RESERVE BANK
OF NEW ZEALAND

DP2008/05

How do Housing Wealth, Financial Wealth and Consumption Interact?

Evidence from New Zealand

Emmanuel De Veirman
and Ashley Dunstan

February 2008

JEL classification: C22, C32, E21

Discussion Paper Series

DP2008/05

How do housing wealth, financial wealth and consumption interact?

Evidence from New Zealand[+]

Emmanuel De Veirman and Ashley Dunstan[*]

Abstract

This paper characterises the relationship between wealth and consumption in New Zealand. We find that there exists a long-run cointegration relation between household consumption, income, housing wealth and net financial wealth. Permanent shocks account for most of the variation in wealth. This implies that our cointegration estimates accurately capture the effect of most wealth changes, in contrast with the findings of Lettau and Ludvigson (2004) for the United States. Our estimates suggest that consumption has adjusted sluggishly to restore long-run equilibrium, but also that consumption booms have anticipated equilibrium-restoring increases in housing wealth. Furthermore, we estimate two alternative econometric models which are more robust to instability in the long-run relationship. All three of our models suggest that permanent changes in wealth have economically important effects on consumption. The dollar-for dollar-effect of financial wealth exceeds that of housing wealth.

JEL Classification: C22, C32, E21
Keywords: consumption wealth effects, cointegration, error-correction.

[+] We thank Phil Briggs, Andrew Coleman, Felix Delbruck, Aaron Drew, Khoon Goh, Arthur Grimes, David Hargreaves, Cath Sleeman, as well as participants of the December 2007 RBNZ and New Zealand Treasury workshop on "The business cycle, housing and the role of policy" for helpful comments and suggestions. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Reserve Bank of New Zealand.

[*] Address: Economics Department, Reserve Bank of New Zealand, 2 The Terrace, PO Box 2498, Wellington, New Zealand.

1. Introduction

In real terms, the net worth of the average New Zealander has doubled since 2001. This increase in wealth is mostly due to New Zealand's most recent housing boom: annual house price inflation was twelve percent on average in the period 2001-2006, compared with an average of four percent in the period from 1991-2000. Associated with this boom has been a rapid increase in outstanding household liabilities and substantial declines in income-based measures of household savings rates. The Reserve Bank of New Zealand, amongst others, has closely monitored these developments. From a monetary policy perspective, the increase in housing wealth is thought to have been a key driver of consumer spending, aggregate demand and consumer price inflation over the recent business cycle. From a financial stability perspective, the rise in the value of New Zealanders' housing stock has been associated with an increased degree of household leverage, which potentially elevates financial stability risk.

In this paper we examine the relationship between household net worth and consumption in New Zealand using a range of empirical models. Broadly in line with previous New Zealand research on wealth and consumption, our key finding is that changes in housing wealth have a large impact on consumption spending relative to international studies. We attribute the larger housing wealth elasticity in New Zealand to the fact that housing assets tend to be a larger fraction of households' portfolios in New Zealand, rather than their being a larger marginal propensity to consume out of housing wealth per-se.

Internationally, the empirical literature is split between two different views on how household wealth relates to consumer spending. The traditional view implies that there is a short-run response of consumption to changes in wealth, but that wealth itself does not depend on movements in consumption. This strand of the literature implicitly assumes that any change in wealth is due to exogenous shocks which are 'permanent' in the sense that their effect does not wither away over time. Under this assumption, standard cointegration estimates accurately capture the long-run effect of any change in wealth on consumption.

More recent empirical work, starting with Lettau and Ludvigson (2004), suggests that changes in consumption tend to anticipate changes in wealth. In a sense, this implies that wealth is endogenous to (or at least can be predicted by) past changes in consumption. Lettau and Ludvigson found that in the United States, only a tiny fraction of the variation in wealth is due to permanent shocks. Since estimates of a long-run cointegration relationship between consumption and wealth only measure the effect of permanent changes in wealth on consumption, their findings suggest that standard cointegration estimates are irrelevant for the effect of most changes in wealth.

In the baseline, we implement the systems-based framework developed by Lettau and Ludvigson (2004) to estimate the long-run cointegrating relationship between consumption and wealth, as well as the implied short-run dynamics. However, we extend the Lettau-Ludvigson model by splitting household wealth into housing and net financial wealth to investigate the separate effects of these different wealth components. Investigating separate effects is particularly important for New Zealand: real per-capita housing wealth has tripled in the two decades since the 1987 stock market crash, while the total value of financial assets has remained relatively stagnant. An additional advantage of splitting household net worth into housing wealth and financial wealth net of household liabilities is that it allows us to examine the influence of house price fluctuations on household borrowing.

Our results suggest that in New Zealand, wealth and consumption co-moved in a way which is mostly in line with the traditional view that wealth changes are permanent. This implies observed changes in wealth tend to cause permanent changes in consumption levels. However, as transitory shocks account for a large part of the variation in consumption the adjustment of consumption to changes in wealth is fairly sluggish. These results differ from Lettau and Ludvigson and plausibly reflect the fact that the composition of household net worth in New Zealand differs from that in the United States. For instance, Lettau and Ludvigson pointed to stock market cycles to explain their finding that most changes in wealth are transitory. Yet in New Zealand, direct equity constitutes only a small fraction of financial wealth.

However, some of our evidence supports the hypothesis that wealth in New Zealand has varied in accordance with the Lettau-Ludvigson view. Although permanent shocks explain most of the variation in housing wealth, there is still substantial transitory variation in housing wealth. Moreover, these transitory changes in housing wealth have tended to restore the long-run equilibrium relationship. In particular, periods in which consumption is high relative to current wealth have tended to anticipate future increases in housing wealth. This is in line with the interpretation that households' expectations of future housing wealth increases tend to imply periods in which consumption is high relatively to contemporaneous wealth.

We find some evidence for instability in the point estimates for the long-run elasticities of consumption to changes in wealth and income. Among other factors, we attribute this to instability in the shares of housing and financial wealth in New Zealanders' overall resources. In part because of this instability, we implement two alternative econometric techniques which are more robust to instability in the long-run relationship. The first is inspired by Aron and Muellbauer (2007), and the second is a variant of Carroll, Otsuka, Slacalek (2006). Both techniques involve single-equation models of consumption which impose the assumption that all movements in wealth are permanent.

Both alternative techniques suggest that the dollar-for-dollar effect of net financial wealth is larger than that of housing wealth. However, housing wealth constitutes a much larger share of total household wealth, which implies a large effect from any one-percent increase in housing wealth on consumption. Finally, our results suggest that increases in housing wealth tend to imply decreases in net financial wealth. Among other interpretations, this is in line with a collateral effect by virtue of which increased house values translate into increased household borrowing.

Our paper is structured as follows. In section 2, we describe mechanisms through which movements in wealth and its components can affect consumer spending, and discuss why the strength of the wealth-consumption linkage is likely to vary over time. In section 3, we discuss our data choices. After summarising the theoretical underpinnings of our estimation framework, section 4 discusses our estimates of the long-run relationship between consumption, income and the components of wealth and the short-run dynamics between these variables. Section 5 discusses and implements two alternative econometric techniques which are more instructive for learning about the relative strength of the dollar-for-dollar effects from financial wealth and housing wealth. Section 6 concludes and offers suggestions for future research.

2. Background

2.1 How do changes in wealth cause changes in consumption?

Studying the effect of changes in household wealth on consumption has a long history in economics, dating back to the early work on the lifecycle model by Ando and Modigliani (1963) and Modigliani and Brumberg (1954). In a typical lifecycle model, any household's consumption responds to changes in its permanent income, which is the annuity value of expected lifetime resources. These resources are made up of two elements: physical wealth (current housing and financial wealth), and human wealth (current labour income plus the present discounted value of expected future labour income).

In such models, changes in wealth affect consumption through the wealth effect: any increase in wealth will increase consumption through its impact on expected lifetime resources. When the value of these resources increases, the household can shift its consumption schedule upward without violating its budget constraint. In the standard lifecycle model, the household increases spending in every remaining period of its lifetime by a constant equal to the increase in permanent income. As such, the change in wealth has both a short-run and a long-run effect on consumption. In the standard lifecycle model, this effect is the same regardless of which component of lifetime resources actually increased.[2]

A number of features of reality are not captured by standard lifecycle/permanent income models. Accounting for some of these features uncovers other potential causal linkages between consumption and wealth.

The first of these features is asymmetric information in credit markets, which leads to the so-called collateral effect on consumption. Imperfect information about borrower characteristics implies that banks typically require borrowers to provide collateral as an insurance against default risk. When a household experiences an increase in wealth, the value of the collateral it can offer will also increase. This generally means that banks will then be willing to increase the supply of loans to such household. If the household was previously credit-constrained (i.e. it could not borrow enough to finance its desired consumption level), then it is likely to actually borrow more in order to finance extra consumption.[3] In New Zealand, as in other countries, almost all household debt is secured against housing wealth, such that the collateral effect applies almost exclusively to changes in housing wealth.

The lifecycle/permanent income models of consumption also do not encapsulate a meaningful role for uncertainty. In the face of uncertainty about future income and asset values, households may choose to hold a buffer stock of wealth, in order to mitigate the downward impact on consumption of low-probability negative income shocks. This is known as the precautionary motive for saving, and acts much like a self-imposed credit constraint. An increase in a household's wealth will increase the value of its buffer stock, and thereby relax the need for precautionary saving. This allows the household to increase spending without eating into its buffer stock of wealth.

---

[2] Some of the more specific implications of the lifecycle / permanent income hypothesis have tended to be rejected by the data. For example, versions of the model with rational expectations imply that predictable changes in wealth do not affect consumption at the time of the wealth change. Much of the subsequent consumption literature has focused on testing this implication of the lifecycle model, with little evidence in its favour. Muellbauer and Lattimore (1995) provide a good overview of this literature.

[3] This is not to imply that the collateral effect is the sole reason for a positive relation between house prices and liabilities. This relation is also likely to reflect the fact that, at times when house prices are high, households will need to borrow more in order to purchase a house.

In both the collateral effect and the precautionary saving effect, the increased liquidity associated with an increase in wealth allows households to consume more today, without necessarily implying an increase in their expected lifetime consumption. That is to say, these effects are associated with an inter-temporal redistribution of spending rather than with an upward shift in the entire consumption schedule. However, collateral and precautionary saving effects may affect consumption beyond the immediate short-run, and may therefore affect our estimates of the long-run effect of wealth changes on consumption.

2.2 Effect from housing wealth vs. financial wealth

In empirical terms, it is questionable whether changes in different components of wealth would have the same implications for spending. We follow a number of recent papers[4] in considering whether changes in housing wealth affect consumption differently than changes in financial wealth. We have already discussed one reason why the effects could differ: housing wealth changes are likely to have a much larger impact on households' marginal borrowing capacity through the collateral effect. Below, we discuss a number of other reasons why the spending implications from these two components of wealth are likely to differ.

One key difference between housing and financial wealth is that households do not merely use their homes to store wealth, but also obtain a service from them: the benefit of living in a house. Although homeowners experience a direct wealth gain when house prices rise, the implicit cost of consuming these housing services will rise as well. For homeowners intending to increase their consumption of housing services (e.g by moving into a more expensive home), or renters wanting to enter the housing market, the net effect on wealth will be negative. In contrast, the wealth gains are likely to offset the increase in the cost of consuming housing services for those households who intend to trade down in the housing market. The size and sign of the aggregate effect of changes in housing wealth on consumption therefore depends on the fraction of households in each of these categories, as well as on the relative size of their spending responses to changes in housing wealth. In contrast, these ambiguities do not apply to financial wealth.

