thus z_{it} is a stacked column vector of national observations across districts and series, A_0 is column vector of unique constants for each dependent variable, A_s is a block matrix of parameters on lagged variables (at lag length s), whilst e_{it} is a stacked column vector of shocks that have zero mean, a constant covariance matrix and no serial correlation. The C superscript on the model coefficients implies the subnational coefficients can vary across the possible categories of subnational TLAs (All, Rural, Quasi-Rural or Urban). As above, P observations are lost from each series for estimation, as this is number of lagged regressors present the system.

Equation (2) is estimated under the assumption that the errors are serially uncorrelated and the system is stable, for a given autoregressive order. We consider how to choose the order, and the reasonableness of these assumptions, in the following subsection.

4.4. Lag Length and System Stability

A requirement of VAR analysis is that the residuals are serially uncorrelated (Fry and Pagan, 2011). As a result, the model must include sufficient lags to account for the serial correlation. For example, increased farmer returns could potentially stimulate building activity, either locally or across urban areas. Such effects may be slow to be observed due to administrative processes or spending patterns; failure to account for such relationships could generate significant serial correlation in the residuals.

To identify serial correlation we estimate the subnational model (for greater precision) under increasing lag lengths and consider the significance of the autocorrelation in the estimated residuals from each equation as our first indicator of the optimal lag length. The results (which are not presented) suggest the presence of serial correlation in at least one equation for all lag lengths less than 4.

We must also consider the efficiency of the model, which could decrease if lagged series with insufficient predictive power are included in the model. Given at least 4 lags are required for validity, we consider whether additional lags should be considered by reference to the Akaike Information Criterion (AIC) or the SBIC. Table 4 presents these statistics for each equation and TLA sample, where the column header defines the equation by the transformed dependent variable. For each dependent variable, and across all samples, each information criterion suggests that as few lags as possible be included in the model. The two results together suggest the use of 4 lags, which is consistent with the use of quarterly data.

The national estimation with an autoregressive order of 4 yields a coefficient matrix with negative eigenvalues, thus we obtain a stable solution for the system which allows for inference of the resulting impulse response functions (IRFs). Estimation of the analogously defined subnational model also yields a stable solution. Thus we can proceed to simulate a shock to the change in commodity prices from within the model and estimate the effect on our set of outcomes.

5. National VAR Results

Let us now consider the changes in national aggregate outcomes due to changes in international commodity prices, controlling for exchange rate movements as one potential transmission mechanism, through the IRFs that result from the national level VAR modelling.[2]

[2] As exchange rate movements are not central to this paper, other than as a transmission mechanism, we omit the graphs for the effect of a change in commodity prices on the exchange rate. The same is true of the subnational analysis.

We introduce a positive one-standard deviation shock to the natural log of the national commodity price and consider the implications through the system, with 200 Monte Carlo simulations conducted to generate confidence bands around the estimates.

The impact of such innovations on the national house price index and the national housing investment rate is portrayed in Figure 9. The implications for national farm prices per hectare are displayed in Figure 10.

The left panel of each figure depicts the time path of the national commodity price index changes as a result of this shock. This panel is not materially different between alternative models, which is to be expected as commodity prices are the shock variable and we expect little, if any, feedback from the remaining variables in the model to future international commodity prices (consistent with our exogeneity assumption). The dynamic path (represented as a dashed line in the graph) reflects the impact of this shock on the change in the log of the national commodity price series, relative to its pre-shock value. By attaching a 95% confidence interval around the dynamic response, depicted as the shaded region, we can consider the statistical significance of responses. We see the dynamic response rises by a statistically significant 0.75% in period 0, the time of the shock, followed by some insignificant oscillation prior to settling down to zero.

By considering the cumulative effect, which is the sum of the dynamic responses, depicted as a solid line, we estimate the implied long-run change in the level of the national commodity price index. The coefficients from the housing (farm) outcomes model imply that a 0.75% increase in the commodity price in period 0 leads to a 1.19% (1.11%) point-estimate increase in the commodity price index in period 40. The fact that the long-run effect differs from the immediate effect suggests a degree of persistence in the changes which reflects, in part, the use of a monthly average commodity price series (so that a shock that is experienced mid-month is reflected more fully in subsequent months than in the initial month of the shock). International prices are expected to follow a random walk, or one could profit from exploiting any such pattern. However given the partially offsetting effects in periods 4 to 8, and the precision of all dynamic point-estimates, the long-run effect is not significantly different from the short-run effect. Thus, by imposing only a minimal set of restrictions, our two models produce results that are consistent with theory.

Although the graphs are omitted, the impact on the exchange rate follows expectations. The analysis suggests that a one standard deviation increase in the change in log commodity prices immediately leads to an appreciation in the exchange rate, with no evidence of feedback from exchange rate movements to future commodity price levels, a result consistent with New Zealand acting as a price-taker in international commodity markets. The long-run evidence suggests that a 1% increase in commodity prices leads to a slightly less than 1% increase in the exchange rate as export returns are bid up. The greater increase in world commodity prices relative to the exchange rate means that the commodity price shock does feed through positively to the income of domestic commodity producers, but other sectors (e.g. households and firms using imported raw materials) may also gain as the cost of their inputs reduce.

5.1. Impact on National Housing

The central panel of Figure 9 portrays the impact of the price innovation on the change in the (log) New Zealand house price index. The initial response in national house prices to the commodity price increase is an insignificant increase in the change in national house prices by 0.12%. There is no significant reaction in either of the following two periods, although the modelling suggests there are significant reductions in national house prices (relative to baseline) between the 3rd and 6th quarters after the shock.

The cumulative impulse response function suggests that long-run house prices are 3.37% lower than their preshock value as a result of the positive commodity price shock. Given the depiction of commodity prices in Figure 1 and the comovements in the change in each series in Figure 5, we believe it is improbable that a 1.19% increase in commodity prices leads to a 3.37% fall in national house prices in a commodity producing country, indicating difficulties with national level analysis of such a shock. We check the validity of this result through considering the local area effects in Section 6.

The panel on the right of the figure captures the impact on national housing investment from higher commodity prices. The estimates suggest a 0.75% initial increase in international commodity prices has a small insignificant effect on the contemporaneous national housing investment rate, which is unsurprising given that there exists a lag between observing the shock, the decision to build, application of a housing consent, and finally official approval. The dynamic series thereafter displays a pattern similar to that of house prices, which may be the key mechanism by which commodity prices affect local housing investment; there is a significant reduction in consents in the 5th to 8th quarters after the commodity price increase, with little significance observed at other lag lengths. The cumulative response point estimate suggests that the national housing stock (i.e. the cumulative housing investment response) will be permanently 0.27% lower as a result of a permanent increase in commodity prices.

Consequently a national level analysis indicates that a permanent commodity price increase is expected to permanently reduce national house prices at a factor of 2.8 times the size of the commodity price shock, with a small reduction in the housing stock also. However, a change in the price of commodities is most relevant to commodity-producing areas, while an offsetting change in the exchange rate may be most relevant to urban areas; thus we may lose the direct effect on such areas by aggregating the responses to the national level. Section 6 considers the response to international price innovations across multiple sub-samples of communities to decompose such effects.

5.2. Impact on Farm Prices

Finally at the national level we consider the impact of a one standard deviation increase in commodity prices on the national farm price per hectare, depicted in the right hand panel of Figure 10.

A 0.75% increase in commodity prices has a small negative effect on national farm prices initially, as agents evaluate whether the shock is permanent or transitory; the point-estimate on the contemporaneous farm price change in is -0.90%. Over time, the increased commodity price leads to an estimated appreciation of long-run farm prices by 0.55% above their pre-shock price level. While the direction of the aggregate farm price response is as anticipated, the dynamic response is insignificant at all quarters following the shock, which may be due to a high degree of noise in the underlying farm price data. Consequently, this analysis does not find evidence of a significant effect from national commodity prices on farm prices.

5.3. National Level Conclusions

The national level VAR analyses suggest that permanent commodity price increases reduce both house prices and the housing stock (relative to baseline) around the country, whilst there is little evidence of any effect on farm sales prices. Given the aforementioned caveats around national level analysis we next investigate the extent to which subnational relationships differ from those described above.

6. Subnational Panel VAR Analysis

To consider subnational relationships we estimate an analogous VAR model over the period 1991Q1–2011Q3, which includes a cross-sectional dimension of observations from every TLA in New Zealand.

One explanation for the lack of power in identifying effects at a national level is that there exist significant effects amongst a subsample of TLAs; however, as we aggregate over the entire country these area-specific sensitivities wash out. To test such a hypothesis we conduct Panel VAR analysis, selecting on subsamples where TLAs are defined as rural (where the average proportion of land value in commodity production over the period exceeds 44%), quasi-rural (where the average proportion of land value in commodity production over the period exceeds 20% but is less than or equal to 44%), or urban (all others). Table 5 lists all TLAs, along with their average proportion of land value attributed to commodity production over the sample period and the corresponding classifications. The boundary weights are partly chosen to place a similar number of the 72 TLAs in each classification; there are 23 urban TLAs, 21 quasi-rural TLAs and 28 rural TLAs. This is supplemented by ensuring the classification of each TLA is consistent with the generally accepted notion of a TLA's primary productive activity.

The subnational analysis follows a similar structure and presentation as that laid out in the national VAR analysis; we consider the impact of a one standard deviation innovation to the change in the subnational commodity price on the change in local house prices and the level of the housing investment rate, where all subnational series are transformed to remove heterogeneity.

6.1. All TLAs

The impact of a permanent one standard deviation shock to the transformed subnational commodity price on local commodity price indices, house prices and the housing investment rate over all TLAs is depicted in Figure 11.

The left panel shows the effect that the shock has on the subnational commodity price. The shock is equivalent to a 1.72% increase in the subnational commodity price for all districts. Because the series is stationary in changes, we see the dynamic response tends back towards zero over the next 20 quarters with a small degree of oscillation. This leads to a post-shock commodity price that, in the long run, is 2.47% above the pre-shock value.

With this permanent increase in local commodity prices, we estimate a small negative contemporaneous effect on the local house price level across all TLAs, as shown in the central panel of Figure 11, although this effect is not significant. However, in the second, fourth and fifth periods after the shock we find that there is a significant increase in local house prices, a lag which could be consistent with the time taken to realise the associated gains and for higher income levels to feed through to housing market activity. The dynamic process suggests the higher long-run subnational commodity prices drive a statistically significant 0.21% increase in long-run subnational house prices across all districts. This is a very different result from that estimated using national aggregates (which showed a reduction in national house prices). The disagreement in results highlights the benefits of conducting this analysis at a subnational level.

