

Does Wealth Variation Matter for Consumption?*

Emmanuel De Veirman and Ashley Dunstan[†]

Reserve Bank of New Zealand

January 25, 2010

Abstract

The consumption wealth effect is often interpreted as capturing the effect of unexpected permanent movements in wealth on consumption. Relatively recently however, Lettau and Ludvigson (2004) found that in the United States, changes in wealth are mostly transitory, implying that only a small fraction of total wealth variation actually has a measurable effect on consumption. We use data for New Zealand, where transitory stock market cycles have not dominated variation in household net worth. We find that most changes in housing and financial wealth have been permanent, and that a typical wealth change does imply a lagged consumption response. In addition, we estimate a model which explicitly characterizes the role of household debt. Our evidence suggests that there is a positive relation between gross household wealth and debt, in line with the hypothesis that liquidity constraints or precautionary saving influence the strength of the relation between wealth and consumption.

JEL codes: C22, C32, E21

Keywords: Consumption wealth effect, permanent-transitory decomposition.

*We thank Andrew Coleman, Felix Delbrueck, Karen Dynan, Khoon Goh, Arthur Grimes, David Hargreaves, Andreas Lehnert, Cath Sleeman, Clive Thorp, as well as members of the Reserve Bank's Monetary Policy Committee for comments and discussion. Any errors are our own. 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.

[†]Emmanuel De Veirman is a senior economist in the research team of the Reserve Bank of New Zealand. Ashley Dunstan is an economist in its modelling team. Corresponding author: deveirman@yahoo.com, Reserve Bank of New Zealand, PO Box 2498, Wellington 6140, New Zealand.

1 Introduction

The early work on the permanent income hypothesis by Friedman and Modigliani has constituted a framework for much of the subsequent literature on the effect of wealth and income changes on consumer spending.¹ In a lifecycle framework, consumers alter their spending plans in response to changes in permanent income, i.e. the annuity value of expected lifetime resources.

Typically, papers estimating the consumption wealth effect empirically have focused on measuring the consumption response to permanent changes in wealth. A common approach to characterizing the wealth effect is to estimate a long-run consumption function characterizing the trend comovement of consumption, wealth, and income, and to then estimate the short-run response of consumption to any deviations from the common trend.² This approach implicitly treats all wealth variation as exogenous and permanent.

Relatively recently however, Lettau and Ludvigson (2004) found that changes in US household wealth were mainly transitory and unrelated to changes in consumption. Bearing in mind that typical estimates of the wealth effect only capture permanent wealth changes, Lettau and Ludvigson conclude that such estimates greatly overstate the response of consumption to a typical change in wealth.

Lettau and Ludvigson explain the importance of transitory shocks for household wealth by pointing to transitory but persistent cycles in stock market wealth. Assuming that households anticipate mean reversion in stock market wealth, it is indeed plausible from a lifecycle model's perspective that surprise changes in stock wealth would not cause households to greatly alter their spending plan.

¹ For seminal work, see Friedman (1957) and Ando and Modigliani (1963). For some influential tests of the permanent income hypothesis, see Hall (1978), Flavin (1981), Campbell and Mankiw (1989) and Carroll (1994).

² See Davis and Palumbo (2001) for a review of this approach.

If observed fluctuations in US household wealth are indeed dominated by boom-bust cycles in stock markets, and if one accepts that this particular type of wealth movements does not induce large consumption responses, the question of interest then becomes whether any other wealth change actually has a bearing on consumer spending. Arguably, Lettau and Ludvigson's main finding does not speak for the relation between underlying wealth changes and consumption, since any such relationship would be masked by the dominance, in terms of volatility, of one particular type of wealth unrelated with consumer spending.

This paper examines whether changes in wealth other than those related to stock cycles affect consumption in a way that is consistent with the permanent income hypothesis. To do so, we use data for New Zealand, where boom-bust cycles in stock markets have not been as predominant over our sample period in terms of their direct effect on household wealth variation. This is largely due to households' portfolio composition: over our sample period, housing on average accounts for two thirds of household assets, while direct equity constitutes only 5% of total assets.