Housing wealth is also arguably less liquid than financial wealth, since the transaction costs associated with trading up or down in the housing market are relatively high. A related point is that the bequest motive may be more important for housing wealth, implying that households are more reluctant to sell their house and are thus less likely to transform any house price increase into liquid assets ready for consumption. These factors tend to reduce the strength of the linkage between housing wealth and consumption. However, financial market liberalisation in New Zealand has made it much easier to obtain debt secured against housing wealth, which has increased the liquidity of housing wealth. (See the next subsection for more on financial market liberalisation.)

On the other hand, housing wealth may have a stronger impact on consumption. For instance, housing wealth tends to be more persistent than financial wealth, such that any change in a household's housing wealth will have a relatively large impact on the household's expected lifetime resources. Given that there are arguments from both sides, it is clear that the relative strength of the spending effects from financial and housing wealth needs to be resolved empirically.

---

[4] Recent influential examples include Case et al (2005) and Carroll et al (2006).

2.3 Time-varying strength of the consumption-wealth linkage

The strength of the relationship between changes in wealth and consumption is likely to change over time. Here, we consider a few key reasons why this could occur. Throughout the paper, the stability of the relationships we are estimating will remain a key focus.

Firstly, financial market liberalisation is likely to have altered the strength of the linkage between wealth and consumption.[5] We can interpret financial innovation as an increase in the availability of credit for given creditor characteristics (such as the creditor's current level of indebtedness). There are at least three ways in which financial innovation affects consumer spending. Firstly, financial liberalisation tends to increase the size of loans relative to (housing) wealth. This means that changes in house prices will lead to larger fluctuations in marginal borrowing capacity, and potentially consumption, through the collateral effect. Secondly, the increase in loan supply associated with financial innovation may reduce the fraction of credit-constrained households. Since the collateral effect is only in play for liquidity-constrained households, this tends to weaken the aggregate linkage of wealth with consumption. Thirdly, any increase in loan-to-value or loan-to-income ratios associated with financial innovation implies that households no longer need to save as much for their first home, which tends to increase the level of consumption in the short run.

Secondly, the strength of the linkage between housing wealth and debt may depend on the direction of change in house prices. Increases in housing wealth tend to allow households to increase the amount borrowed from the very moment of the housing wealth increase (subject to the loan-to-income constraint not binding). On the other hand, even if a decrease in housing wealth implies that some households' leverage ratios increase beyond the level which banks deem optimal, banks typically do not have the option of unilaterally withdrawing credit before the end of the loan contract.[6] This mutes the short-run effect of decreases in housing wealth on borrowing and consumption. Similarly, increases in house prices transmit to increase aggregate liabilities through passive equity withdrawal. In particular when house prices have been increasing, home buyers need to take on larger mortgages than the previous owners were holding. This effect also works in the opposite direction. But because turnover in the housing market is much higher during a boom period, this channel works faster when there have been increases in house prices than when there have been decreases.

2.4 Existing empirical evidence

A long-standing literature has investigated the marginal propensity to consume from overall wealth or, sometimes, from financial wealth. Paiella (2007) reviews recent estimates from this literature which place the long-run marginal propensity to consume (MPC) from wealth in the range of 3-8 cents and the long-run MPC from financial wealth in the broader range of 2-16 cents. More recent papers have also investigated the MPC from housing wealth. A widely cited study is Case *et al.* (2001), which concluded, using a panel of US states as well as a panel of OECD countries, that the MPC from housing wealth was around 4-9 cents, and substantially larger than the MPC from financial wealth. Carroll *et al.* (2006) obtain similar results using time series data for the US. However, Dvornak and Kohler (2003) find that, for a panel of Australian states, the MPC from stock market wealth is 2-3 times as large as the MPC from housing wealth, while Bertaut (2002) fails to find any consistent pattern across OECD countries.

---

[5] See Hull (2003) for a summary of changes in New Zealand and Alemeida *et al*. (2005) for international evidence on liquidity effects.

[6] This reasoning assumes that a non-negligible fraction of loans are written for a few years, which is the case in New Zealand. See Hull (2003) for a discussion of typical loan terms in New Zealand.

Several recent papers, stemming from the influential paper of Lettau and Ludvigson (2004), have questioned the relevance of such estimates of the long-run MPC from wealth. Lettau and Ludvigson (2004), Fernandez-Corugedo et al. (2003), Fisher *et al.* (2005), Chen (2006) and Pichette and Tremblay (2003) all find that a substantial component of wealth is transitory, for the US, UK, Australia, Sweden and Canada. The implication is that that previous studies overstated the long-run correlation between wealth and consumption. As yet, it is unclear which component of wealth drives these transitory movements, and it may be that the answer to this question varies from country to country. Lettau and Ludvigson (2004) and Pichette and Tremblay (2003) find that transitory movements in wealth are due to stock market wealth, whereas Fisher *et al*. (2005) and Chen (2006) both find instead that they are due to housing wealth.

In New Zealand, there is relatively little existing evidence on the linkage between wealth and consumption. Using annual data with a sample starting in 1978, Hull (2003) finds a strong long-run relationship between house prices and consumption. However, she did not include financial wealth in her regressions or examine the dynamic interaction between consumption and wealth. Downing and Goh (2002), using quarterly data from 1992-2002, investigated the effects of both financial and non-financial (primarily housing) wealth on consumption. They concluded that housing wealth has virtually no impact on consumption in the long run, while the effect from financial wealth is sizeable (but smaller than the estimate of Hull (2003) for housing wealth). However, they do find that housing wealth has an important short-run impact on consumption. In our paper, we implement a systems-based approach to investigating the consumption-wealth linkage, in which all variables are allowed to interact with each other and adjust to restore our estimated long-run equilibrium.

3. Data

Our dataset consists of real per-capita household consumption, wealth[7], and labour income[8] for the period 1982Q2-2006Q1. We discuss the most substantive dataset choices below, and refer to Appendix 1 for further details.

We focus on a measure of total household consumption that includes both non-durables and durables spending. Many papers in the literature have instead used a measure of non-durables consumer spending, on the grounds that households derive utility from durable goods over many periods. This means that standard macroeconomic models that assume instantaneous and time separable utility from consumption have little to say about this type of spending. Therefore, non-durables consumption is the appropriate measure for studies which focus on testing the specific implications of these models. In this paper, we are mainly interested in characterising how aggregate demand responds to wealth changes, which is why we include durable goods in our measure of consumer spending.

In all our data choices, we have ensured that the estimated equations are consistent with the aggregate implications of the household budget constraint identity. Rudd and Whelan (2006) show how failure to account for these considerations has led to misleading results in some recent US studies. For example, our consumption measure counts durable goods as consumption rather than wealth. Therefore, wealth should be measured net of any

[8] The nominal budget constraint implies that labour income – rather than a measure of disposable income that includes property income – is appropriate. In addition, income flows from the stock of wealth may partly be captured by the capital gain component of wealth (which enters in our measure of wealth) so that including property income would amount to double-counting.

accumulated durables goods. Another important choice is the procedure used to deflate the variables. The budget constraint is a relationship between nominal quantities, but in practice we are interested in relationships between real variables. To preserve the nominal relationships, it is important that we use the same index to deflate all variables. We use the CPI index for this purpose.

We decompose net worth into gross housing wealth and net financial wealth, defined as gross financial wealth minus total household debt. While the division into housing and financial wealth springs from our interest in estimating separate effects on consumption, there is little theoretical guidance on how to treat liabilities. On the one hand, it might make sense to subtract mortgage debt from housing wealth, given that this type of debt is secured on the housing stock. On the other hand, mortgage debt and financial wealth have some characteristics in common, such as the fact that neither is treated as a consumption good by households. Most importantly, if we were to define housing wealth net of housing debt we would not be able to explore interactions between housing wealth and liabilities, such as those implied by the collateral effect. As a robustness check, we also present some results in Appendix 2 that define housing wealth net of mortgage debt and define financial wealth net of non-mortgage liabilities.

Figures 1 and 2 graph the main series used in this paper. Figure 1 plots consumption and labour income. The figure shows that, since 1992, consumption growth has been remarkably strong. This trend has occurred in the context of a widening gap between consumption and labour income. Figure 2 shows household net worth, along with its components: gross housing wealth, gross financial wealth and liabilities. One clear implication of this graph is that the variations in household net worth have largely been driven by movements in housing wealth, at both a high and a low frequency. Another notable feature of the graph is the massive increase in housing wealth since 2001. This development has generated large increases in the share of housing wealth in household net worth. Meanwhile, increases in mortgage debt over this period led to a substantial erosion of net financial wealth.