The right panel of Figure 11 shows the effect of the permanent increase in commodity prices on local housing investment, encompassing both the direct effect and the indirect effect as a result of the increased local house prices. The analysis suggests that the permanent increase in local commodity prices has no significant contemporaneous impact on local housing investment rates, as expected given the time required to observe and act upon a shock. However we estimate a significant increase in housing investment beginning four periods after the shock, which is consistent with the time required to receive official consent following reaction to price changes. The housing stock continues to rise relative to baseline (indicated by the cumulative effect) as the change in underlying housing investment slowly converges back towards zero. The cumulative response estimates suggest that the local housing stock across all districts, 40 quarters after the shock, is expected to be 0.03 percentage points higher as a result of increased local commodity prices. This is indicative of a permanent increase in population (i.e. net inward migration) following a positive commodity price shock, affecting national incomes.

Analysis at a subnational level therefore suggests a positive dynamic effect of commodity prices on subnational house prices and housing investment exists, which was not the case under the national analysis. That is, commodity price innovations are internalised by local economies. In the discussion that follows we attempt to identify whether the nature of the areas affects the manner in which this response is observed.

6.2. Rural TLAs

Figure 12 depicts the analogous results for the sub-sample of rural districts, being the TLAs with the highest degree of commodity production. A one standard deviation shock to the change in rural commodity prices is equivalent to a 2.23% contemporaneous increase, or an increase of 3.26% in long-run commodity prices. Note that this increase is greater than that seen across all TLAs, as our commodity price index attributes a constant real price to the proportion of land value not in commodity production.

The dynamic process of rural house prices is very similar to that witnessed across all TLAs, albeit with a lack of significance, as captured in the central panel of Figure 12. Rural house prices respond positively in the period following the shock, and the impact remains positive for the following three periods, although estimates for all quarters remain well within the 95% confidence interval. As a result the long-run house price response across rural districts has a point estimate which is positive but insignificant.

The panel on the right of Figure 12 depicts the impact of a commodity price shock on the housing investment rate in rural districts. We find that a permanent increase in local commodity prices has a small negative effect on housing investment in the contemporaneous period, as well as the two subsequent quarters, but such effects are well within a 95% confidence interval. In the five quarters that follow we estimate an increase in the housing investment rate, with a significant positive estimate in the fifth quarter after the shock. Given the dynamic changes, the point estimate for the long-run housing stock response is less than 0.01 percentage points higher that the pre-shock value.

The broad trend in point-estimate house price changes and housing investment rates are similar between rural communities and the larger sample of all TLAs. However, it appears as though rural communities contribute less than other TLAs to the long-run response in subnational housing outcomes. Furthermore, the reduced sample size compared to the all TLAs sample and relatively more volatile consents data in rural districts means that the dynamic responses are mostly insignificant following the shock, albeit with some significance for the housing investment rate. This analysis offers an interesting insight; we find little evidence that commodity price innovations are internalised by the local economy in areas dominated by commodity producers.

6.3. Quasi-Rural TLAs

Figure 13 displays the results of the analysis for the subsample of quasi-rural districts. A one standard deviation shock to the change in the commodity price index for this group of TLAs is equivalent to an immediate 0.99% increase in local commodity prices, leading to a 1.18% increase in long-commodity prices.

The direct effect of commodity price changes on local house prices appears volatile and imprecise, with the point estimate fluctuating above and below zero for the 10 quarters after the price innovation. The cumulative estimates suggest that the long-run house price amongst this subset of communities falls by 0.08%; however, this impact is not significant.

The impact of the shock on quasi-rural housing investment differs from that seen across rural districts, with no significant change in any quarter and a tiny (insignificant) negative long-run impact. Thus we fail to find any evidence that commodity price fluctuations are passed on to quasi-rural economies.

6.4. Urban TLAs

Finally, we consider the impact of a permanent commodity price shock on urban housing outcomes – as depicted in Figure 14. A one standard deviation increase in the change in log commodity prices in urban areas is equivalent to a 1.55% increase in the local index, leading to a long-run commodity price index that is 2.24% greater than its pre-shock value.

A permanent increase in commodity prices has a broadly similar effect on urban housing outcomes as that estimated for rural districts, although the magnitude and precision of the urban estimates are much greater. While there is little contemporaneous effect on house prices, in each of the following six quarters we estimate a highly significant increase in urban house price changes, after which the impact reverts towards zero. This combination of dynamics leads to a long-run house price level that is 1.02% above the baseline case after 40 quarters. The difference in house price responses across TLA classifications is marked. While we cannot preclude a small effect on long-run house prices of rural commodity producing areas, we find evidence of a pronounced increase in urban house prices.

The impact of a commodity price shock on urban housing investment is also relatively pronounced, which is unsurprising given the effect of house prices on housing construction. The permanent commodity price innovation increases the quarterly housing investment rate above baseline in each of the following 40 quarters. Such long lived dynamics are consistent with the findings of Grimes and Hyland (2013), which shows that it can take 8 years for the housing stock to fully adjust to a 5% population (migration) shock. The cumulative changes show a long-run increase in the housing stock of 0.18%, an effect consistent with a population increase in urban areas following an increase in commodity prices.

6.5. Urban Extension

The difference between the impact of commodity prices on urban housing outcomes and those of all other TLA classifications encourages some checking of the robustness of the result. One possible concern is that the commodity activity of the small proportion of land value in each urban TLA is driving the result. For example, if all urban commodity producers were engaged in the same activity, and housing outcomes were particularly responsive to that commodity's price, we would bias the regional comparisons. In response to this concern, we construct an alternative commodity price index for urban areas, where the total weight placed on commodity prices still reflects the proportion of a TLA's land value that is attributed to commodity production (relative to non-commodity production), but the relative weights on each component price reflects the commodity production composition of the Regional Council (RC) in which the urban TLA resides. Regional Councils are generally larger than TLAs and so include a greater proportion of agricultural land than do urban TLAs. The construction of the alternative index is discussed further in the Appendix.

The results of the associated analysis are not discernibly different from those discussed in Subsection 6.4, and for this reason are not presented here. We find that permanent commodity price increases lead to sustained increases in urban house prices and housing construction whether the set of prices reflects those from the urban TLA itself or from the wider region. The results of sub-sections 6.4 and 6.5 lead us to conclude that the performance of housing outcomes in urban areas is strongly connected with the experiences of local and regional commodity producers.

6.6. Subnational Level Conclusions

To ease the comparison of results, we can frame the VAR results discussed above in terms of long-run elasticities for house prices and the dwelling stock. This follows from considering the impact of a percentage change in commodity prices on the percentage change in house prices and on the housing investment rate, which approximates the percentage change in the housing stock each quarter. Table 6 reports the long-run point estimate of these elasticities for the national and subnational samples.

As suggested in the discussion of the national level analysis, the elasticity of house prices at the national level is (dubiously) negative. As we consider the various subnational samples we find that house prices respond positively across all districts on average, but the response is rather inelastic. This moderately inelastic response is driven by a moderate sensitivity in urban areas, with an inelastic but positive and negative response in rural and quasi-rural areas, respectively.

The story is similar for the elasticity of the housing stock. The national analysis suggests that an increase in commodity prices produces a relatively large reduction in the housing stock (relative to baseline), which is unlikely to hold in reality. The subnational analysis yields more plausible results in terms of magnitudes. We find a small positive impact of commodity prices on the housing stock across all districts, concentrated principally in urban areas.

The subnational analysis indicates that rural districts are broadly insulated from commodity price shocks, an initially counter-intuitive result. However, the potential for employment growth may be heavily constrained in rural districts, so there is less potential for an income effect to be widely distributed through the immediate community. Instead, the income effect may be realised as a change in the rural area's demand for goods and services, such as new farm equipment, motor vehicles and professional or business services. The provision of these goods and services may be mostly from urban areas; thus it is the latter areas that benefit most in terms of economic activity from a rural commodity price increase. The exchange rate channel may exacerbate this effect. Both the national and subnational models indicate that the exchange rate appreciates following an international commodity price increase, as predicted by theory. The exchange rate appreciation reduces the price of imported raw materials and machinery which may stimulate domestic (urban) investment. The appreciation also reduces the price of traded goods such as petrol, enabling consumers in all areas to increase consumption of non-traded goods. Again, many of these purchases will be sourced from urban areas, thus transferring the income effects of a commodity price increase to urban producers and suppliers.

We conclude that the effect of commodity price innovations on housing outcomes is felt most strongly in urban, rather than rural or quasi-rural, areas. Given the evidence that long-run housing consents is a strong predictor of long-run population change, we conclude also that the relative population growth following a commodity price increase is strongest in urban areas. In turn, this suggests that rural commodity price increases may paradoxically contribute to greater urbanisation.

7. Conclusions

Whilst previous research has considered various effects of rural productivity, this paper is the first to estimate the extent to which rural shocks (to commodity prices) are internalised by rural communities, and the extent to which urban economies experience the effects of these shocks. These questions are inherently important for a small open economy based extensively on commodity production.

To estimate such relationships, we consider the dynamic response of localised outcomes within a small open economy to an exogenous internationally-sourced innovation, across 72 districts in New Zealand between 1991Q2 and 2011Q2. Construction of a cross-sectionally varying commodity price index, which reflects the production bundle of each district, allows for a causal interpretation. Using a minimal set of restrictions and institutional knowledge, we estimate a set of VAR models which identifies the impact of commodity price innovations on national (and subnational) house prices and the housing stock. The latter variable is shown to approximate population changes in the long run. We also examine the impacts of commodity price innovations on farm prices, as a direct indicator of the profitability of commodity production, at the national level.

Prior research into commodity price effects in small open economies has largely centred on national effects (Mendoza 1995, Kose 2002, Cespedes and Velasco 2012). Our analysis, when conducted at the highly aggregated national level, indicates large negative (albeit insignificant) national responses to an exogenous commodity price shock. This may be due to a small sample size, and to ignoring the spatial distribution of benefits that is averaged away under such aggregation.