To find out whether a typical change in wealth affects future consumption, we use a Vector Error Correction Model (VECM) which in principle allows both consumption and the components of wealth to correct for any short-run discrepancy between consumption and wealth levels. Typical empirical estimates of a consumption function derived from a lifecycle model, as reviewed by Davis and Palumbo (2001), only allow for error-correction in consumption, and can therefore be seen as a special case of the model which we implement. Recall that such procedures focus on estimating the effect from permanent and exogenous wealth changes. The approach which we use allows us to verify empirically whether such estimates accurately capture the effect of a typical wealth change in an environment where stock market cycles have little direct influence on household wealth.³

³ Our econometric approach is similar to that of Lettau and Ludvigson (2004). Recently, VECM

Our results suggest that most changes in housing wealth, financial wealth and consumption are permanent, such that our estimated long-run relationship adequately captures the effect of a typical wealth change. The long-run effect of financial wealth turns out to exceed that of housing wealth.

Unlike Lettau and Ludvigson's finding for the United States, we find that there is non-trivial transitory variation in consumption, and that consumption error-corrects to short-term discrepancies between consumption and wealth. This suggests that permanent changes in wealth imply a lagged consumption response. In sum, our results suggest that a typical change in wealth matters for consumption.

While permanent shocks account for most of the variation in housing wealth, we also find that there is non-trivial transitory variation in housing wealth. This reflects the fact that housing wealth adjusts to short-run discrepancies between wealth and consumption. Arguably, endogenous changes in wealth such as these are not the focus of lifecycle theory. Typically, the 'wealth effect' is interpreted as the effect of unexpected wealth movements (or changes in expectations regarding future wealth movements) on consumption, while endogenous changes in wealth are already incorporated in the households' spending plan.

However, our finding that housing wealth responds to short-run deviations from equilibrium does not necessarily imply a rejection of the standard model of the consumption wealth effect. To summarize our argument here, it is likely that in empirical terms, both consumption and house prices depend on expectations regarding future income and wealth movements. If housing wealth is the more persistent variable (which our estimates suggest), changes in expectations will tend to affect consumption

models involving a long-run consumption equation have also been applied to the UK by Fernandez-Corugedo, Price and Blake (2003), to Australia by Fisher, Otto, Voss (2005), to Sweden by Chen (2006), and to Germany by Hamburg, Hoffmann and Keller (2007). In these papers, consumption does not play an important role in restoring equilibrium, in line with Lettau and Ludvigson's results.

before being reflected in housing wealth. The subsequent change in housing wealth will appear to correct for the discrepancy implied by the earlier rise in consumption.

Much of the existing literature on the wealth effect does not characterize household debt explicitly, but instead models debt implicitly as part of net household wealth. In an extension of our model, we also estimate a system in which gross wealth enters separately from household liabilities. This model implies a long-run equation for consumption as well as a separate equation which relates debt to gross wealth and income. We detect a statistically significant relation between trend debt levels, gross wealth and income. Our findings suggest the presence of factors such as liquidity constraints or buffer stock saving behavior, factors which tend to increase the marginal propensity to consume out of wealth.

Furthermore, the gross wealth model serves as a reminder that any finding regarding the importance of permanent and transitory shocks in total wealth variation is model-dependent. In particular, the gross wealth model involves two long-run relations to which variables potentially error-correct, and therefore implies somewhat more transitory variation throughout. Nevertheless, it turns out that in our case, permanent shocks continue to account for most of the variation in all variables.

Finally, this paper uses cointegration methods to estimate the long-run relation between wealth and consumption. While this is a popular approach, it is subject to the criticism that it hinges on the existence of a stable long-run relation between wealth, income and consumption.⁴ To partially address this criticism, we restrict our sample to a relatively stable economic environment, avoiding the large structural reforms that occurred in New Zealand in the 1980s. Furthermore, we document that our long-run coefficient estimates are stable according to the relevant structural break tests with unknown break date.⁵

⁴ See Carroll, Otsuka and Slacalek (2006) and Rudd and Whelan (2006).

⁵ We implement Seo (1998) break tests, which are applicable because we use Full Information Maximum Likelihood techniques to jointly estimate the long-run relation(s) and the VECM.

Perhaps more convincingly, we show that our conclusions regarding the adjustment to discrepancies between consumption and wealth are robust to using an alternative framework by Whelan (2008) which does not require the estimation of a long-run relationship. In analogy with our VECM results, we find that wealth gains have played a less important role in restoring equilibrium than in the United States, while consumption has played a more important role.