Figure 1 Consumption and labour income

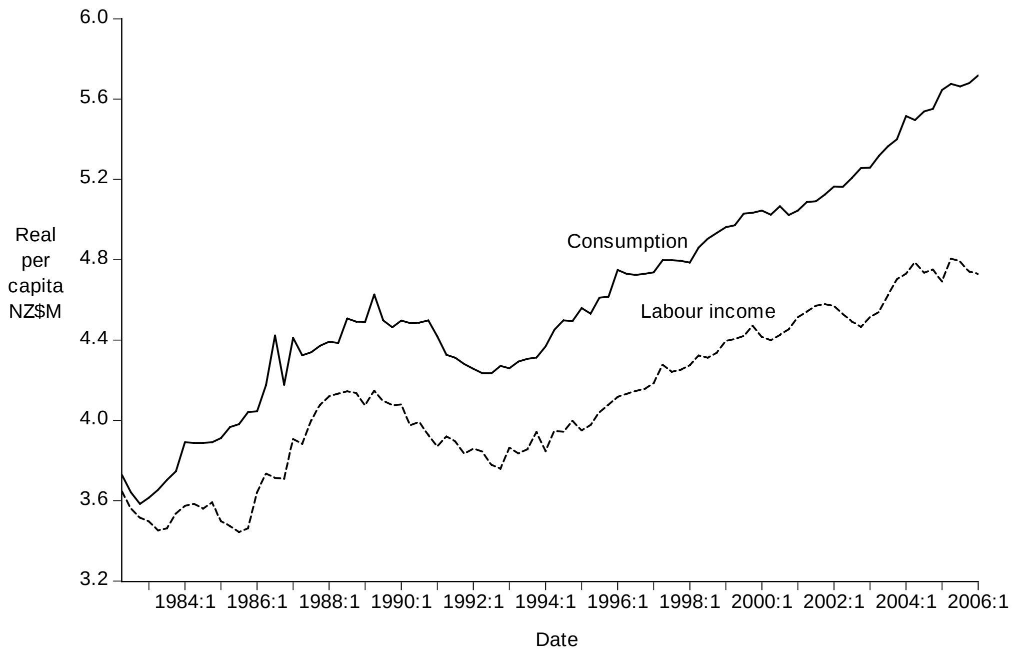

Figure 2 Components of household wealth

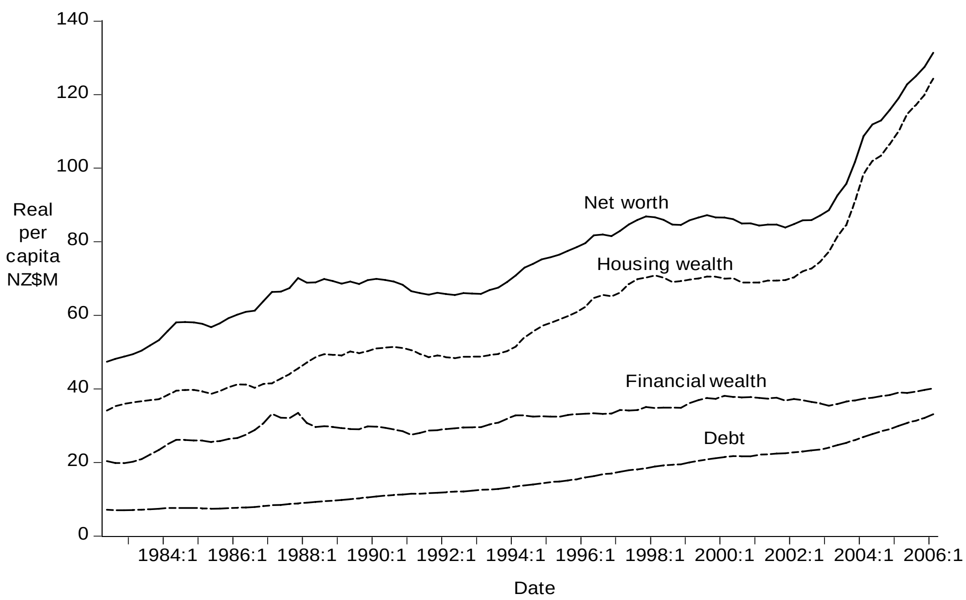

Table 1 presents key summary statistics for the quarterly growth rates in real per-capita consumption, income, housing wealth and net financial wealth. In New Zealand, housing wealth has roughly the same volatility as consumption and income. Net financial wealth is more volatile, but this is a relatively small component of overall household net worth. Both housing and financial wealth growth is far more persistent than growth in consumption or income. Table 1 also reveals some interesting correlations in the data. Housing wealth and income growth are positively correlated with consumption, although net financial wealth is not. There is also a strong negative correlation between housing wealth and net financial wealth. One of the main objectives of our analysis is to sort out the economic relationships which these correlations reflect.

Table 1 Summary statistics

	Mean (%)	Standard deviation	Serial correlation	Correlation matrix			
				Δc	Δh	Δnf	Δy
Consumption growth (Δc)	0.45	0.015	-0.04	1	0.21	-0.07	0.31
Housing wealth growth (Δh)	1.34	0.019	0.63		1	-0.28	0.14
Financial wealth growth (Δnf)	-0.59	0.040	0.38			1	0.10
Income growth (Δy)	0.27	0.014	0.00				1

Notes: For details on the series see Appendix 1. The table contains quarterly standard deviations. The degree of serial correlation is computed from an OLS regression of growth in each series against its lag and a constant. A dummy for GST changes is included in the consumption growth equation.

4. The long-run relationship between consumption and wealth and implied dynamics

We divide our results on the relationship between consumption and wealth into two separate sections. In the present section, we examine the long-run relationship between consumption, housing wealth, net financial wealth and income. We also investigate each of these variables' short-run response to a situation in which the system is not at its long-run equilibrium. However, there are two drawbacks to this particular methodology. Firstly, we find some evidence for instability in our estimated long-run relationship. Secondly, this formulation does not provide a suitable framework for examining the relative spending impact of an increase in housing wealth versus an increase in net financial wealth. For these reasons, section 5 considers two alternative techniques that are robust to these drawbacks, but have less to say about the dynamic interactions between the variables.

4.1 Estimation framework

The long-run relationship between consumption, income and wealth

In this section, we examine the foundations of the long-run relationship that we estimate, and how we can test for the existence of this long-run relationship. We also develop our framework for examining the short-run dynamic interactions between consumption, income and wealth, including the influence of the long-run relationship. Finally, we consider how to decompose each variable into a permanent and transitory component. This will turn out to be useful in interpreting our long-run regression estimates, allowing us to determine whether the long-run coefficient estimates provide an accurate reflection of the long-run correlation between consumption and wealth.

A useful framework to discuss our estimated long-run relationship is the household budget constraint identity. This allows us to motivate the particular relationship that we estimate without having to impose additional preference structure from any particular theory of consumer behaviour, in line with the argument of Lettau and Ludvigson (2001). Following the seminal work of Campbell and Mankiw (1989), we begin with a version of the inter-temporal budget constraint identity for an infinitely lived representative agent. We define wealth in a broad sense so that it includes human wealth:

$$W_{t+1} = (1 + R_{w,t+1})(W_t - C_t) \tag{1}$$

In words, any accumulated wealth (W) that is not used to finance consumption (C) will add to next period's wealth according to the rate of return in that period (R_w).

To examine the implications of this identity, Campbell and Mankiw (1989) consider a log-linear approximation to the dynamics implied by equation (1). By imposing a transversality condition[9] and assuming that the household has rational expectations they obtain:

$$c_t - w_t \approx E_t \sum_{k=1}^{\infty} \rho_w^k \left(r_{w,t+k} - \Delta c_{t+k} \right) \tag{2}$$

Where ρ_w is the reciprocal of one minus the average consumption to wealth ratio, which will typically be positive but less than one (throughout the paper, lower-case letters mean that the variable is expressed in natural logarithms). The equation states that the consumption-wealth ratio today should reflect rational forecasts of the returns on wealth and consumption growth. Intuitively, a rational household can only afford to have consumption in excess of its current wealth, *ex ante*, if its expected future wealth returns more than offset its expected future consumption growth.

Log aggregate wealth w in equation (2) contains the expected discounted value of future labour income, and is therefore unobservable. As such, it is not immediately clear how to test equation (2) empirically. Lettau and Ludvigson (2001) mitigate this issue by providing a set of assumptions that link unobservable aggregate wealth W to observable series on income and wealth. Firstly, aggregate wealth is equal to the sum of human wealth (HU), housing wealth (H), and net financial wealth (NF). As long as the share of each of these series in aggregate wealth is stationary, we can write this identity in log-linear form as $w_t = \omega_y hu_t + \omega_h h_t + (1 - \omega_y - \omega_h) nf_t$, where ω_y and ω_h are the steady state shares of human and housing wealth, respectively, in total lifetime resources. Secondly, if we model income as a dividend paid on the stock of human wealth, then unobservable human wealth can be linked to current labour income (y) according to $hu_t = \upsilon + y_t + z_t$, where z is a mean zero stationary random variable. This implies that income captures the non-stationary component of human wealth. Under these assumptions, the log of aggregate wealth is equal to current housing wealth, net financial wealth and income (all expressed in logs) plus a stationary residual z.

These assumptions yield the following approximate equivalent to equation (2):

$$c_t - \tau - \omega_y y_t - \omega_h h_t - (1 - \omega_y - \omega_h) nf_t = \eta_t \tag{3}$$

[9] The transversality condition implies that the log consumption-wealth ratio is zero in the limit. This ensures that consumption or wealth do not end up becoming an infinite fraction of each other. This certainly makes sense in the context of the individual household, but for the aggregate economy as a whole the consumption-wealth ratio could drift up or down, even in the limit.

Where τ is a constant and η_t is a residual. Under the maintained assumptions that the steady state share of each component of wealth is stationary and current income captures the non-stationary component of human wealth, the implications of equation (2) for the consumption-wealth ratio carry over to equation (3). Put another way, there is a strong argument that the residual η_t should be a stationary series. This residual reflects expectations of future consumption growth, income growth and the returns to housing wealth and financial wealth, as well as the residual *z* from the human wealth approximation. All of these series should in principle be stationary, such that, assuming that consumption, labour income, housing wealth and financial wealth are I(1), η_t is a cointegrating residual reflecting the deviation from the long-run relationship shown on the left-hand side of equation (3). The coefficients of this long-run relationship equal the steady state shares of human wealth, housing wealth and financial wealth in total expected lifetime resources, respectively. Implicitly, this framework assumes that all forms of wealth have the same marginal propensity to consume (MPC), such that the long-run elasticities depend only on wealth shares.[10]

Whelan (2006) provides an alternative means to investigate the implications of the budget constraint identity. Instead of beginning with a constraint involving unobservable expected lifetime resources, he uses the textbook version of the constraint which does not keep track of the accumulation of human wealth over time. The slightly modified budget constraint identity is $A_{t+1} = R_{a,t}\left(A_t + Y_t - C_t\right)$, where A_t is housing plus net financial wealth and $R_{a,t}$ is the time-varying return on these assets. Using a very similar approach to Campbell and Mankiw (1989), he derives the following equation summarizing the dynamic implications of this constraint:

$$xc_t - a_t = E_t \sum_{k=1}^{\infty} \rho_a^k \left(r_{t+k}^a - \Delta xc_{t+k}\right) \tag{4}$$

Where *xc* is the log of excess consumption (the difference between consumption and current labour income). The implications of this equation are much the same as equation (2), except that (by construction) the expected discounted value of future income does not play a role. Intuitively, the ratio of excess consumption to assets, the left-hand side of equation (4), reflects how much households are eating into their existing assets in any given period. A rational household can only afford to have this ratio positive if expected future returns on its housing and net financial wealth more than offset its expected future excess consumption growth. Equation (4) has the important property that it involves only observable variables. This means that we can test the implications of this equation directly by setting up long-horizon regressions exactly as specified by the equation. This provides an alternative test of the existence of the long-run relationship between consumption, income and wealth (albeit one that ignores the role of human wealth).

---

[10] Equation (3) also implies that theoretically, this long-run relationship eliminates any deterministic trends in the variables, given that there are no trends in the cointegrating residual. However, our discussion assumes a representative agent framework, which cannot account for changes in the composition of the population. If this composition is in fact slowly evolving over time, there would be good reason to incorporate a time trend in the long-run regression. For example, Hahn and Lee (2006) argue that if stock market participation slowly increases, this introduces a trend in the consumption-wealth ratio. However, when we augmented equation (3) with a time trend, the trend turned out to be statistically insignificant and including it had only a small effect on the other coefficient estimates.

Examining the implications of the long-run relationship

At any particular point of time the system is unlikely to be exactly at its long-run equilibrium, such that the co-integrating residual η_t is typically non-zero. We now move to examining how equilibrium tends to be restored from such a situation. Notice that equations (3) and (4) give us no guidance on how this occurs. The budget constraint simply tells us that, if households currently have a high consumption-wealth ratio (such that the cointegrating residual is positive) then future wealth gains need to offset future consumption growth to make this ratio sustainable. However, this says nothing about whether consumption, wealth or income will adjust in order to achieve this. If the disequilibrium reflects higher expected future wealth (income) gains and if these expectations are confirmed, then wealth (income) will increase to restore equilibrium. In contrast, if it reflects an expectation of lower consumption growth then consumption will tend to do the necessary adjustment.