As a result of these factors, we extend the analysis through panel VAR estimation, to increase sample size along the cross-sectional dimension, and to allow consideration of different sub-samples, thus identifying region-specific responses. The results of such analysis across all districts in the country are more consistent with theory. We find evidence that commodity price increases lead to a small but significant increase in housing investment and house prices across all communities. We also find that an international commodity price increase leads to an exchange rate appreciation. This appreciation reduces the impact on the local economy, ceteris paribus, which is consistent with Cespedes and Velasco (2012).

Our most important contribution is to disentangle effects across district types. We find, paradoxically, that rural and quasi-rural communities are relatively insulated from international commodity price innovations. One potential explanation for this phenomenon is the short-run constraints on rural employment opportunities. In contrast, we find evidence that urban areas are relatively sensitive to rural profitability shocks. Whilst our analysis cannot identify the mechanism by which price shocks are realised in urban areas, a number of potential explanations exist. Firstly, rural income shocks may affect the demand for goods and services by commodity producers sourced primarily from cities. This could reflect a consumption-investment dichotomy, where consumption in rural areas is relatively constant whilst the associated income shocks induce changes to the rural demand for urban-based durable investment goods and professional services. An alternative transmission mechanism is the endogenous exchange rate appreciation. An exchange rate appreciation reduces the price of imported raw materials and machinery (promoting investment) as well as of traded consumer goods (such as petrol), providing a positive income effect for consumers. Both the positive consumer income effect and the increase in investment are likely to be felt most in urban areas, rather than rural areas.

To check the robustness of our urban results we replicate the analysis with an alternative commodity price index which reflects the composition of commodity production within the wider region which encompasses each urban TLA. This clarifies the channel from rurally concentrated profitability shocks to urban economies which we seek to examine. This analysis produces no discernible differences from those discussed above; there is little evidence that rural commodities internalise commodity price innovations, whilst the performance of urban economies remains sensitive to the profitability of commodity production.

Given that house prices and housing investment are most responsive to commodity price shocks in urban areas, future analysis could consider induced expenditure changes at a disaggregate commodity and spatial level, and the extent to which these changes are explained by exchange rate movements. Furthermore, greater policy consideration could be given to the indirect and redistributive effects associated with exogenous price innovations, which do not necessarily have the greatest effect on those areas that directly experience the shock. One such example for policy consideration may be an examination of the final incidence of climate change policy that initially impacts on the agricultural sector (e.g. through agricultural emissions). Our results suggest that while the initial cost of such a policy may be felt by the producer, it may be the urban areas that are affected most by such a change. Similarly, countries that are considering reducing agricultural support may find that the greatest impact of such changes may be felt by their urban, rather than rural, communities. While analysis of such policy options needs to be conducted explicitly to cater for the exact nature of contemplated policy changes, our results indicate that the dominant incidence of shocks may be felt in entirely different communities to those that are initially directly affected by the policy or other exogenous shock.

References

Bourassa, S., M. Hoelsi and J. Sun. 2006. "A Simple Alternative House Price Index Method", *Journal of Housing Economics* 15, pp. 80–97.

Cespedes, L. F., and A. Velasco. 2012. "Macroeconomic Performance During Commodity Price Booms and Busts", *NBER Working Paper* 18569, National Bureau of Economic Research, Cambridge, MA.

Fry, R., and A. Pagan. 2011. "Sign Restrictions in Structural Vector Autoregressions: A Critical Review", *Journal of Economic Literature* 49:4, pp. 938–60.

Gozgor, G. 2011. "Panel Unit Root Tests of Purchasing Power Parity Hypothesis: Evidence from Turkey", *International Research Journal of Finance and Economics* 61:2011, pp. 136–40.

Grimes, A. 2006. "A Smooth Ride: Terms of Trade, Volatility and GDP Growth", *Journal of Asian Economics* 17:4, pp. 583–600.

Grimes, A., F. Holmes and R. Bowden. 2000. *An ANZAC Dollar: Currency Union and Business Development*. Wellington: Institute of Policy Studies.

Grimes, A., and S. Hyland, with A. Coleman, J. Kerr and A. Collier. 2013. "A New Zealand Regional Housing Model", *Motu Working Paper* 13-02, Motu Economic and Public Policy Research, Wellington.

Grimes, A., D. C. Maré and M. Morten. 2009. "Defining Areas and Linking Geographical Data: An Example from New Zealand," *Population, Space and Place,* 16:2, pp. 165–170.

Hamilton, J. D. 1994.*Time Series Analysis*. Princeton: Princeton University Press.

Hlouskova, J., and M. Wagner. 2006. "The Performance of Panel Unit Root and Stationarity Tests: Results from a Large Scale Simulation Study", *Econometric Reviews* 25:1, pp. 85–116.

Hornbeck, R. 2012. "The Enduring Impact of the American Dust Bowl: Short- and Long-run Adjustments to Environmental Catastrophe," *American Economic Review* 102:4, pp. 1477–507.

Hornbeck, R., and P. Keskin. 2012. "Does Agriculture Generate Local Economic Spillovers? Short-run and Long-run Evidence from the Ogallala Aquifer", *NBER Working Paper* 18416, National Bureau of Economic Research, Cambridge, MA.

Kose, M. A. 2002. “Explaining Business Cycles in Small Open Economies – ‘How Much do World Prices Matter?’”, *Journal of International Economics* 56:2, pp. 299–327.

Love, I., and L. Ziccino. 2006. “Financial Development and Dynamic Investment Behavior: Evidence from Panel VAR”, *The Quarterly Review of Economics and Finance* 46:2006, pp. 190–210.

Mendoza, E. G. 1995. “The Terms of Trade, the Real Exchange Rate, and Economic Fluctuations”, *International Economic Review* 36:1, pp. 101–37.

Sims, C. 1980. “Macroeconomics and Reality”, *Econometrica 48*, pp. 1–48.

Woods, D., and A. Coleman. 2012. “Price, Quality, and International Agricultural Trade”, *Motu Working Paper* 12-08, Motu Economic and Public Policy Research, Wellington.

Wu, J., and M. Gopinath. 2008. “What Causes Spatial Variations in Economic Development in the United States?” *American Journal of Agricultural Economics* 90:2, pp. 392–408.

Table 1: Data Series

Series	Notation	Source
New Residential Building Consents	HC	SNZ
Housing Stock	H	Derived using SNZ data
World Commodity Price Index	$PCom$	Derived using QVNZ and ANZ data
National Farm Price per Hectare	PF	Derived from QVNZ Data
House Sales Price Index	PH	QVNZ
Commodity Price Implied Exchange Rate	XE	Derived using QVNZ and ANZ data

SNZ, QVNZ and ANZ refer to Statistics New Zealand, Quotable Value New Zealand and Australia and New Zealand Banking Group, respectively.

Table 2: National Series Stationarity Tests

	Lag Length	ADF
$\Delta lnPCom_t$	1	-4.7483
		(0.0001)
$\Delta lnXE_t$	1	-5.6874
		(0.0000)
$\Delta lnPH_t$	1	-3.6782
		(0.0044)
$HC_{it}/H_{i,t-1}$	1	-1.8981
		(0.3329)
$\Delta lnPF_t$	2	-7.5564
		(0.0000)

The ADF unit root test statistic is reported with its associated p-value in parentheses under a null hypothesis of non-stationarity. Lag length is chosen to minimise the Schwarz Bayesian Information Criterion.

Table 3: Panel Stationarity Tests

	LLC		IPS	
	Lags	Test Statistic	Lags	Test Statistic
$\Delta\widetilde{lnPCom}_{it}$	0	-34.4726	0	-35.5354
		(0.0000)		(0.0000)
$\Delta\widetilde{lnXE}_{it}$	0	-63.5636	0	-62.7698
		(0.0000)		(0.0000)
$\Delta\widetilde{lnPH}_{it}$	0	-79.9351	0	-80.8958
		(0.0000)		(0.0000)
$\widetilde{HC_{it}/H_{i,t-1}}$	0	-27.1502	0	-29.6338
		(0.0000)		(0.0000)

The Levin-Lin-Chu (LLC) and Im-Pesaran-Shin (IPS) panel unit root tests both have a null hypothesis of nonstationarity. The tests are performed on a cross-sectionally demeaned and Helmert-differenced transformation of the series, denoted by a tilde above the series name, across all TLAs. Test statistics are presented with associated p-values in parentheses.

Table 4: Autoregressive Order Selection, AIC and SBIC Statistics

TLA Sample	Lag Length	$\Delta\widetilde{lnPCom}_{it}$ AIC	$\Delta\widetilde{lnPCom}_{it}$ BIC	$\Delta\widetilde{lnXE}_{it}$ AIC	$\Delta\widetilde{lnXE}_{it}$ BIC	$\Delta\widetilde{lnPH}_{it}$ AIC	$\Delta\widetilde{lnPH}_{it}$ BIC	$\widetilde{HC_{it}/H_{i,t-1}}$ AIC	$\widetilde{HC_{it}/H_{i,t-1}}$ BIC
All	1	-29492	-29465	-38038	-38011	-20292	-20265	-59400	-59373
	2	-29171	-29118	-37615	-37562	-20106	-20053	-58985	-58932
	3	-29064	-28984	-37077	-36998	-19916	-19837	-58421	-58341
	4	-28740	-28634	-36642	-36536	-19745	-19639	-57734	-57628
	5	-28527	-28396	-36156	-36024	-19652	-19520	-56959	-56827
	6	-28167	-28010	-35758	-35600	-19431	-19273	-56173	-56015
	7	-27812	-27628	-35485	-35301	-19201	-19018	-55456	-55272
	8	-27428	-27218	-35001	-34791	-19034	-18825	-54708	-54499
Rural	1	-10367	-10344	-16085	-16062	-6830	-6807	-22949	-22926
	2	-10256	-10210	-15978	-15933	-6789	-6743	-22705	-22659
	3	-10235	-10167	-15763	-15695	-6721	-6653	-22482	-22414
	4	-10121	-10030	-15569	-15479	-6653	-6563	-22188	-22097
	5	-10067	-9954	-15396	-15283	-6636	-6523	-21911	-21798
	6	-9920	-9785	-15240	-15105	-6565	-6430	-21601	-21466
	7	-9815	-9658	-15259	-15102	-6478	-6320	-21297	-21139
	8	-9676	-9496	-15054	-14875	-6409	-6230	-21011	-20832
QRural	1	-10773	-10751	-12358	-12336	-6580	-6558	-18626	-18605
	2	-10671	-10627	-12208	-12165	-6543	-6499	-18522	-18479
	3	-10654	-10589	-12031	-11966	-6467	-6402	-18344	-18279
	4	-10647	-10561	-11913	-11827	-6430	-6343	-18162	-18075
	5	-10641	-10533	-11755	-11646	-6349	-6241	-17916	-17808
	6	-10614	-10484	-11646	-11516	-6258	-6128	-17686	-17556
	7	-10508	-10357	-11551	-11400	-6218	-6068	-17448	-17297
	8	-10344	-10172	-11382	-11210	-6166	-5994	-17229	-17058
Urban	1	-9482	-9461	-10499	-10477	-7977	-7955	-17982	-17960
	2	-9363	-9319	-10366	-10322	-7930	-7886	-17908	-17865
	3	-9282	-9216	-10212	-10147	-7859	-7794	-17716	-17651
	4	-9172	-9085	-10090	-10003	-7810	-7723	-17504	-17418
	5	-9068	-8960	-9947	-9839	-7763	-7655	-17255	-17147
	6	-8943	-8814	-9827	-9697	-7663	-7534	-17002	-16873
	7	-8817	-8666	-9723	-9572	-7570	-7419	-16817	-16666
	8	-8759	-8588	-9595	-9423	-7491	-7319	-16573	-16401