In a similar vein, this paper's working paper version, De Veirman and Dunstan (2008), estimates two additional models which do not rely on the assumption of long-run stability, one inspired by Carroll, Otsuka and Slacalek (2006) and another due to Aron and Muellbauer (2006). Like much of the existing literature on the consumption wealth effect, these techniques focus on the effect from permanent changes in wealth. In any case, they yield long-run MPCs from housing and financial wealth that are qualitatively similar to those in the present paper.

Our paper proceeds as follows. Section 2 discusses the theoretical foundations of the long-run relations which we estimate, and provides economic intuition behind these relations. Section 3 describes the data. Section 4 provides long- and short-run results based on a system involving consumption, net wealth, and income. Section 5 extends the model to allow for an explicit role of household debt. Section 6 tests for long-run stability, and shows that our conclusions regarding adjustment to wealth-consumption discrepancies are robust to using an alternative framework which does not assume long-run coefficient stability. Section 7 concludes.

2 Estimation framework

2.1 Foundations and intuition of the long-run consumption function

In a typical lifecycle model, an unexpected increase in wealth will increase consumption through its impact on the expected lifetime resources of the household. The increase in lifetime resources, and therefore in permanent income, allows the household to 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: the *wealth effect*.

The wealth effect is typically interpreted as capturing the response of consumption to exogenous (present or future) changes in asset wealth and income. In the present paper, we estimate a model which in principle allows wealth and income to systematically adjust to short-term discrepancies between consumption and wealth, so that wealth may play a role in ensuring that the intertemporal budget constraint is satisfied. As part of that model, we will estimate a long-run consumption function similar to that which Campbell and Mankiw (1989) derived from the intertemporal budget constraint of a rational representative household.

According to the intertemporal budget constraint, wealth, defined in a broad sense so that it includes human wealth as well as asset wealth, follows the following law of motion:

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

In words, any accumulated wealth (W_t) that is not used to finance consumption (C_t) will add to next period's wealth according to next period's gross rate of return ($R_{w,t+1}$).

To see the economic implications of this identity, consider a log-linear approximation to equation (1). By imposing a transversality condition⁶ and taking expectations we obtain:

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

where ρ_w is positive but less than one, and equals one minus the exponential of the average log consumption-wealth ratio. α is a constant. Throughout the paper, lower-case letters represent variables expressed in natural logarithms.

The implication of equation (2) is that the current log consumption-wealth ratio should reflect rational forecasts of future returns on wealth and consumption growth. Intuitively, a household can only afford to consume in excess of its current wealth, *ex ante*, if its expected future wealth returns more than offset its expected future consumption growth. In this framework, the household will respond to an expected increase in wealth by consuming more today or by increasing its planned future consumption growth.

Log aggregate wealth w_t in equation (2) contains the expected discounted value of future labor income, and is therefore unobservable. Hence, we cannot directly test equation (2) empirically. Lettau and Ludvigson (2001) mitigate this issue by providing a set of assumptions that link unobservable aggregate wealth W_t to observable series on income and asset wealth.⁷ Firstly, aggregate wealth is equal to the sum of human wealth (HU_t) and observable asset wealth (A_t). As long as the share of each of these series in aggregate wealth is stationary, we can write the log-linear approximation of that equality as:

⁶ The transversality condition implies that the log consumption-wealth ratio is zero in the limit. This ensures that neither consumption nor wealth end up becoming an infinite fraction of the other variable.

⁷ In Section 6 we will consider an alternative approach to examining the implications of the household budget constraint developed by Whelan (2008). The advantage of this alternative approach is that the relationship predicted by the model involves only observable variables, and hence there is no need to make any assumptions about unobservable variables before testing the model.

$$w_t = \omega_y hu_t + \omega_a a_t \quad (3)$$

where ω_y and ω_a are the steady-state shares of human and asset wealth, respectively, in total lifetime resources. Secondly, Lettau and Ludvigson (2001) write human wealth as a function of discounted current and expected future income, and discuss that a log-linear approximation of that equation links log human wealth to the log of contemporaneous labor income (y_t), as follows:

$$hu_t = \nu + y_t + z_t \quad (4)$$

where ν is a constant and z_t is a mean-zero stationary random variable. According to equation (4), current income captures the non-stationary component of human wealth.