For this reason, we use an empirical framework that allows for the possibility that any or all of the variables might play a strong role in the error-correction process. The Vector Error Correction model (VECM) provides such a framework. The generic VECM is:

$$\Delta X_t = A_0 + A(L)\Delta X_t + \alpha\eta_{t-1} + \varepsilon_t \qquad (5)$$

Where A_0 is a vector of deterministic terms, ε_t is a vector of stationary random variables, $A(L)$ is a polynomial in the lag operator reflecting the short-run dynamics of the system, and the vector α determines how disequilibria get ground out of the system. In this paper, $X_t = (c_t, y_t, h_t, nf_t)'$. The VECM is a four-equation system in consumption growth, income growth, housing wealth growth, and net financial wealth growth, their lags and the disequilibrium in the long-run relationship (η_{t-1}).

The VECM estimates, along with the estimates of the long-run relationship, allow us to decompose each variable into a permanent and a transitory component. The existence of a long-run relationship between the variables implies that at least some shocks must be transitory, otherwise the long-run equilibrium would never be restored. We perform the permanent and transitory decomposition using the approach developed by Gonzalo and Ng (2001). Our aim is to find a matrix *G* that transforms the reduced-form VECM residuals ε_t into a set of permanent and transitory shocks *u*:

$$u_t = G\varepsilon_t \qquad (6)$$

Gonzalo and Ng (2001) show that the following formulation for *G* will achieve this:

$$G = \begin{bmatrix} \alpha_{\perp}' \\ \beta \end{bmatrix} \qquad (7)$$

Where β captures the long-run relationship between the variables of the system. Equation (3) implies that $\beta = (1, \omega_y, \omega_h, (1-\omega_h-\omega_y))$. The matrix $\alpha_{\perp}$ is orthogonal to α such that $\alpha_{\perp}'\alpha = 0$. In our paper, ε_t has four elements and there is only one long-run relationship between the variables in the system. This implies that the vector β contains only one row, such that the fourth entry of the vector u_t is a single transitory shock. In contrast, $\alpha_{\perp}'$ has three rows, such that the first three entries of u_t represent three permanent shocks. We do not attempt to identify these shocks separately. Instead, we decompose the variation in each variable into a proportion due to the transitory shock and a proportion due to the permanent shocks.

4.2 Results

Existence of a long-run relationship

For the concept of a long-run relationship like equation (3) to be meaningful, the variables in the long-run relationship should be non-stationary. Table A2.1 in Appendix 2 presents the results from Augmented Dickey-Fuller (ADF) tests for stationarity using various orders of augmentation in the testing regression. For the orders of augmentation chosen by the Akaike and Schwarz information criteria, the null hypothesis of a unit root cannot be rejected for any of the four variables that potentially enter the long-run relationship.

In light of these results, Table A2.2 in Appendix 2 presents tests for cointegration between these variables. First, we implement Engle-Granger tests for cointegration. These tests consist of applying ADF tests to the residuals of a static OLS regression of consumption on income, housing wealth and financial wealth. At the order of augmentation selected by the Akaike and Schwarz information criteria, we can reject the null hypothesis of a unit root in these residuals at the 10% level. This constitutes moderate support in favour of the existence of a cointegrating relationship. Secondly, we apply Johansen tests based on the rank of the matrix $\alpha\beta'$ representing the long-run relationships in the VECM representation of the system. These tests reject the null hypothesis of no cointegration at the 5 percent level. In sum, bearing in mind that the ADF test has weak power for rejecting a unit root, we find substantial evidence for the existence of our hypothesized long-run relationship.

We can get some additional evidence on the existence of this long-run relationship by directly testing for the dynamic responses required to maintain this relationship according to the budget constraint identity. Table 2 presents results from regressing cumulative future gaps between asset returns and excess consumption growth on the log ratio of excess consumption to assets for various forecast horizons up to 5 years, as suggested by equation (4). For all forecast horizons, the coefficient on the ratio of excess consumption to assets is positive and statistically significant at the 5% level. This means that the ratio of excess consumption to net worth forecasts future movements in asset returns and excess consumption growth that keep the ratio from drifting too far up or down going into the future. This provides additional evidence in favour of the idea of a long-run relationship between consumption, income and wealth.

Table 2 Evidence for a long-run relationship based on long-horizon regressions

	Forecast Horizon, N (quarters)				
	1	4	8	12	20
λ	0.15**	0.36**	0.70**	0.96**	0.95**
	(2.61)	(2.46)	(4.98)	(7.77)	(8.30)
$\bar{R}^2$	0.09	0.22	0.50	0.77	0.83

Notes: This table presents long-horizon regressions of the form:

$$\sum_{k=1}^{N}\left(r_{t+k}^{a}-\Delta xc_{t+k}\right)=\tau+\lambda\left(xc_t-a_t\right)+e_t$$

Where *N* is the forecast horizon. Our definition of the rate of return on net wealth follows Whelan (2006). It is:

$$R_{a,t}=\frac{A_{t+1}}{\left(A_t+Y_t-C_t\right)}$$

t- statistics based on Newey-West standard errors are in brackets. * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level.

Estimates of the long-run relationship

Having provided substantial evidence for the existence of our hypothesized long-run relationship, we now estimate this long-run relationship. Table 3 shows these estimates, which were obtained using the Dynamic OLS (DOLS) estimator.[11] The long-run income elasticity is estimated to be 0.51 and significant at the five percent level. Furthermore, the long-run parameter for housing wealth is substantially larger than that for net financial wealth. A permanent one-percent increase in real per-capita housing wealth is associated with a 0.19% increase in real per-capita consumption in the long run, an effect which is statistically significant at the five percent level. In contrast, the elasticity from net financial wealth is small and statistically insignificant.

Table 3 DOLS estimates of the long-run relationship

	Housing wealth	Net financial wealth	Income
Estimate	0.20**	0.02	0.51**
	(6.21)	(0.63)	(3.93)

Notes: t-statistics are in parentheses. Estimates based on the Dynamic OLS estimator include differenced explanatory variables from *t-3* to *t+3*, and the standard errors were corrected for serial correlation according to the asymptotic distribution of the DOLS estimator.

The hypothetical scenario of equal MPCs is a useful benchmark for understanding our estimates. Under this scenario, the estimated long-run parameters reflect the respective steady-state wealth shares of human wealth, housing wealth and net financial wealth. It is plausible that human wealth accounts for most of aggregate wealth (as suggested by the relatively large estimate of the long-run parameter for income). In addition, the relatively large size of the housing wealth share is consistent with the fact that New Zealand households hold a large fraction (as high as 75 percent in recent years) of their physical wealth in the form of housing. For the same reason, it is not surprising that the steady state share of *net* financial assets is estimated at close to zero, given that this variable is constructed by

[11] DOLS differs from Static Ordinary Least Squares (SOLS) in that it adds leads and lags of differenced explanatory variables to the regression. It has two advantages over its more traditional variant, Static OLS (SOLS). Firstly, the SOLS estimator is biased in finite samples because it omits the short-run dynamics between the integrated variables. Secondly, unlike the SOLS estimator, the DOLS t-statistics can be rescaled to have a conventional asymptotic distribution, which makes inference possible. The SOLS results (not reported in the paper) are comparable to the DOLS results.

subtracting total household liabilities from financial assets, and as such is a very small number in comparison with housing wealth.

In Table A2.4 in Appendix 2, we show the main results of the paper under an alternative wealth decomposition, which defines housing wealth net of mortgage debt. This table shows that financial wealth defined net of non-mortgage debt has a larger parameter in the long-run relationship at 0.17, while the housing wealth parameter falls to 0.15 when it is defined net of mortgage debt. Thus, the percentage change in long-run consumption from a change in net housing wealth is comparable to that from a change in financial wealth net of non-mortgage debt. It is not surprising that financial wealth net of non-mortgage debt has a relatively large elasticity: it has recently been six times as large as our baseline definition of net financial wealth. However, Table A2.3 shows that the alternative wealth decomposition has almost no effect on our VECM estimates. As mentioned above, the advantage of our baseline decomposition is that it allows us to examine the interaction between housing wealth and liabilities.

VECM estimates of the dynamics

One of our main objectives is to understand how the equilibrium implied by this long-run relationship gets restored. To investigate this further, Table 4 presents a condensed version of the results of an estimated fourth-order VECM for the variables in the long-run relationship (this is a condensed version of these results; Table A2.3 in Appendix 2 presents a complete set of results).[12] In both the consumption and the housing wealth equations, the adjustment parameter on the cointegrating residual η_{t-1} is statistically significant at the five percent level. Starting from a point where the cointegrating residual is positive (so that consumption is above its trend level implied by income and housing wealth), consumption tends to fall and wealth tends to rise, working to restore equilibrium. As for labour income, the adjustment parameter is only significant at the ten percent level, but otherwise suggests that income moves in a way to restore equilibrium. Finally, the net financial wealth adjustment parameter is negative and quite large, but statistically insignificant. We discuss the interpretation of this estimate below.

Table 4 also shows some important interactions between the growth rates of the variables. Our results suggest that consumption growth is a positive function of past changes in housing wealth and income. Housing wealth and, to a lesser extent, net financial wealth are predictable by their own lags and seem to be very persistent. We also find that changes in either of the components of wealth predict future income growth. Finally, there are strong negative effects of housing wealth on net financial assets in subsequent quarters. Among other possible interpretations, this could reflect a positive linkage between housing wealth and liabilities (either through the collateral effect or passive equity withdrawal). Alternatively, households might be more willing to run down their gross financial assets following an increase in their housing wealth, which would be in line with buffer stock saving behaviour as in Carroll (1997).

---

[12] The Akaike and Schwarz information criteria suggested that we include only one lag in the VECM. However, the resulting specification suffered from serial correlation. We continued to include more lags until the errors were close to white noise. Further support for the fourth order specification was found from lag exclusion tests, which suggest that the higher order dynamics play a particularly important role in some of the equations, notably the equation for financial wealth.

Table 4 VECM estimates conditional on the DOLS long-run relationship

	Equation			
	Δc_t	Δh_t	Δnf_t	Δy_t
$\sum_{i=1}^{4}\Delta c_{t-i}$	0.26 (0.39)	-0.26 (0.93)	1.45** (4.94)	-0.01 (1.29)
$\sum_{i=1}^{4}\Delta h_{t-i}$	0.13** (3.82)	0.93** (14.02)	-0.95** (4.41)	0.21* (2.15)
$\sum_{i=1}^{4}\Delta nf_{t-i}$	0.00 (0.21)	0.08* (2.18)	0.27** (5.40)	0.12 (1.13)
$\sum_{i=1}^{4}\Delta y_{t-i}$	0.14 (1.33)	-0.10 (0.38)	0.35* (2.34)	-0.13 (0.27)
$\hat{\beta}'Z_{t-1}$	-0.18** (-2.19)	0.30** (2.58)	-0.37 (-1.46)	0.19* (1.75)
R^2	0.61	0.57	0.55	0.29

Notes: This table reports a condensed version of the VECM estimates conditional on the long-run relationship shown in table 2, reporting only the sum of the coefficients on the lagged growth rates from 1 to 4 quarters. The full table is in Appendix 2. The F-statistic for a test of the hypotheses that all four coefficients in the sum are zero is reported in parentheses for the growth rate terms (5% critical value = 2.50); t-statistics are shown in parentheses for the adjustment parameters. All equations include a dummy for GST changes in 1986/1989 (not reported). * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level.