Table 5: Categorisation of New Zealand Territorial Local Authorities (TLAs)

Sub-sample	TLA Name	Land Value Proportion in Commodities	TLA Name	Land Value Proportion in Commodities
Urban	North Shore City	0.0012	Wellington City	0.0015
	Lower Hutt City	0.0017	Auckland City	0.0018
	Christchurch City	0.0106	Tauranga District	0.0114
	Upper Hutt City	0.0117	Nelson City	0.0155
	Porirua City	0.0167	Hamilton City	0.0169
	Waitakere City	0.0221	Palmerston North City	0.0383
	Manukau City	0.0463	Napier City	0.0491
	Kawerau District	0.056	Kapiti Coast District	0.0569
	Invercargill City	0.0579	Papakura District	0.0633
	Thames-Coromandel District	0.0755	Dunedin City	0.0766
	Queenstown-Lakes District	0.0769	Rodney District	0.1341
	New Zealand	0.1593		
Quasi-Rural	Taupo District	0.2045	Whangarei District	0.2088
	Rotorua District	0.2265	Tasman District	0.2428
	Marlborough District	0.2509	Wanganui District	0.2871
	New Plymouth District	0.2889	Grey District	0.2929
	Buller District	0.2942	Kaikoura District	0.2947
	Timaru District	0.3421	Far North District	0.3444
	Masterton District	0.3846	Westland District	0.3864
	Hastings District	0.4014	Western Bay of Plenty District	0.405
	Central Otago District	0.4051	Franklin District	0.4071
	Gisborne District	0.4337	Waimakariri District	0.4348
	Horowhenua District	0.4397		
Rural	Opotiki District	0.4434	Whakatane District	0.472
	Mackenzie District	0.4759	Selwyn District	0.4815
	Waipa District	0.5446	Wairoa District	0.546
	Waitaki District	0.5581	Kaipara District	0.559
	Waikato District	0.5613	Ruapehu District	0.6116
	Hurunui District	0.6201	Ashburton District	0.6249
	Hauraki District	0.6259	South Wairarapa District	0.6302
	Manawatu District	0.6516	Gore District	0.6578
	Carterton District	0.6594	Central Hawke's Bay District	0.7027
	Waitomo District	0.7049	Matamata-Piako District	0.7116
	Southland District	0.718	Clutha District	0.7378
	Stratford District	0.7758	South Waikato District	0.783
	Rangitikei District	0.7893	Otorohonga District	0.8007
	Tararua District	0.801	South Taranaki District	0.832
	Waimate District	0.8454		

Table 6: Long Run Elasticities across Differing Samples

Sample	$\varepsilon_{PH,PCom}$	$\varepsilon_{H,PCom}$
National	-2.8276	-0.2250
All	0.0867	0.0123
Rural	0.0360	0.0021
QRural	-0.0679	-0.0076
Urban	0.4542	0.0794

$\varepsilon_{PH,PCom}$ ($\varepsilon_{H,PCom}$) denotes the elasticity of house prices (the housing stock) with respect to commodity prices, where the housing stock changes are approximated through the cumulated housing investment rate responses.