Combining equations (3) and (4), we see that the log of aggregate wealth is a linear combination of log current observable asset wealth a_t , log income y_t and a term proportional to the stationary variable z_t . Given that information one can derive the following approximate equivalent to equation (2):

$$c_t - \tau - \omega_y y_t - \omega_a a_t = \eta_t \quad (5)$$

where τ is a constant and η_t is a residual, which should be stationary for reasons we are about to explain. Given the foregoing assumptions, the residual η_t is a function of expectations of future consumption growth, of future income growth and of future returns to observable asset wealth. Assuming these variables are all stationary, η_t should be stationary. As long as the levels of consumption, labor income, and observable net asset wealth are non-stationary, η_t can be interpreted as a cointegrating residual. In that

case, equation (5) captures the long-run relation between consumption, income, and wealth. As we are about to discuss, we will estimate this equation in modified form.

2.2 Distinguishing between housing and financial wealth

The framework underlying equation (5) implicitly assumes that a single coefficient captures the long-run elasticity of the different components of household wealth. In our paper, we relax that assumption by estimating separate elasticities for housing wealth and financial wealth. In doing so, we follow the papers by Case, Quigley and Shiller (2005) and Carroll, Otsuka and Slacalek (2006).

To generalize the model in this way, we append the estimation framework with an additional log-linear approximation to an identity for observable asset wealth:

$$a_t = \psi_{nh} nh_t + \psi_{nf} nf_t \quad (6)$$

where net housing wealth nh_t is gross housing wealth minus mortgage debt, net financial wealth nf_t is gross financial wealth minus non-mortgage liabilities, and ψ_{nh} and ψ_{nf} are their respective steady-state shares in net asset wealth. This yields the following long-run consumption function:

$$c_t - \tau - \omega_y y_t - \omega_{nh} nh_t - \omega_{nf} nf_t = \eta_t \quad (7)$$

where $\omega_{nh} = \psi_{nh} \omega_a$ is the share of net housing wealth in total lifetime resources, and $\omega_{nf} = \psi_{nf} \omega_a$ is the corresponding share for net financial wealth.

This equation is part of the model which we will estimate in section 4. In theory, the coefficients equal the real steady-state shares of each form of wealth. This implies that theoretically speaking, the long-run elasticities can only differ across wealth types by

virtue of differences in the sizes of the stocks of the different types of wealth. That is, in theory each type of wealth has the same dollar-for-dollar impact on consumption, i.e. it has the same marginal propensity to consume (MPC). Empirically however, MPCs are likely to differ across wealth types. This implies that the coefficient estimates from equation (7) may well be different from the steady-state wealth ratios.

There are a number of reasons why housing wealth and financial wealth could affect consumption with different MPCs. The first key difference between housing and financial wealth is that households not only use their homes to store wealth, but also obtain housing services from them. This means that homeowners can experience both a wealth gain and an increase in the implicit cost of consuming housing services when house prices rise. The net effect on wealth will vary across households depending on their position in the housing market. For example, renters will experience a decrease in wealth, while those homeowners looking to decrease their consumption of housing services will experience an increase in wealth. The size and sign of the aggregate wealth effect 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 link between housing wealth and consumption. On the other hand, there are reasons for housing wealth to have a stronger impact on consumption. For instance, households may believe that house values are more persistent than the prices of financial assets, such that a given change in housing wealth will have a relatively large impact on expected lifetime resources, and therefore on consumption.

2.3 Incorporating an explicit role for household debt

In the equations above, the level of household debt played only an implicit role as a component of household net worth. In this section, we consider a variant of equation (7) in which household debt enters separately from gross wealth. We call this the gross wealth model, as opposed to the net wealth model which we have discussed so far. Entering debt as a separate variable allows us to explicitly track how gross wealth and income are related to debt levels. As we will discuss, this will help us to assess whether factors such as liquidity constraints or precautionary saving behavior are present. Since these factors tend to increase the marginal propensity to consume out of wealth,⁸ knowing whether they are present is important in order to interpret the estimated wealth elasticities in the consumption equation.

This time, we generalize equation (5) by writing total net asset wealth as the sum of gross housing wealth and gross financial wealth minus household debt, corresponding to the following log-linear approximation:

$$a_t = \psi_h h_t + \psi_f f_t + \psi_d d_t \quad (8)$$

where h_t is gross housing wealth, f_t gross financial wealth, and d_t stands for household debt. ψ_h and ψ_f represent the steady-state ratios of gross housing wealth and gross financial wealth, respectively, to total net asset wealth, and ψ_d represents the *negative* of the corresponding steady-state ratio for household debt.