Asymmetry in the equilibrium dynamics and the behaviour of net financial wealth

The point estimate for the adjustment parameter on net financial wealth implies that transitory movements in this variable tend to move the system further away from equilibrium. This means that when consumption is above its equilibrium given housing wealth and income, net financial wealth actually tends to decline. One potential reason for this effect is the strong impact of housing wealth on household liabilities. Although rising housing wealth has often helped to bring the system back to equilibrium following a positive cointegrating residual, it has also been associated with increases in liabilities that move the system further away from equilibrium.

However, we have already argued that the effect of housing wealth on liabilities could be asymmetric (see section 2, ‘Time varying strength of the consumption-wealth linkage’). Table 5 presents some tests for the implied asymmetry in the response of net financial wealth to the cointegrating residual. Our approach is to allow the adjustment parameters to vary depending on the value of a threshold dummy variable, such that equilibrium correction is allowed to be asymmetric depending on the sign of the threshold variable. Firstly, we allow the adjustment parameters to depend on the sign of the error correction residual, corresponding to the standard model of threshold cointegration investigated by Enders and Siklos (2001). Our estimates, as reported in the top rows of Table 5, suggest that net financial wealth responds far more negatively to the error correction term when consumption is below its equilibrium level than when it is above equilibrium. By itself, this implies that a negative cointegrating residual will take longer to get ground out of the system than a positive one. The bottom rows of Table 5 sheds further light on this asymmetry by allowing the adjustment parameters to depend on the sign of housing wealth growth.[13] These estimates suggest that net financial wealth responds much more negatively to the cointegrating residual when housing wealth growth is positive than when it is negative. Taken together, our results support the hypotheses

[13] Gonzalo and Pitarakis (2006) extend the threshold cointegration framework to allow for any exogenous variable to act as the threshold variable.

that increasing or above-equilibrium housing wealth quickly translates into increases in household debt, but that declining or below-equilibrium housing wealth does not immediately lead to declining liabilities.

Table 5 also reveals some important asymmetries in the adjustment parameters for the other variables in the system. From the top rows, we can see that the equilibrium correction due to housing wealth has been much stronger when the cointegrating residual was positive. This means that our estimate of a positive and significant adjustment parameter for housing wealth is primarily attributable to housing wealth increases, rather than decreases. The bottom schedule reveals that households are more willing to reduce (increase) their consumption to offset a positive (negative) cointegrating residual if the growth in housing wealth has recently been negative. This possibly reflects the negative impact of housing wealth growth on future consumption. Finally, income displays stronger equilibrium correction when housing wealth has recently been positive, perhaps because when housing wealth growth has been positive households might be more willing to borrow against their expected future income growth.

Table 5 Estimated adjustment parameters allowing for asymmetric error-correction

1. Asymmetry conditional on the sign of the error correction term				
	Equation			
	Δc_t	Δh_t	Δnf_t	Δy_t
η_{t-1} if $\eta_{t-1} > 0$	-0.20* (-1.66)	0.40** (2.44)	-0.12 (-0.36)	0.14* (0.89)
η_{t-1} if $\eta_{t-1} < 0$	-0.17 (-1.08)	0.12 (0.56)	-0.82* (-1.77)	0.25 (1.24)
R^2	0.61	0.57	0.56	0.29
2. Asymmetry conditional on the sign of real per-capita housing wealth growth				
	Equation			
	Δc_t	Δh_t	Δnf_t	Δy_t
η_{t-1} if $\Delta h_{t-1} > 0$	-0.09 (-1.01)	0.30** (2.28)	-0.51* (-1.82)	0.24* (1.90)
η_{t-1} if $\Delta h_{t-1} < 0$	-0.44** (-2.92)	0.26 (1.19)	-0.05 (-0.11)	0.03 (0.14)
R^2	0.63	0.56	0.56	0.30

Notes: This table reports VECM estimates conditional on the long-run relationship shown in table 2. The table only reports the results for the coefficients on the adjustment parameters, but the estimated VECM includes four lagged growth rates of all variables and a dummy for GST changes in 1986/1989. t-statistics are in parentheses. * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level.

The error-correction process in motion

To refine our interpretation of the dynamics of the system, Figure 3 plots the impulse responses to a one standard deviation negative transitory shock. For this experiment, we have set the adjustment parameters on net financial wealth and income, which were not significant at the 5% level, to zero. The negative transitory shock implies shocks to consumption and housing wealth: housing wealth rises 2 percent above its steady-state level, while consumption falls by 1 percent relative to its steady-state growth path. After the shock, housing wealth continues to increase for a few quarters, reflecting its high degree of persistence. In contrast, consumption immediately begins to increase towards equilibrium. The persistently elevated level of housing wealth causes a decline in net financial wealth (which reaches its trough at 4% below equilibrium) and also drags both consumption and income above equilibrium. Eventually, housing wealth does begin to decline, and the system gradually converges to equilibrium. This takes quite some time, with housing and financial wealth remaining a long way from equilibrium even three years after the shock.

Figure 3 Impulse responses from a one standard deviation transitory shock

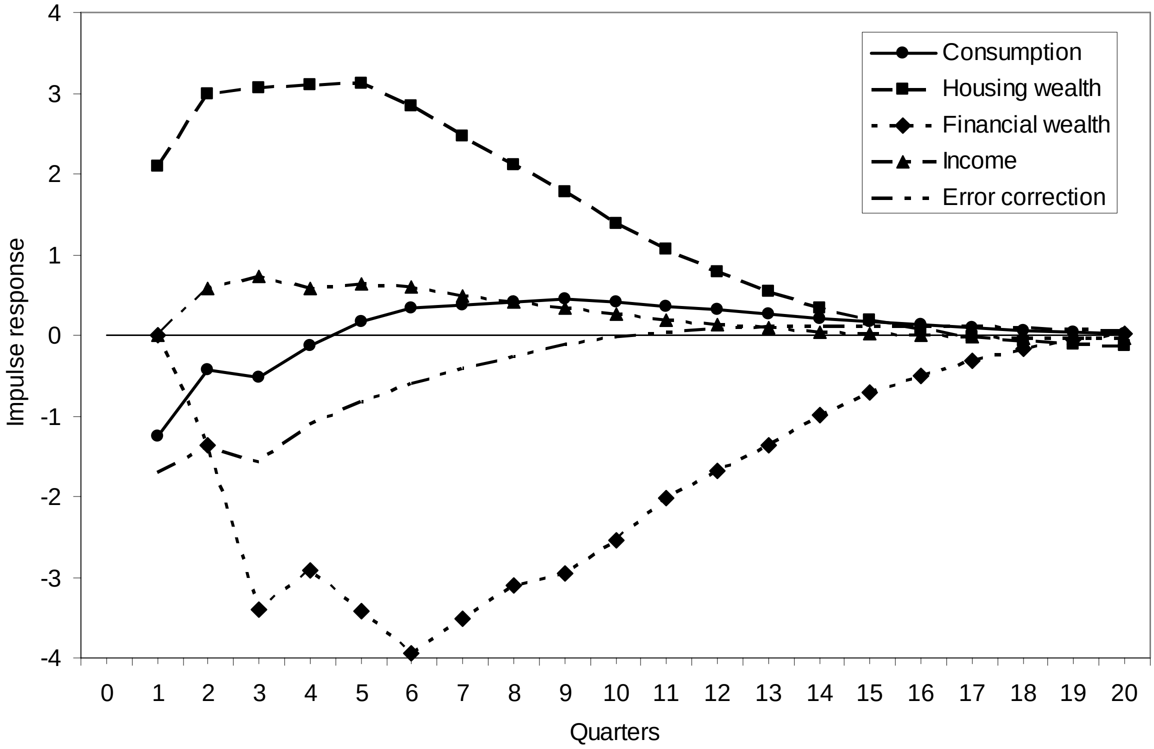

Permanent-transitory decomposition

Using the estimates of the long-run relationship along with the VECM estimates, we can explicitly decompose the forecast error variance for each variable into a portion that is due to the transitory shock and a portion that is due to the combined effect of the permanent shocks. Since the permanent shocks and the transitory shock are not necessarily orthogonal to each other, we also account for the portion of the forecast error variance that can be explained by the covariances among the shocks. We perform the variance decomposition both for the case where we set the financial wealth and income adjustment parameters equal to zero, as suggested by Gonzalo and Ng (2001), and for the case where they are set at their estimated values.

For each of the system's four variables, Table 6 decomposes the *h*-quarter ahead forecast error into a fraction due to the transitory shock, a fraction due to the permanent shocks, and a fraction due to the covariances between transitory and permanent shocks. For consumption, the fraction of the forecast error variance attributable to transitory movements is of a similar order of magnitude as that due to permanent shocks. This implies that consumption tends to gradually adjust to permanent changes in wealth or income, and that this adjustment has constituted a quantitatively important fraction of the overall movement in consumption. Although the transitory shock initially explains almost as much of the forecast error in housing wealth as the permanent shocks, permanent shocks are the dominant driver of housing wealth after a year. Thus, when multiple shocks are hitting the economy at different points in time, such that they effectively get averaged across forecast horizons, a typical movement in housing wealth will mainly be permanent. Meanwhile, income and financial wealth are also dominated by permanent shocks. (For income, this is especially true when its adjustment parameter is set equal to zero). This implies that our estimated long-run relationship gives a reasonably accurate measure of the elasticity of consumption to most changes in housing wealth, financial wealth and income.

Table 6 Forecast error variance decomposition by persistence of shocks

	$\Delta c_{t+h} - E_t \Delta c_{t+h}$			$\Delta h_{t+h} - E_t \Delta h_{t+h}$			$\Delta nf_{t+h} - E_t \Delta nf_{t+h}$			$\Delta y_{t+h} - E_t \Delta y_{t+h}$		
Horizon *h*	*T*	*P*	*P*,*T*	*T*	*P*	*P*,*T*	*T*	*P*	*P*,*T*	*T*	*P*	*P*,*T*
						$\alpha_{nf}, \alpha_y = 0$						
1	78	47	-25	190	222	-312	0	100	0	0	100	0
4	85	51	-36	104	183	-188	24	113	-36	10	105	-15
8	84	58	-42	75	147	-122	24	112	-36	10	104	-14
12	82	59	-41	76	137	-113	25	109	-34	10	103	-14
∞	82	59	-40	76	136	-113	26	108	-35	10	103	-14
						α_{nf}, α_y estimated						
1	51	69	-20	75	135	-110	17	92	-9	18	55	27
4	64	59	-23	39	122	-61	16	95	-11	18	57	25
8	61	63	-23	30	111	-41	16	94	-10	19	57	24
12	59	63	-23	29	109	-38	17	93	-10	19	57	24
∞	59	63	-22	29	108	-37	17	93	-10	19	57	24

Notes: This table reports the percentage of the forecast error variance at horizon *h* that is attributable to permanent and transitory shocks, as well as two times the covariance between them. It does so for two cases: one where we set adjustment parameters that are not significant at the 5 percent level equal to zero, and one where we leave them at their estimated values.

Instability in the long-run relationship

Any application of the cointegration framework assumes that there exists a stable long-run relationship between the I(1) variables in that relationship. While we need a sufficiently long span of data to appropriately disentangle the long-run and short-run relationships, in practice the relationship can hardly be expected to be stable over long horizons. The reason for this is that the structural factors which influence the strength of the relationships are rarely stable over longer horizons. For instance, the share of housing wealth in households' portfolio has been increasing, especially over the last six years, which will lead to instability if this reflects an increasing steady-state share of housing in total wealth. The question we turn to now is whether this instability is large enough to invalidate our main conclusions.