Figure 1: Component Commodity Price Indices

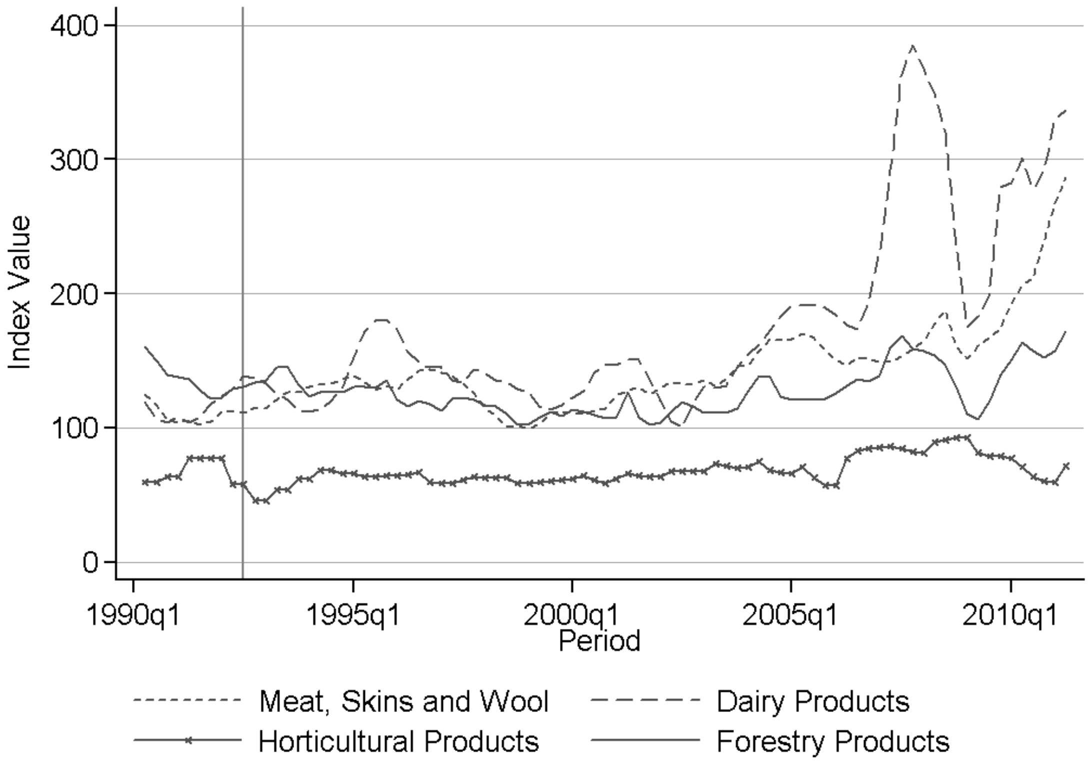

Figure 2: Housing Consents Issued vs. Population Change, by TLA

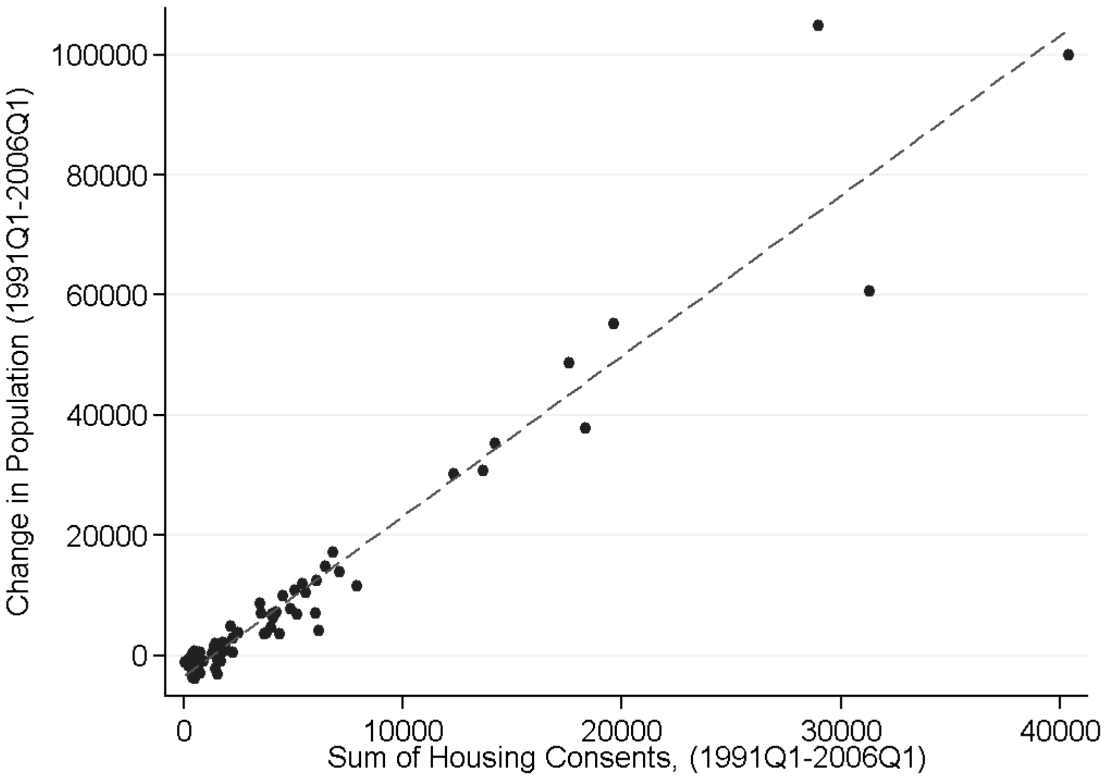

Figure 3: Alternative Exchange Rate Comovements

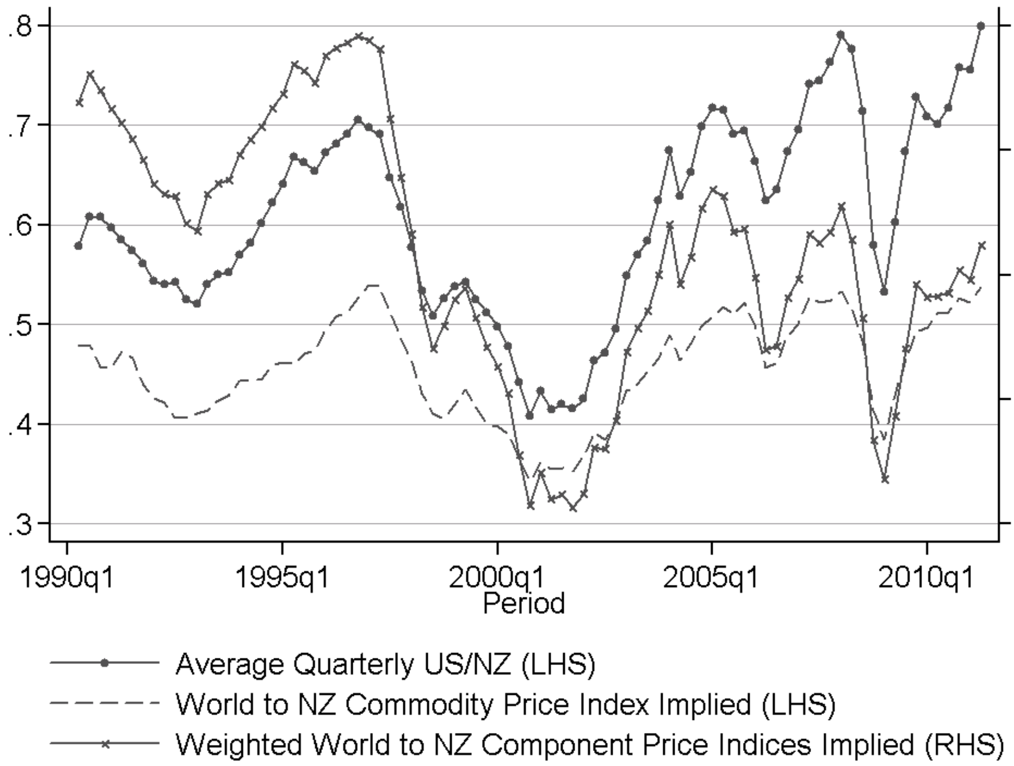

Figure 4: National Commodity Price Index and Derived Exchange Rate Comovements

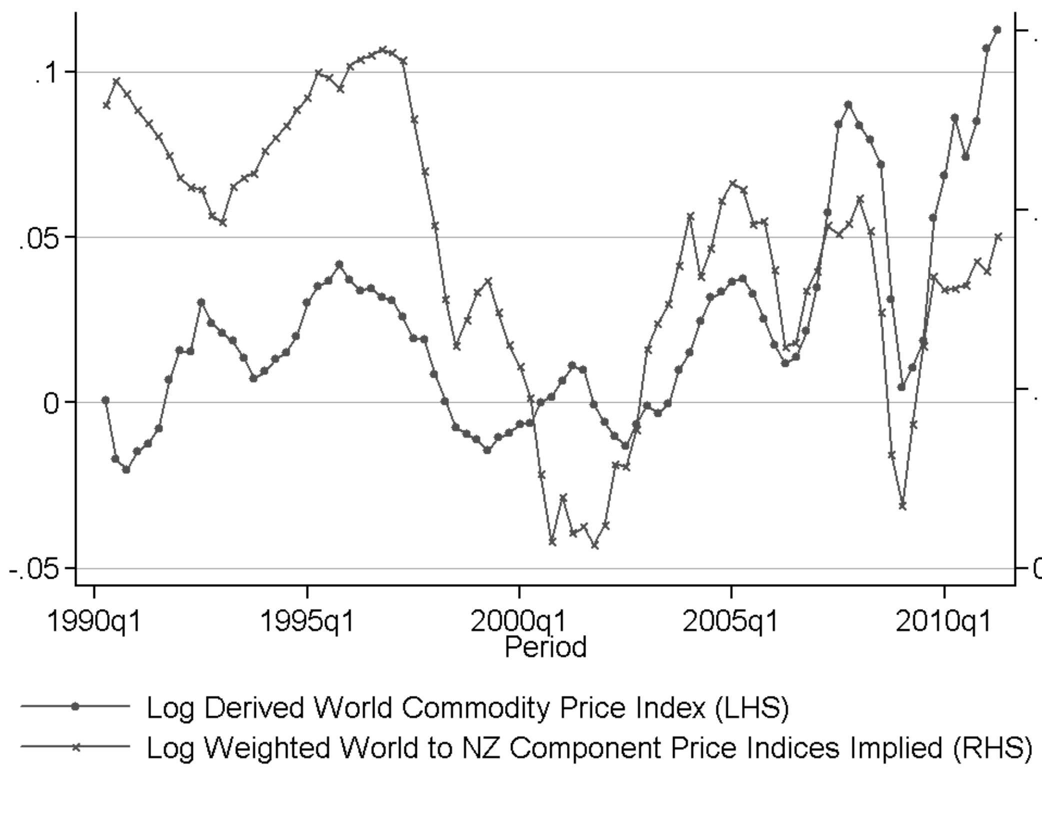

Figure 5: Comovements in National Commodity Price and House Price Indices

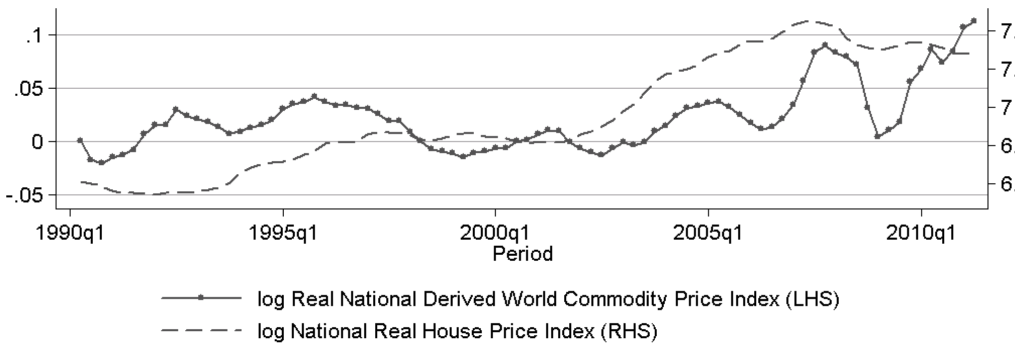

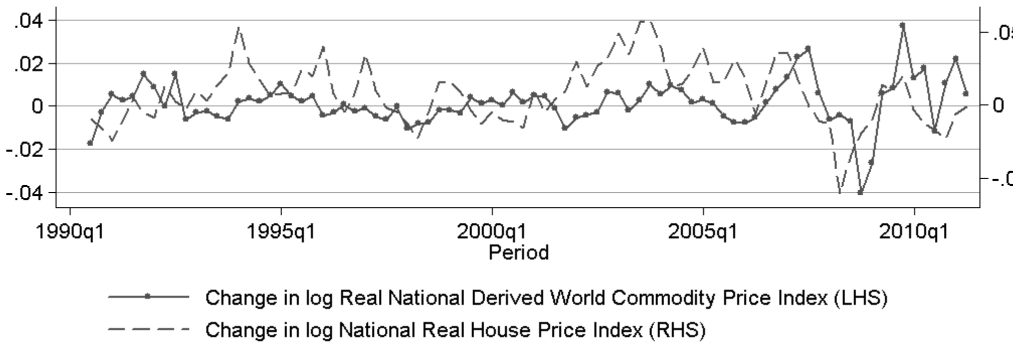

Figure 6: Comovements in National Commodity Price Index and Housing Investment Rate