Substituting this in equation (5) yields:

$$c_t - \tau_1 - \omega_y y_t - \omega_h h_t - \omega_f f_t - \omega_d d_t = \eta_{1t} \quad (9)$$

⁸ Carroll (2001) makes this point for precautionary saving, and states that the precautionary saving motive can generate behavior that is virtually indistinguishable from that implied by liquidity constraints.

where $\omega_h = \psi_h \omega_a$, $\omega_f = \psi_f \omega_a$, and $\omega_d = \psi_d \omega_a$ are the contributions of gross housing wealth, gross financial wealth, and household debt, respectively, to total lifetime resources. In theory, ψ_d and ω_d are negative. Theory therefore predicts that changes in the level of debt are, in the long run, negatively related with consumption. The intuition behind this is that household debt allows households to bring their consumption forward at the cost of paying interest on the debt, which means that higher long-run debt levels leave fewer resources available for consuming.

In section 5, we find empirical evidence for the existence of two cointegrating relations in the gross wealth model. Since we only detect a second long-run relation once we separate debt from gross asset wealth, we interpret the second equation as capturing a previously implicit relationship between precisely those variables. As we will explain in section 5, we normalize the second relation such that it expresses household debt as a linear function of gross household wealth and income:

$$d_t - \tau_2 - \theta_y y_t - \theta_h h_t - \theta_f f_t = \eta_{2t} \quad (10)$$

The coefficients θ_y , θ_h and θ_f represent the percentage effect of permanent changes in income, the aggregate value of the housing stock and the value of households' financial assets on the long-run level of household debt.

The purest form of the permanent income hypothesis says that an unexpected permanent increase in wealth leads households to increase consumption by the annuity value of the wealth increase, but is silent about the implications for household debt. As we are about to explain, modifying the life cycle model by accounting for liquidity constraints or incorporating a precautionary saving motive can imply that asset wealth increases tend to increase household borrowing. Equation (10) can be seen as a test of whether the data reflect any of these mechanisms. We discuss three reasons why a permanent increase in asset wealth may be associated with higher aggregate trend debt levels.

The first reason is the possibility that risk-averse households engage in precautionary saving, as formalized by Carroll (1997). We interpret that theory as meaning that, in the face of uncertainty about future income and asset values, households desire to hold a buffer stock of wealth in order to mitigate the risk of having to consume very little in case of a low-probability large negative shock to income or wealth. An unexpected increase in a household's wealth or income will increase the value of its buffer stock, and thereby relax the need for precautionary saving. One possible way in which the household can adjust its net worth back down towards its optimal level is by borrowing more to finance consumption. Therefore, a permanent increase in aggregate asset values can allow households to permanently run higher debt levels.

Another reason why aggregate household liabilities may be positively related to asset wealth lies in the likelihood that at any point of time, there is bound to be a fraction of the population which is liquidity-constrained. Given imperfect information about borrower characteristics, banks typically require borrowers to provide collateral to insure against default risk. When the value of households' collateralizable assets increases, banks will tend to be willing to lend more, or at more generous terms. Those households which were credit-constrained before the wealth increase are then likely to borrow more in order to bring consumption forward. In the aggregate, higher gross wealth levels may therefore imply higher average debt levels. In New Zealand as in other countries, almost all household debt is secured against housing wealth, such that this collateral effect likely applies almost exclusively to changes in housing wealth.

The collateral effect as described above primarily applies to existing homeowners. That effect may work in tandem with a mechanism known as passive equity withdrawal, which applies primarily to new home buyers. If house prices are high today relative to the past, a new home buyer will tend to take on a larger mortgage in order to finance a house purchase, while the sale might allow the previous owner to pay off the outstanding amount on a smaller mortgage. In a lifecycle sense, this mechanism implies that a house

price increase increases debt levels of younger cohorts, where the increase in cohort-specific debt levels makes its way through the age distribution as time progresses. In conclusion, passive equity withdrawal would imply that a permanent increase in asset prices has a long-run effect on household debt.