We test for stability of the cointegration relationship using two tests for unknown breakpoints: the Andrews (1993) and Andrews-Ploberger (1994) SupF and the AveF test, using p-values provided by Hansen (1992).[14] Figure 4 graphs the F-statistic for every observation. The F-statistic peaks in 1986Q4, and then quickly falls down to very low levels. Based on these F-statistics, the SupF test strongly suggests the existence of a statistically significant break at the 5 percent level, while the AveF test does not provide evidence for a statistically significant structural break at the 5 percent level, but is close to rejecting at the 10 percent level. Thus, there is evidence for instability in the cointegrating relationship in the mid-eighties.[15]

Figure 4 Sequence of Chow tests, with critical values for the *SupF* and *MeanF* tests.

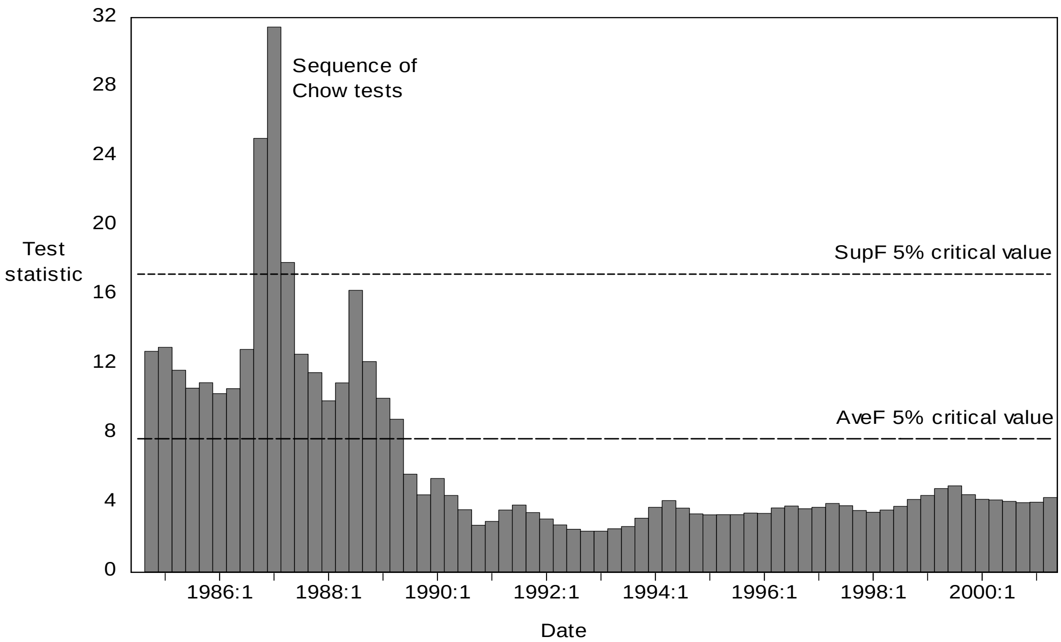

[14] For every observation, the test involves computing the F-statistic for the null hypothesis that there is no structural break in the coefficients at that time. The SupF test is based on the largest of those F-statistics, and as such tests against the alternative of a single structural break. The AveF test is based on the average of all F-statistics, and as such tests against the alternative of the coefficients following a martingale, without yielding an estimate for a particular break date.

[15] As Hansen (1992) notes, this does not mean that we should conclude in favour of the alternative of a single structural break in the cointegrating relationship. There are many assumptions underlying the linear cointegration model, and the rejection of the SupF test could reflect a rejection of any of them. The relevant conclusion is instead that there is some evidence that the linear cointegration model is misspecified.

It is interesting to note, however, that the break occurs around the time of the 1987 stock market crash. This might have caused households to be more cautious about acquiring financial assets. Another factor behind the absence of an increase in net financial wealth is that financial market liberalisation since the mid-eighties has tended to result in an increase in per-capita household debt. This change in the underlying economic structure may partly explain our rejection of the cointegration model around that time. Another candidate explanation is that the method used to construct national accounts data in New Zealand – which our data for consumption and, partly, income are based on – changes in 1987. Table 7 shows that, when we exclude data through 1986 from the sample, this really only affects the income coefficient. Hence, a first reading might suggest that the changes in the long-run relationship are relatively small, and therefore may not undo our main results. Nevertheless, this evidence provides us with more than enough reason to investigate the consumption-wealth linkage under an alternative set of assumptions than the cointegration model. We now turn to this task.

Table 7 – DOLS estimates of the long-run relationship with a split sample

	Housing wealth	Financial wealth	Income
1987-2006	0.16**	0.03	0.82**
	(3.03)	(1.39)	(6.13)
Full sample	0.21**	0.02	0.55**
	(6.21)	(0.63)	(3.93)

5. Alternative approaches: housing wealth vs. financial wealth effect

We consider some alternative models of the consumption-wealth linkage for two reasons:

- The alternative techniques directly estimate dollar-for-dollar effects rather than elasticities. Our previous estimates of the elasticity of spending to each form of wealth confound differences in wealth shares with differences in the strength of the dollar for dollar spending effect.

- Specifying the relationship between consumption and wealth in dollar-for-dollar terms also makes the estimates more robust to the instability in the long-run relationship. The first of the two approaches in the present section assumes that the dollar-for-dollar spending effects are stable over time, and thus not assume that wealth shares are stable. This seems like a more palatable assumption given the huge movements in the composition of household wealth over recent years. In our second alternative technique, the short-run estimates do not hinge on the existence of a stable long-run relationship.

At the outset, we highlight two key features of the alternative approaches:

- In order to directly estimate the MPC from each form of wealth, we work with stationary transformations of the components of wealth. In the first approach, wealth is specified as a ratio to income while in the second approach it is specified as a first difference divided by lagged consumption. In contrast to our model based on cointegration, this will make these models susceptible to endogeneity bias (we can no longer rely on the principle of a cointegrating regression where the variables are of a higher order of integration than the error term). For this reason, we need to include controls in the regression for variables that might simultaneously increase both wealth and consumption.

- The alternative approaches rely on single equation models where consumption growth is the only variable that evolves endogenously to the long-run relationship. Thus, the effects computed by these alternative techniques have a very specific interpretation: the increase in long-run consumption that would result after an increase in wealth, holding wealth and income constant thereafter. Thus, in this section we analyse the effect of permanent shocks to wealth and income, but do not allow for transitory variation in those variables.

Econometric framework for estimating the MPC from different forms of wealth

Like the cointegration approach, our first alternative approach incorporates a long-run relationship between consumption and the components of wealth. However, instead of the log asset formulation used in the VECM, assets enter the long-run relationship as a ratio to income. Furthermore, a number of control variables are used to estimate this long-run relationship. Following Aron and Muellbauer (2007), we use a partial adjustment specification for the dynamics of consumption. This implies that households gradually adapt their consumption level to reflect the gap between their target level (i.e. the level implied by the long-run relationship) and their current level of consumption. This is consistent with the idea that consumption growth responds sluggishly to permanent shocks, as occurs in models of consumption where households have habit persistence or sticky expectations. We estimate the following model:

$$\begin{aligned} \Delta c_t &= \alpha\left(c_t^* - c_{t-1} + \theta_{sr} Z_{sr,t}\right) + u_{1,t} \\ c_t^* &= y_{t-1} + \gamma_h \frac{H_{t-1}}{Y_{t-1}} + \gamma_f \frac{NF_{t-1}}{Y_{t-1}} + \theta_{lr} Z_{lr,t} \end{aligned} \tag{8}$$

Where c_t^* is the target level of consumption, α governs how quickly consumers adjust their consumption to reflect changes in the long-run target level and *Z* is a vector of controls (which is split into those controls that influence the long-run relationship and those that do not).

In practice, the speed of adjustment α can be estimated as the parameter on $y_{t-1} - c_{t-1}$ (the two variables that enter the dynamic equation with a parameter equal to α), and the other parameters are taken outside of the brackets and are estimated separately as short-run effects. For example, our estimates of the parameters on housing and financial wealth will reflect a short-run MPC. Dividing these estimated parameters by the speed of adjustment α gives the implied long-run MPC from each form of wealth, γ_h and γ_f. For the vector of controls, we use a parsimonious set of explanatory variables that have explanatory power for consumption growth: the real effective mortgage rate, the unemployment rate and income growth. Income growth should only have a short-run effect on consumption, but the other controls could enter the long-run relationship.

In our second approach, which is inspired by the work of Carroll, Otsuka, Slacalek (2006), the short-run estimates do not rely on any long-run information. Hence, the short-run estimates are robust to the presence of instability in the long-run relationship. The approach relies on dynamic models of consumption to compute the long-run MPCs implied by the estimates of the initial impact of wealth on consumption. The particular dynamic model of consumption used in this approach is again based on the idea that households sluggishly adjust their consumption to permanent shocks:

$$\Delta c_{t+1} = b_0 + \chi \Delta c_t + u_{2,t} \tag{9}$$

Consumption growth adjusts sluggishly as long as χ is positive and smaller than one. Given an estimate of χ and an estimate for how any initial shock affects current consumption, we can forecast the 'long-run' effect of the shock on future consumption levels by dividing the short-run effect by $(1-\chi)$.

However, consumption growth is measured with error, and so OLS estimation of equation (9) would produce a downward-biased estimate of χ. Similar to Carroll, Otsuka, Slacalek (2006), we use instrumental variables estimation to overcome this problem. This approach is valid so long as we choose instruments that are not correlated with the measurement error in consumption, but are correlated with the non-measurement error component of consumption growth. Changes in housing and financial wealth plausibly fulfil these requirements. We instrument for consumption growth by means of the following equation:[16]

$$\Delta c_t = a_0 + a_1 \frac{\Delta H_{t-1}}{C_{t-2}} + a_2 \frac{\Delta NF_{t-1}}{C_{t-2}} + a_3 Z_{t-1} + u_{3,t-1} \tag{10}$$

Not only does this two-step approach allow us to obtain consistent estimates of the sluggishness in consumption growth in the presence of measurement error, it also provides us with estimates of the short-run impact of changes in housing and financial wealth. Because wealth changes are defined relative to twice-lagged consumption in our regressions (which is a very similar denominator as that of the dependent variable), the coefficients a_1 and a_2 capture the short-run marginal propensity to consume from a change in housing and net financial wealth, respectively. *Z* refers to the same vector of controls as above. By means of the short-run MPCs from equation (10) and the estimates of the sluggishness of consumption growth from equation (9), we can compute the long-run MPC from changes in housing or financial wealth.

Estimates of the MPC from housing and financial wealth

Table 8 presents our estimates of the dynamic specifications for consumption growth suggested by our two alternative approaches.[17] In both models, the control variables enter with plausible signs, although they are not always significant. Turning to the estimates of the MPCs from housing and financial wealth, we find that both MPCs are reasonably large and economically and statistically significant. Our first alternative approach suggests that a one-

[16] The literature on estimating IV specifications of consumption growth suggests that the instruments used in this procedure should be lagged by at least two periods. This explains why all of the explanatory variables in (10) are lagged one period (resulting in them being determined two periods before C_{t+1}). Our specification is also slightly different from that of Carroll, Otsuka and Slacalek (2006), because these authors enter a weighted average of past wealth changes, rather than last quarter's wealth growth, into the regression. The weights constitute a geometrically declining series based on the estimates of sluggishness from equation (9). We found that the dynamics from wealth changes in our sample did not follow a geometrically declining pattern, so that we did not wish to impose this restriction.