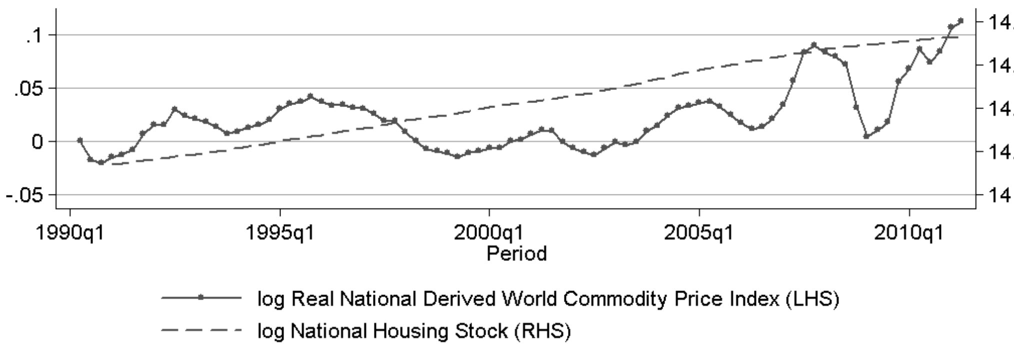

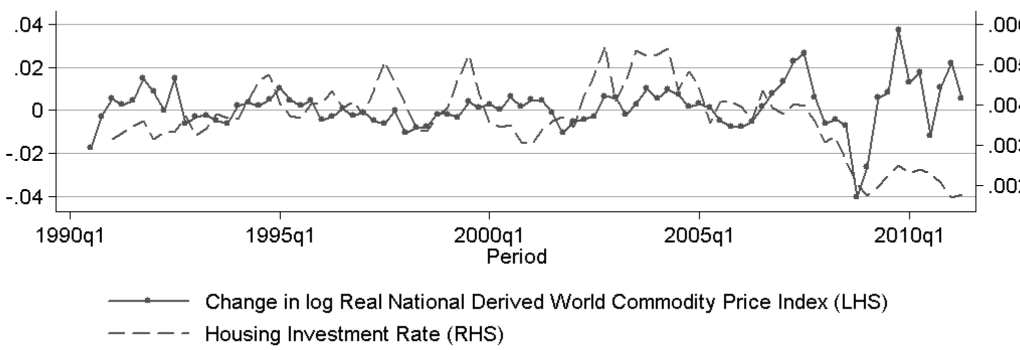

Figure 7: Comovements in National House Prices and Housing Stock

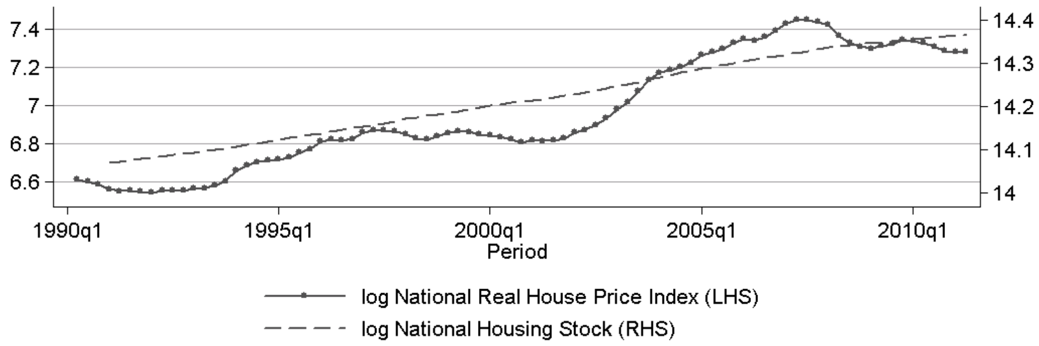

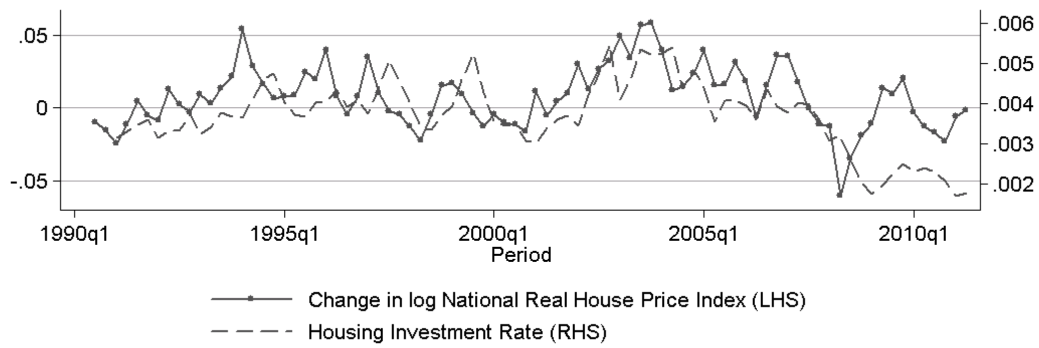

Figure 8: Comovements in National Commodity Price and Farm Price

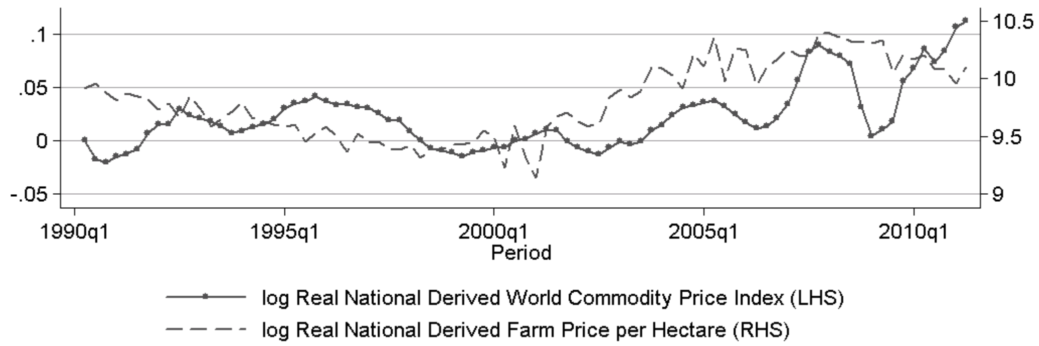

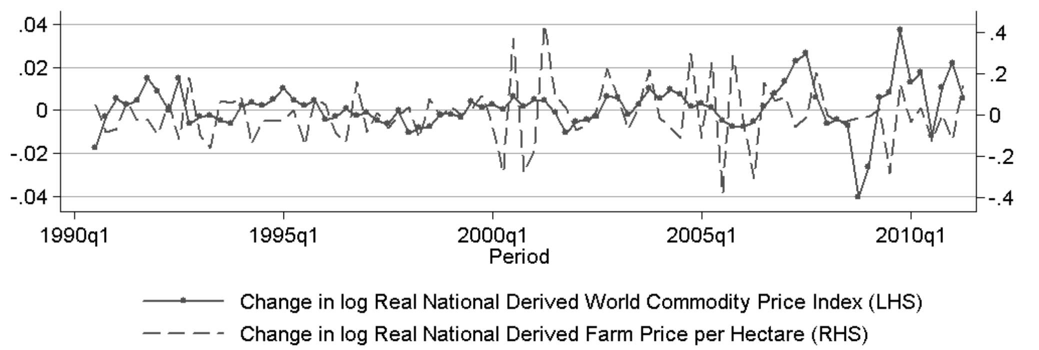

Figure 9: Result of a shock to the change in log commodity prices on housing outcomes, National.

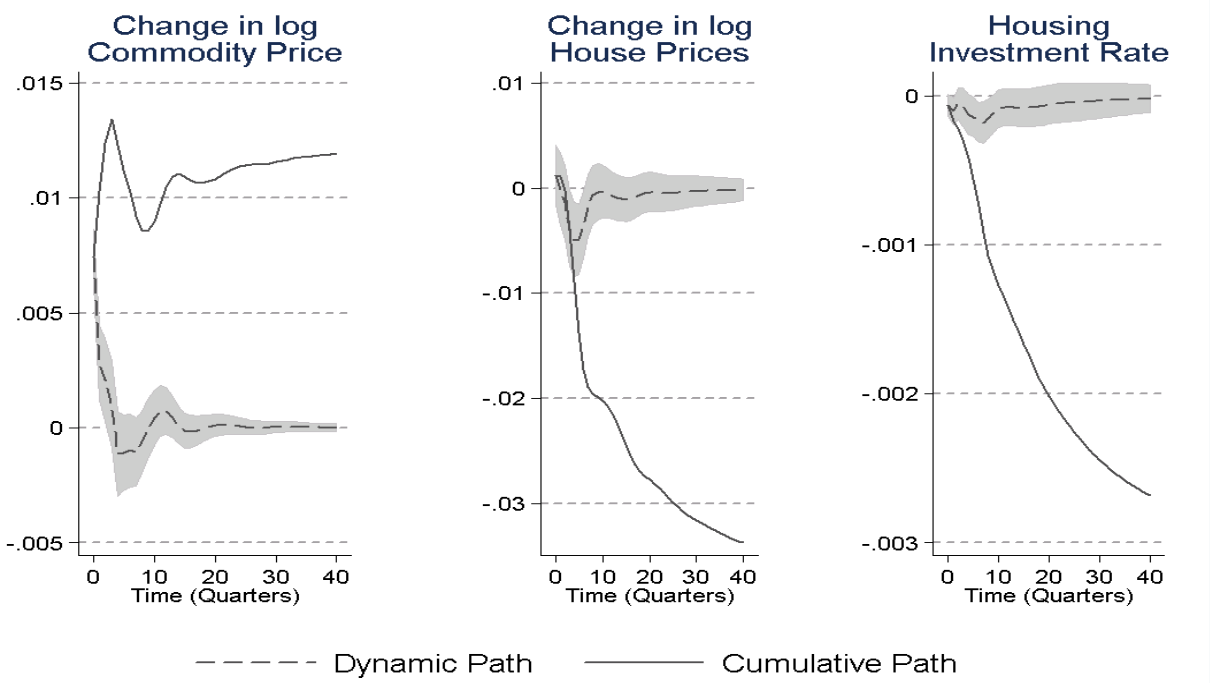

In

In this and all subsequent figures, the shaded region represents a 95% confidence interval around the dynamic path, whilst the cumulative path is the sum of all preceding dynamic adjustments.

Figure 10: Result of a shock to the change in log commodity prices on farm prices, National.

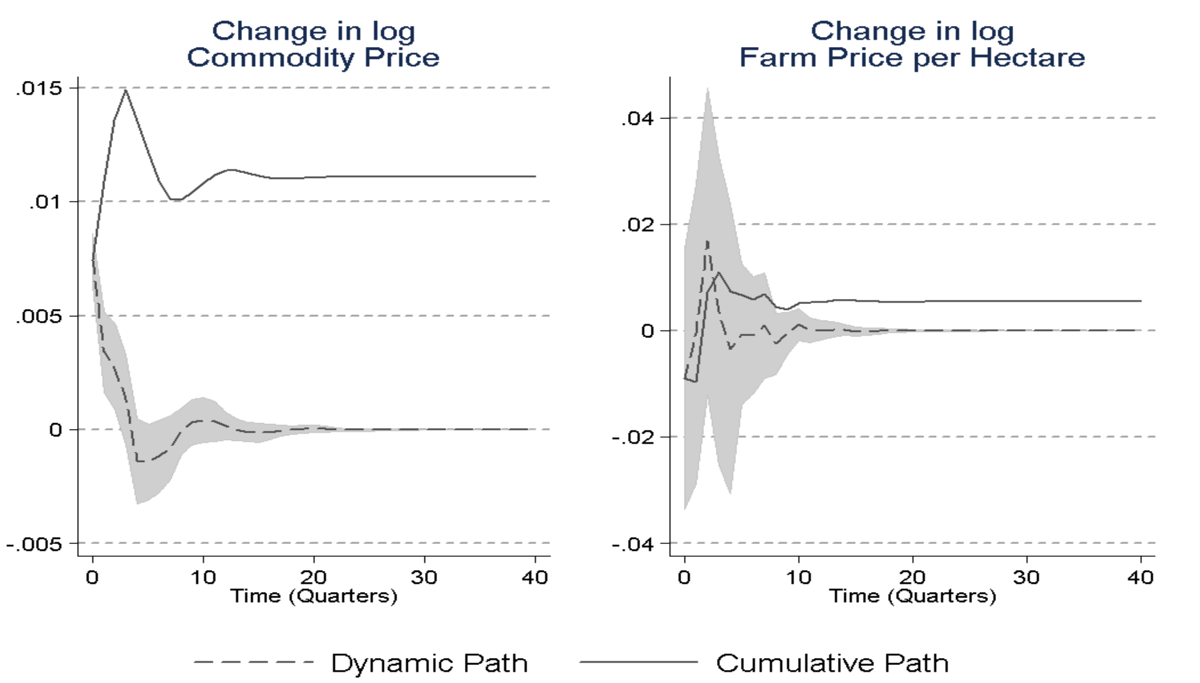

Figure 11: Result of a shock to the change in log Commodity Prices, All TLAs.

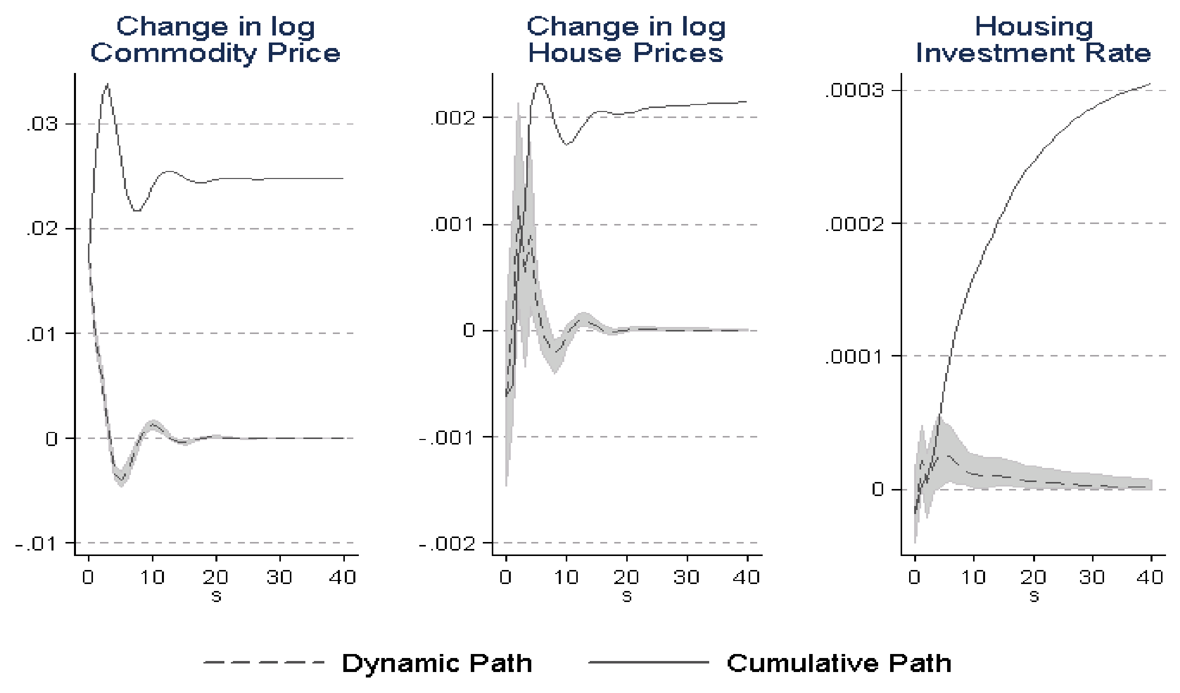

Figure 12: Result of a shock to the change in log Commodity Prices, Rural TLAs.

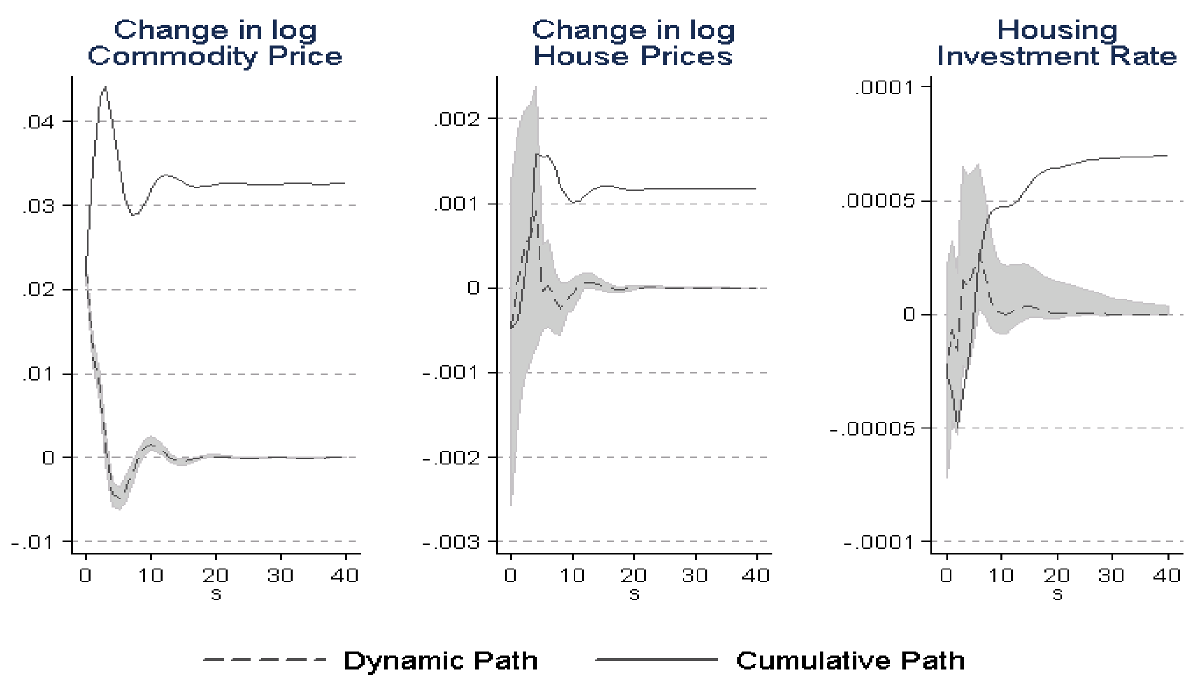

Figure 13: Result of a shock to the change in log Commodity Prices, Quasi-Rural TLAs.

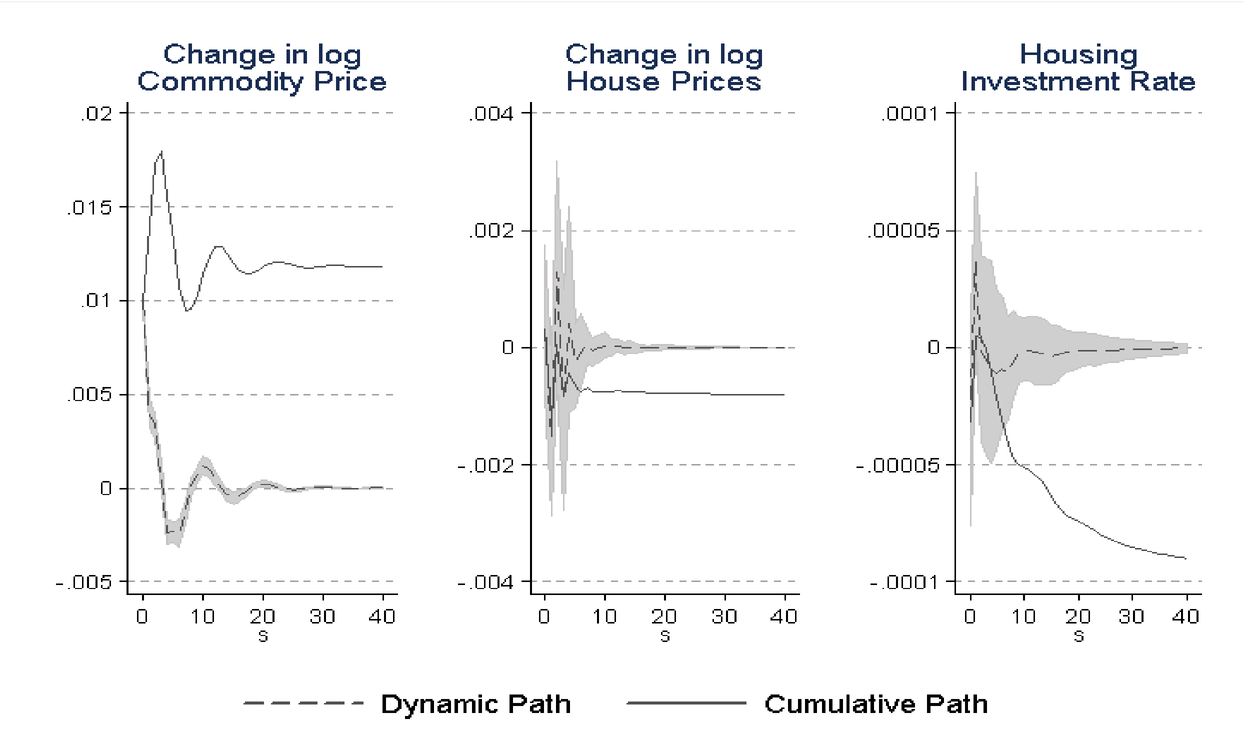

Figure 14: Result of a shock to the change in log Commodity Prices, Urban TLAs.