[17] A few comments on the general properties of these equations are in order. Firstly, the fit of the two models is comparable. However, the second alternative approach suffers from serial correlation in the residuals, while there is no evidence of this in the first approach. This justifies the use of Newey-West standard errors in the second approach, and suggests to us that the exclusion of the levels information results in some degree of dynamic misspecification.

dollar increase in housing wealth leads to an annualized short-run increase in consumption of 0.9 cents, and a long-run increase of 5.4 cents. Our second alternative approach suggests a short-run effect of 2.4 cents, and a long-run effect of 7.5 cents. For net financial wealth, the short-run MPC is 2.4 cents and the long-run MPC is 14.1 cents for the first alternative approach, while for the second alternative approach these figures are 5.8 cents and 18.1 cents, respectively.[18] The difference between the housing and financial wealth effects is only significant at the 10 percent level in the first approach, and is not significant in the second approach. Nevertheless, the results provide substantial evidence that consumption responds more strongly to financial wealth than it does to changes in housing wealth.

Table 8: Alternative estimates of the consumption wealth linkage

	Aron and Muellbauer (2007)	Carroll, Otsuka, Slacalek (2006)
Speed of adjustment/sluggishness	0.17**	0.68**
	(2.69)	(2.45)
Unemployment rate	-0.0010	-0.0013
	(-1.41)	(-1.07)
Interest rate	-0.0009*	0.0000
	(-1.66)	(0.18)
Income growth	0.29**	0.20
	(3.48)	(1.41)
Housing wealth	0.0023**	0.0060*
	(2.63)	(1.77)
Financial wealth	0.0060**	0.0145**
	(2.43)	(2.16)
Implied short-run, long-run MPC housing wealth	[0.9, 5.4]	[2.4, 7.5]
Implied short-run, long-run MPC Financial wealth	[2.4, 14.1]	[5.8, 18.1]
p-value of test for equal MPC's	0.07*	0.21
p-value for serial correlation	0.25	0.05**
R^2	0.46	0.47

Notes: This table reports estimates of the alternative specifications designed to estimate the MPC from each form of wealth. Alternative approach one provides estimates of equation (9) by OLS. Alternative approach two is estimated by using equation (11) as the first-stage regression for a two-stage least squares procedure that produces estimates of equation (10). Approach Two was found to suffer from serial correlation so Newey-West standard errors were used. All regressions include a constant and a 1/-1 dummy for the implementation of changes to GST in 1986 and 1989. We tested for 2nd order serial correlation using a Breush-Godfrey test (based on the significance of lagged residuals in predicting consumption growth). * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level.

[18] We can compare the long-run estimates with the implied MPCs from our long-run estimates in table 5. Rescaling the elasticity by the average ratio of consumption to the relevant form of wealth over the period gives an MPC from housing of 7.2 cents and an MPC from net financial wealth of only 2.4 cents.

6. Conclusions and further research

This paper has revisited the relationship between wealth and consumption in New Zealand. We have noted a number of important pieces of evidence on the relationship. Firstly, we find evidence for a long-run relationship between consumption, income, housing and financial wealth. Secondly, we find that the majority of wealth movements have an economically significant impact on long-run consumption. We have also considered some alternative models that are more robust to instability in the long-run relationship and directly estimate the dollar-for-dollar effect from housing and financial wealth. These estimates suggest that the marginal propensities to consume from both components of wealth are economically important, with the MPC from net financial wealth exceeding that from housing wealth.

In our baseline approach, our results in this paper hinge on the assumption that the long-run relationship between consumption and wealth is stable. In the two alternative models, our results rely on the assumption that changes in wealth can be treated as permanent. Ideally, one would like to avoid making either of these assumptions. A fruitful avenue for future research in this direction would be implement a cointegration/VECM procedure in which the long-run coefficients explicitly vary over time as a function of the determinants of the long-run relationship. Many of these determinants, such as the degree of financial liberalisation and the composition of household portfolios, are likely to change only slowly over time within any given country, but are likely to exhibit substantial variation across countries. Therefore, a cross-country approach would be particularly useful for this exercise.

References

Alemeida, H, M Curill and C. Liu (2006), “The Financial Accelerator: Evidence from International Housing Markets,” *Review of Finance,* 10:3, 321-352.

Ando, A and F Modigliani (1963), “The ‘Life Cycle’ Hypotheses of Savings: Aggregate implications and tests,” *American Economic Review,* 53(1), 55-84.

Andrews, D (1993): “Tests for Parameter Instability and Structural Change With Unknown Change Point”, Econometrica, 61(4), 821-856.

Andrews, D, and W Ploberger (1994): “Optimal Tests when a Nuisance Parameter is Present Only Under the Alternative'”, Econometrica, 62(6), 1383-1414.

Aron, J, and J Muellbauer (2007), “Housing wealth, credit conditions and UK consumption,” *CEPR Working Paper* (forthcoming).

Bertaut, C (2002), “Equity Prices, Household Wealth, and Consumption Growth in Foreign Industrial Countries: Wealth Effects in the 1990s,” *Federal Reserve Board IFDP Working Paper*, 724.

Campbell, J and G Mankiw (1989), “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence,” in *NBER macroeconomics annual,* eds O Blanchard and S Fischer, Cambridge, MA.

Carroll, C (1997), “Buffer-stock saving and the life cycle/permanent income hypothesis,” *Quarterly Journal of Economics,* 112(1), 1-55.

Carroll, C, M Otsuka and J Slacalek (2006), “How large is the housing wealth effect? A new approach,” *NBER Working Paper,* 12746.

Case, K, J Quigley and R Shilller (2001), “Comparing wealth effects: the stock market versus the housing market,” *Cowles Foundation Discussion Paper,* 1335.

Chen, J (2006), “Re-evaluating the association between housing wealth and aggregate consumption: New evidence from Sweden,” *Journal of Housing Economics,* 15(4), 321-348.

Dvornak, N, and M Kohler (2003), “Housing Wealth, Stock Market Wealth and Consumption: A Panel Data Analysis for Australia,” *Reserve Bank of Australia, Discussion Paper,* 2003/7.

Enders, W, and P Siklos (2001), “Cointegration and threshold adjustmentm,” *Journal of Business and Economic Statistics,* 19(2), 166-176.

Fernandez-Corugedo, E, S Price and A Blake (2003), “The Dynamics of Consumer Expenditure: The UK Consumption ECM Redux,” *Bank of England Working Paper,* 204.

Fisher, L, G Otto, and G Voss (2005), “Housing Booms, Non-Financial Wealth and Consumption: Lessons from the Australian Experience,” *University of Victoria,* mimeo.

Goh, K, and R Downing (2002), “Modelling New Zealand consumption expenditure over the 1990s,” *New Zealand Treasury working Paper,* 2002/19.

Gonzalo, J, and S Ng (2001), "A Systematic Framework for Analyzing the Dynamic Effects of Permanent and Transitory Shocks," *Journal of Economic Dynamics and Control* 25, 1527–46.

Gonzalo, J, and J Pitarakis (2006), "Threshold Effects in Multivariate Error Correction Models," in *Palgrave Handbook of Econometrics: Volume I: Econometric Theory*, eds C Mills and K Patterson, Palgrave Macmillan.

Hahn, J, and H Lee (2006), "Interpreting the predictive power of the consumption–wealth ratio," *Journal of Empirical Finance,* 13(2), 183-202.

Hansen, B (1992), "Tests for Parameter Instability in Regressions with I(1) Processes," *Journal of Business and Economic Statistics*, 10(3), 45-59.

Hull, L (2003), "Financial regulation and household indebtedness", *Reserve Bank of New Zealand Discussion Paper,* DP2003/01.

Lettau, M, and S Ludvigson (2001), "Consumption, Aggregate Wealth, and Expected Stock Returns," *The Journal of Finance*, 66(3), 815–49.

Lettau, M, and S Ludvigson (2004), "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption," *American Economic Review*, 94(1), 276-299.

Modigliani, F, and R Brumberg (1954), "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," in *Post Keynesian Economics,* ed K Kurihara, New Brunswick, N.J.

Muellbauer, L, and R Lattimore (1995), "The consumption function: a theoretical and empirical overview," in *Handbook of Applied Econometrics*, eds H Pesaran and M Wickens, Oxford, Blackwell.

Paiella (2007), "The stock-market, housing and consumer spending: A survey of the evidence on wealth effects," mimeo.

Pichette, L, and D Tremblay (2003), "Are Wealth Effects Important for Canada?," *Bank of Canada Working Paper,* 2003/30.

Rudd, J, and K Whelan (2006), "Empirical Proxies for the Consumption-Wealth Ratio", *Review of Economic Dynamics*, 9(1), 34–51.

Whelan, K (2006): 'Consumption and Expected Asset Returns without assumptions about Unobservables", *Central Bank & Financial Services Authority of Ireland Research Technical Paper*, 2006/4.

Appendix 1 Data

This appendix gives more details on the main variables in our dataset.

Consumption

We take our data on total private consumption from the quarterly national accounts release. We use total household consumption, excluding the consumption of private non-profit organisations.

Labour Income

There is no direct measure of quarterly labour income in New Zealand. Our measure is similar to that used by Khoon and Goh (2002). We define labour income as:

$$y_t = WAGES - TAX + TRANSFERS$$

WAGES is quarterly wage income, constructed by multiplying the average hourly earnings (including overtime payments) for each quarter from the Quarterly Employment Survey (QES) by hours worked from the Household Labour Force Survey (HLFS). The HLFS includes hours worked by agricultural sectors and self-employed workers, but the QES excludes these sectors. Therefore, the measure we construct is a proxy for compensation of employees plus entrepreneurial income, under the assumption that the earnings of self-employed and farmers are similar to the rest of the economy. Finally, we added back a smoothed difference between the calendar year total for this proxy and the annual national accounts measure on a pro rata basis. This was done to ensure that our quarterly proxy was consistent with the best available annual data.

Quarterly tax payments (*TAX*) are constructed using our measure of wage income and a smoothed version of the implied effective tax rate from the annual national accounts. Downing and Goh (2002) instead matched average income payments to observable tax schedules. Our approach has the advantage that the effective tax rate will reflect the interaction between tax schedules and the income distribution. This should be an effective method to capture movements in tax payments as long as tax schedules and the income distribution don't move much between years, so that quarter-quarter movements in tax payments can be proxied for by movements in gross incomes.

Where possible, transfer payments (*TRANSFERS*) were linked to observable benefit rates and an estimate of number of beneficiaries. However, this was possible only for the unemployment benefit and the pension benefit, where a measured sector of the population receives the benefit. Benefit rates for these groups were obtained from representative rates compiled by the Ministry of Social Development. These benefits account for less than half of transfer payments. However, the other components – such as the domestic purposes benefit and the invalid and sickness benefit – are unlikely to vary much from quarter to quarter, so we just add back a smoothed residual between the annual and quarterly proxies to account for these payments.

Housing values

The value of the housing stock prior was obtained from official RBNZ estimates (available from the RBNZ website, although we use unrounded values). The primary input into these estimates is the Quotable Value New Zealand (QVNZ) house price index. QVNZ obtains capital values from local authorities who conduct periodic revaluations for the purpose of levying rates. The ratio of actual sales prices to the capital values can be used to produce a price index for each local authority area, which can then be aggregated.

For the period 1995 to 2004, the RBNZ has obtained official annual estimates of March quarter total residential dwelling values from QVNZ. By using Statistics New Zealand data on dwelling consents (ie measuring the total number of dwellings) and the house price index, quarterly housing values data were constructed. Prior to 1995, quarterly housing values were obtained based purely on dwelling consents and the house price index.

Financial assets and liabilities

Data on financial assets and liabilities are taken from RBNZ estimates. The data on financial liabilities are official series that can be split into mortgage and ‘other’ debt. Quarterly data on financial assets are taken from estimates that are currently not publicly available. They are constructed with the same methodology that is currently used to produce the annual estimates of financial assets that are available publicly from the RBNZ website (for more on this dataset see Ung and Thorp (2002)).

However, the data for financial assets is only available after 1995. With the exception of directly held equities, most of the components of financial assets – including deposits, fixed interest securities and mutual funds – didn’t exhibit much volatility from quarter to quarter. Interpolation using a cubic spline provided an excellent fit, as could be verified by comparing the interpolated data with the actual after 1995. To construct a quarterly measure of directly held equities, we matched the growth rates in the series each quarter with a weighted average of the New Zealand and Australian capital price indices (converted to NZ$; the weights were equal to the share of overseas and domestic directly held equities in each year). The series was then manipulated to ensure it always equalled actual directly held equities for the fourth quarter of the year.

As a final caveat, note that in New Zealand, unincorporated businesses (including farms) are not accounted for in the available statistics on household net worth and its components.

Appendix 2 Additional results

Table A2.1 ADF tests for stationarity

	Lag length						
	0	1	2	3	4	8	12
Consumption	-2.10	-1.71	-2.12	-2.70[S]	-2.64	-2.44[A]	-3.58**
Housing wealth	0.12	-1.93[A,S]	-2.02	-2.43	-2.26	-2.97	-1.90
Financial wealth	-0.09	-0.78	-0.91	-1.58	-0.79[A,S]	-1.45	-2.55
Income	-2.43[A,S]	-2.20	-2.21	-2.61	-3.12	-3.02	-3.54**
5,10 percent critical value	-3.45, -3.15						

Notes: The table reports tests of the null hypotheses that the series is I(1). All of the tests include both a trend and a constant, given that all series exhibit strong growth over our sample period. Tests to exclude the trend were undertaken, and none of them rejected. 'Lag length' refers to the order of augmentation of the testing regression. Superscripts with A,S refer to the lag length selected by the Akaike and Schwarz information criteria, respectively. * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level.

Table A2.2 Tests for cointegration

	Lag length						
	0	1	2	3	4	8	12
Residual based ADF test	-4.02*[A,S]	-3.45	-3.46	-3.42	-3.37	-3.06	-2.09
5, 10 percent critical value	-4.16, -3.84						
L-max test for cointegration	34.3**	33.5**[AS]	41.8**	47.9**	27.2*	44.3**	44.7**
5 percent critical value	27.58						
Trace test for cointegration	76.0**	63.5**[AS]	73.3**	66.5**	64.3**	84.9**	110.4**
5 percent critical value	47.85						

Notes: The table reports tests of the null hypotheses that there is no cointegration between consumption, income, housing wealth and financial wealth. Residual based ADF tests test for stationarity in the residuals of an OLS regression of consumption on the other components of the hypothesized long-run relationship. The L-max and Trace tests are based on the rank of the parameters of the lagged levels matrix in the VECM representation of the system. The L-max test uses the alternative of exactly one cointegrating relationships, while the Trace test uses the alternative of at least one cointegrating relationship. 'Lag length' refers to the number of lagged differenced terms to the order of augmentation for the ADF tests, and the number of lagged differences included in the VECM for the L-max and Trace tests. Supercripts of A,S refers to the lag length chosen by the Akaike and Schwarz information criteria, respectively. Asymptotic critical values for the Residual based ADF test are taken from Hayashi (2002). The critical values for the Johansen tests are from E-views. * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level.

Table A2.3 Complete estimates of the VECM model

	Equation			
	Δc_t	Δh_t	Δnf_t	Δy_t
Δc_{t-1}	-0.00 (-0.08)	0.05 (0.30)	0.01 (0.02)	-0.18 (-1.15)
Δc_{t-2}	0.11 (1.01)	-0.13 (-0.90)	1.19** (3.81)	0.14 (1.01)
Δc_{t-3}	0.06 (0.54)	-0.02 (-0.13)	0.51 (1.50)	0.12 (0.77)
Δc_{t-4}	0.09 (2.76)	-0.19 (-1.37)	-0.29 (0.98)	-0.11 (-0.81)
Δh_{t-1}	0.24** (3.02)	0.67** (5.98)	-0.65** (-2.69)	0.28** (2.60)
Δh_{t-2}	-0.33** (-3.20)	-0.05 (-0.32)	-0.32 (-1.06)	-0.01 (-0.11)
Δh_{t-3}	0.18 (1.69)	0.31** (2.12)	-0.21 (-0.67)	-0.02 (-0.16)
Δh_{t-4}	0.04 (0.38)	-0.02 (-0.18)	0.28 (1.01)	-0.05 (-0.37)
Δnf_{t-1}	0.02 (0.52)	0.09* (1.94)	0.27** (2.44)	0.08* (1.73)
Δnf_{t-2}	-0.02 (-0.67)	0.07 (1.47)	-0.07 (-0.61)	0.05 (1.02)
Δnf_{t-3}	0.02 (0.61)	-0.03 (-0.51)	0.29** (2.72)	-0.00 (-0.04)
Δnf_{t-4}	-0.01 (-0.31)	-0.07 (-1.47)	-0.16 (-1.56)	-0.01 (-0.27)
Δy_{t-1}	0.20** (2.13)	-0.04 (-0.28)	0.69** (2.46)	-0.09 (-0.70)
Δy_{t-2}	0.03 (0.37)	0.08 (0.62)	0.14 (0.49)	0.02 (0.14)
Δy_{t-3}	-0.06 (0.09)	-0.11 (-0.88)	0.02 (0.08)	0.02 (0.18)
Δy_{t-4}	-0.05 (-0.54)	0.01 (0.04)	-0.56** (-2.10)	-0.06 (-0.50)
$\hat{\beta}' Z_{t-1}$	**-0.18** (-2.19)**	**0.30** (2.58)**	**-0.37 (-1.46)**	**0.19* (1.75)**
R^2	0.61	0.57	0.55	0.29

Notes: This table reports VECM estimates conditional on the long-run relationship shown in table 2. All equations include a dummy for GST changes in 1986/1989 (not reported). * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level.

Table A2.4 Estimates of the cointegrated model with housing wealth defined net of mortgage debt

1. DOLS estimates of the long-run relationship

	Housing wealth	Financial wealth	Income
Estimate	0.15**	0.17	0.57**
	(2.64)	(1.46)	(3.82)

Notes: T-statistics are in parentheses. Estimates based on the Dynamic OLS estimator include differenced explanatory variables from *t-3* to *t+3*, and the standard errors were corrected for serial correlation according to the asymptotic distribution of the DOLS estimator.

2. VECM estimates conditional on the DOLS long-run relationship

	Equation			
	Δc_t	Δh_t	Δf_t	Δy_t
Δc_{t-1}	0.01	0.07	-0.25	-0.17
	(0.06)	(0.31)	(-1.18)	(-1.00)
Δc_{t-2}	0.15	-0.26	0.63**	0.14
	(1.31)	(-1.32)	(3.28)	(0.98)
Δc_{t-3}	0.04	-0.04	0.08	0.14
	(0.35)	(-0.21)	(0.41)	(0.93)
Δc_{t-4}	-0.03	-0.17	-0.42	-0.05
	(-0.32)	(-0.95)	(-2.42)	(-0.39)
Δh_{t-1}	0.17**	0.66**	-0.18	0.20**
	(2.59)	(5.70)	(-1.54)	(2.28)
Δh_{t-2}	-0.23**	-0.08	-0.16	-0.03
	(-2.89)	(-0.54)	(-1.09)	(-0.27)
Δh_{t-3}	0.13	0.23	0.07	-0.08
	(1.52)	(1.57)	(0.51)	(-0.76)
Δh_{t-4}	0.06	0.01	0.26	-0.03
	(0.83)	(0.12)	(2.14)	(-0.32)
Δf_{t-1}	0.06	0.11	0.44**	0.08
	(1.04)	(1.17)	(4.65)	(1.10)
Δf_{t-2}	-0.07	0.12	-0.08	0.04
	(-1.15)	(1.17)	(-0.84)	(0.57)
Δf_{t-3}	0.06	0.00	0.28**	-0.02
	(1.13)	(0.08)	(2.88)	(-0.21)
Δf_{t-4}	-0.03	-0.07	-0.17*	-0.02
	(-0.62)	(-0.70)	(-1.79)	(-0.21)
Δy_{t-1}	0.18*	-0.05	0.49**	-0.06
	(1.83)	(-0.29)	(2.89)	(-0.46)
Δy_{t-2}	-0.015	0.22	0.07	0.09
	(-0.15)	(1.27)	(0.43)	(0.13)
Δy_{t-3}	-0.13	-0.02	0.02	0.10
	(-1.38)	(-0.12)	(0.16)	(0.83)
Δy_{t-4}	-0.11	0.09	-0.37**	-0.02
	(-1.20)	(0.55)	(-2.30)	(-0.14)
$\hat{\beta}'Z_{t-1}$	**-0.20****	**0.31****	**0.23**	**0.15**
	(-2.23)	**(1.95)**	**(1.48)**	**(1.23)**
R^2	0.61	0.57	0.55	0.29

Notes: This table reports VECM estimates conditional on the long-run relationship shown in Table 2. All equations include a dummy for GST changes in 1986/1989 (not reported). * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level.

Table A2.5 Estimates of the alternative models with housing wealth defined net of mortgage debt

	Approach One	Two
Speed of adjustment/sluggishness	0.22**	0.67**
	(2.69)	(2.44)
Unemployment rate	-0.0007	-0.0018*
	(-1.08)	(-1.66)
Interest rate	-0.0002*	0.00014
	(-0.46)	(0.44)
Income growth	0.32**	0.18
	(3.98)	(1.35)
Housing wealth	0.0016**	0.0048
	(2.05)	(1.35)
Implied short-run, long-run MPC	[0.6, 2.9]	[1.9, 5.8]
Financial wealth	0.0084**	0.0168**
	(3.28)	(2.16)
Implied short-run, long-run MPC	[3.4, 15.3]	[5.8, 20.4]
R^2	0.46	0.47

Notes: This table reports estimates of the alternative specifications designed to estimate the MPC from each form of wealth. Alternative approach one provides estimates of equation (9) by OLS. Alternative approach two is estimated by using equation (11) as the first-stage regression for a two-stage least squares procedure that produces estimates of equation (10). Approach Two was found to suffer from serial correlation so Newey-West standard errors were used. All regressions include a constant and a 1/-1 dummy for the implementation of changes to GST in 1986 and 1989. We tested for 2nd order serieal correlation using a Breush-Godfrey test (based on the significance of lagged residuals in predicting consumption growth). * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level.