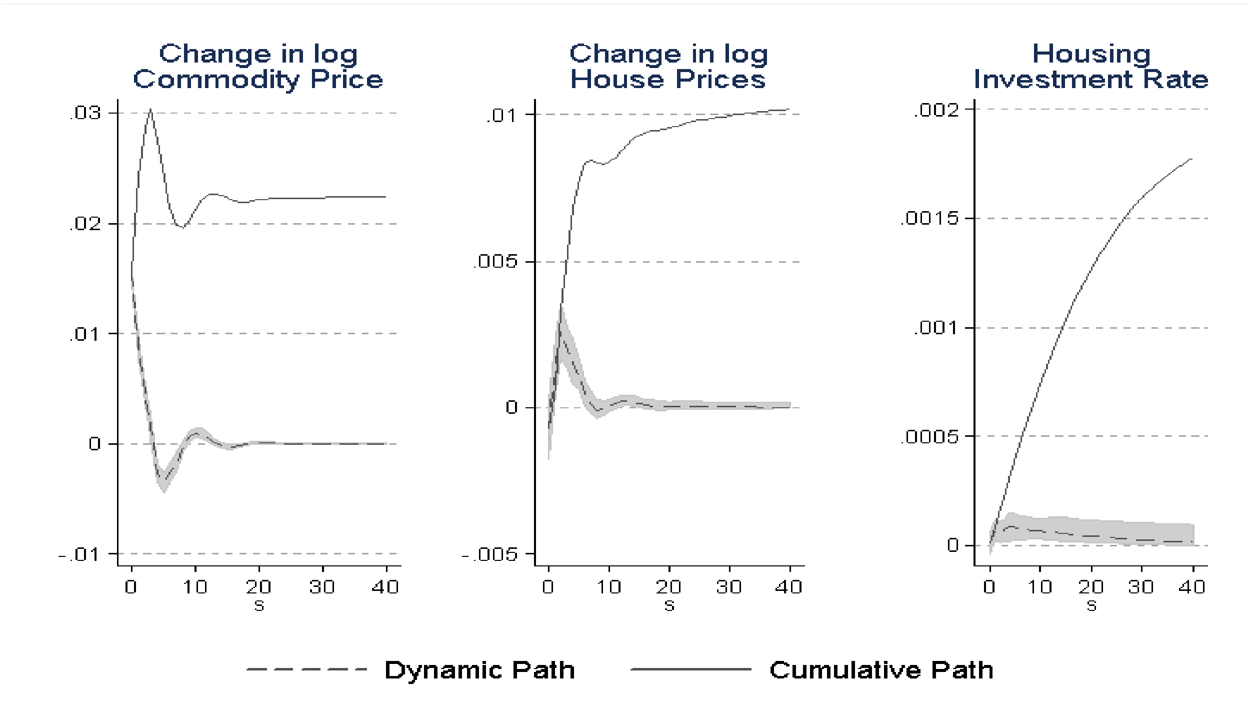

Appendix: Data Series

Quotable Value New Zealand (QVNZ) compiles a quality-adjusted house price index for each TLA and New Zealand as a whole, all of which we deflate by the CPI. This comprises our first measure of local (or national) outcomes, as house prices capitalise the productivity and amenity values of the surrounding locality. The second housing market outcome concerns the growth of the housing stock, which reflects employment opportunities and location preferences. To approximate the quarterly housing stock, a series which is not officially available at a subnational level, we adjust the five-yearly censal final count of occupied private dwellings from Statistics New Zealand (SNZ) by the number of new residential building consents issued by each TLA within a quarter (also from SNZ) as in Grimes and Hyland (2013). The second outcome is then defined as the ratio of housing consents in a quarter, relative to the previous quarters housing stock; a measure we term the housing investment rate. The building of a new dwelling requires obtaining a building consent from the TLA, in accordance with the Building Act 2004. Consequently, we can be confident in the dynamics of the numerator of this measure.

For a series of real commodity price movements we draw upon the ANZ Commodity Price Index dataset, which is available monthly since January 1986. The dataset contains an aggregate index, as well as indices for the component categories: meat, skins and wool; dairy: horticultural products; forestry products; seafood; and aluminium. The aggregate series are expressed both in New Zealand and US dollars, whilst the component series are expressed either in New Zealand dollars or as an internationally-denominated series which is calculated from prices in 5 different currencies depending on the locality of each commodity's market. We construct a local world-price denominated commodity price index by weighting the component indices by the relative value of that commodity to a district. We multiply the remaining share of value by the domestic Consumer Price Index (CPI) since 1991Q1, and then deflate the aggregate weighted index by the CPI. Given that the component indices are defined in various currencies, a weighted average of the ratio of world to domestic component price series provides our district-specific commodity-weighted exchange rate.

QVNZ also collects the sales price and corresponding land area for farms disaggregated across TLAs, with a description of primary farm activity, as well as the total land value in a TLA attributed to a wide range of activities. We use the national average sales price per hectare, to control for the effect of changing lot sizes on the sales price, across each ANZ Commodity Price Index component category. The weighted average of the average sales price per hectare across commodity components (deflated by the CPI) represents our farm price measure.

To construct the weights used in the three previous series we obtain land value data from QVNZ, which is the result of (usually) three-yearly valuations of all properties for the purposes of setting rates (local property taxes). The value of land attributed to production of a specific commodity in a TLA is used as a proportion, relative to total TLA land value, to derive weights that explain the intensity to which a community is attached to different commodity markets.

Thus, the area-specific commodity price index, exchange rate and sales price per hectare are defined as follows:

$$PCom_{it} = \sum_{j \in ANZ} \left(\frac{LV_{ijt}}{LV_{it}} . \frac{ANZ_{jt}^{W}}{CPI_t} \right) + \left(1 - \sum_{j \in ANZ} \frac{LV_{ijt}}{LV_{it}} \right) \left(. \frac{CPI_t}{CPI_t} \right) \qquad (A1)$$

$$XE_{it} = \sum_{j \in ANZ} \left(\frac{LV_{ijt}}{\sum_{j \in ANZ} LV_{ijt}} . \frac{ANZ_{jt}^{W}}{ANZ_{jt}^{NZ}} \right) \qquad (A2)$$

$$PF_{it} = \sum_{j \in ANZ} \left(\frac{LV_{ijt}}{\sum_{j \in ANZ} LV_{ijt}} . \frac{FSP_{jt}/FSA_{jt}}{CPI_t} \right) \qquad (A3)$$

where ANZ represents the set of commodities in the ANZ Commodity Price Index, LV_{ijt} is the land value attributed to the production of commodity j in TLA i, period t, LV_{it} is the total land value across all activities in TLA i, period t, CPI_t denotes the New Zealand Consumer Price Index, ANZ_{jt}^{W} and ANZ_{jt}^{NZ} denote the ANZ Commodity Price Index for the j^{th} component in world and New Zealand denominated prices respectively, whilst FSP_{jt} and FSA_{jt} denotes the sum of farm sales prices and area in hectares, respectively, for all farms sold in period t, engaged in production of commodity j.

To consider the robustness of our urban results we construct an alternative urban commodity price index which reflects the proportion of an urban TLAs land value in commodity production, and attributes this according to the composition of commodity production in the encompassing RC. Thus the alternative index is defined as follows:

$$\begin{aligned} \widehat{PCom}_{it} = & \left(\sum_{k \in ANZ} \frac{LV_{ikt}}{LV_{it}} \right) \sum_{j \in ANZ} \left(\frac{\sum_{l \in RC_i} LV_{ljt}}{\sum_{l \in RC_i} \sum_{k \in ANZ} LV_{ikt}} \right) \frac{ANZ_{jt}^{W}}{CPI_t} \\ & + \left(1 - \sum_{k \in ANZ} \frac{LV_{ikt}}{LV_{it}} \right) \left(. \frac{CPI_t}{CPI_t} \right) \end{aligned} \qquad (A4)$$

where RC_i denotes the set of TLAs within the same RC as TLA i and all other series and indices are defined as above. Importantly, none of our urban TLAs extend across RC boundaries (see Grimes, Maré and Morten, 2009), which simplifies the construction of this index.

Recent Motu Working Papers

All papers in the Motu Working Paper Series are available on our website www.motu.org.nz, or by contacting us on info@motu.org.nz or +64 4 939 4250.

13-09 Allan, Corey, Arthur Grimes and Suzi Kerr. 2013. "Value and Culture."

13-08 Maré, David C., and Richard Fabling. 2013. "The Incidence and Persistence of Cyclical Job Loss in New Zealand".

13-07 Grimes, Arthur, and Nicholas Tarrant. 2013. "A New Zealand Urban Population Database".

13-06 Fabling, Richard, and David C. Maré. 2013. "Firm-Level Hiring Difficulties: Persistence, Business Cycle and Local Labour Market Influences".

13-05 Crichton, Sarah, and David C. Maré. 2013. The Impact of Wage Subsidies on Jobseekers' Outcomes and Firm Employment".

13-04 Crawford, Ron, and David C. Maré. 2013. "Investigation of Options for a New Longitudinal Household Survey: Issues and Options Paper".

13-03 Dixon, Sylvia, and David C. Maré. 2013. "The Costs of Involuntary Job Loss: Impacts on Workers' Employment and Earnings".

13-02 Grimes, Arthur, and Sean Hyland, with Andrew Coleman, James Kerr and Alex Collier. 2013. "A New Zealand Regional Housing Model".

13-01 Fabling, Richard, and Lynda Sanderson. 2013. "Export Performance, Invoice Currency, and Heterogeneous Exchange Rate Pass-Through".

12-14 Motu Economic and Public Policy Research. 2012. "Roadmap for Implementing a Greenhouse Gas Emissions Trading System in Chile: Core Design Options and Policy Decision-Making Considerations".

12-13 Fabling, Richard, Arthur Grimes and David C. Maré. 2012. "Performance Pay Systems and the Gender Wage Gap."

12-12 Kerr, Suzi. 2012. "The Economics of International Policy Agreements to Reduce Emissions from Deforestation and Degradation."

12-11 Coleman, Andrew. 2012. "Pension Payments and Receipts by New Zealand Birth Cohorts, 1916–1986."

12-10 Tímár, Levente. 2012. "Spatial and Temporal Responses to an Emissions Trading System Covering Agriculture and Forestry: Simulation Results from New Zealand."

12-09 Grimes, Arthur, Les Oxley and Nicholas Tarrant. 2012. "Does Money Buy Me Love? Testing Alternative Measures of National Wellbeing."

12-08 Woods, Darian, with Andrew Coleman. 2012. "Price, Quality, and International Agricultural Trade."

12-07 Olssen, Alex, Wei Zhang, David Evison, and Suzi Kerr. 2012. "A Forest-Profit Expectations Dataset for New Zealand, 1990–2008".

12-06 Kerr, Suzi, and Alex Olssen. 2012. "Gradual Land-Use Change in New Zealand: Results from a Dynamic Econometric Model".

12-05 Abramitzky, Ran, and Isabelle Sin. 2012. "Book Translations as Idea Flows: The Effects of the Collapse of Communism on the Diffusion of Knowledge".

12-04 Fabling, Richard, and David C. Maré. 2012. "Cyclical Labour Market Adjustment in New Zealand: The Response of Firms to the Global Financial Crisis and its Implications for Workers".

12-03 Kerr, Suzi, and Adam Millard-Ball. 2012. "Cooperation to Reduce Developing Country Emissions".

12-02 Stillman, Steven, Trinh Le, John Gibson, Dean Hyslop and David C. Maré. 2012. "The Relationship between Individual Labour Market Outcomes, Household Income and Expenditure, and Inequality and Poverty in New Zealand from 1983 to 2003".