3 Data

We use the following data: real per-capita total household consumption, household wealth and its components, and after-tax labor income,⁹ for the period 1990Q1-2006Q1. In this section, we summarize substantive aspects of our data choices and characterize the data. We refer to the appendix of De Veirman and Dunstan (2008) for details on data sources.

Throughout, we have ensured that the estimated equations are consistent with the aggregate implications of the household budget constraint, in line with the recommendations of Rudd and Whelan (2006). For example, our measure of household consumption includes spending on both non-durable and durable goods. That is, we treat the purchase of durable goods as consumption rather than as wealth accumulation. To be consistent, we measure wealth net of durable goods. A second point in a similar vein is that, to preserve the relationship between nominal quantities inherent in the budget constraint, it is important that we use the same index to deflate all variables. We use the CPI index for this purpose.

Figures 1 and 2 graph the main series which we use in this paper. All series are in logarithms, and in real per-capita terms.

⁹ After-tax labor income differs from total disposable income in that it excludes property income. According to the intertemporal budget constraint, income flows from the stock of wealth -capital gains as well as dividends and interest income- are captured by the rate of return on wealth, not by the measure of income which enters human wealth.

Figure 1 Consumption and labor income




Figure 2 Components of household wealth

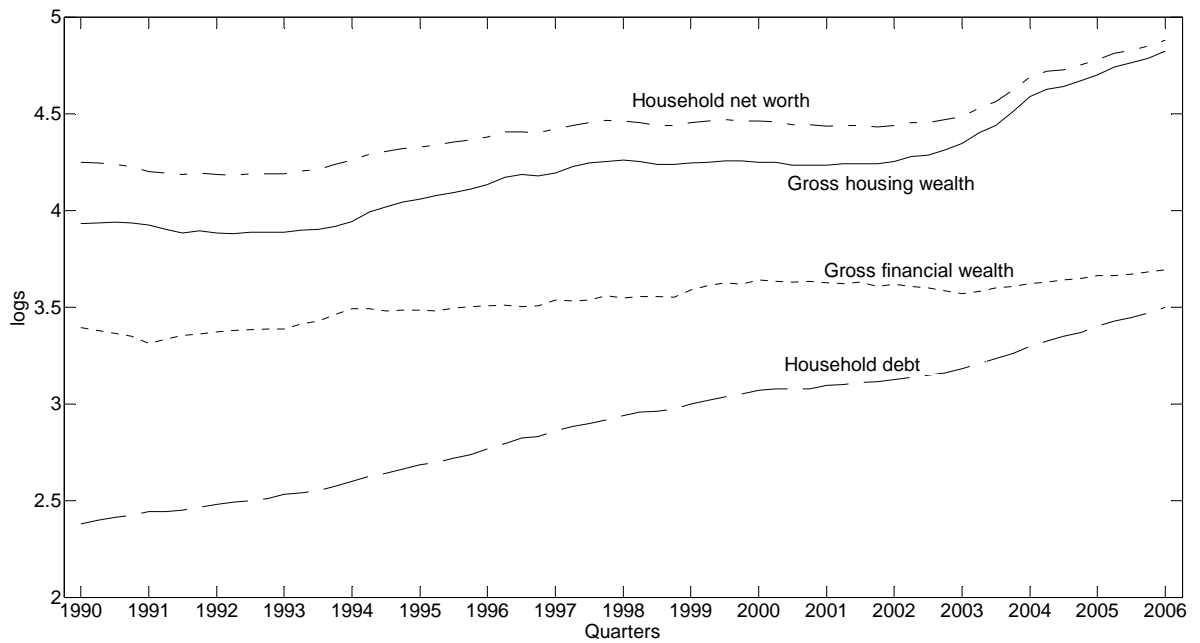


Figure 1 plots consumption and labor income. Consumption growth has been strong for most of our sample, a trend which has occurred in the context of a widening gap between consumption and labor income.

Figure 2 shows household net worth, along with its components: gross housing wealth, gross financial wealth and household liabilities. In New Zealand, housing wealth represents about two thirds of total household assets. In our sample, the share of housing in total assets has never been below 60%. As a result of the rapid increase in housing wealth since 2001, housing wealth constitutes about 75% of total gross wealth towards the end of the sample.

Furthermore, note that financial wealth has not trended upward as strongly as the other components of net worth. Also, financial wealth growth has not been particularly volatile over the sample in terms of standard deviation (see below). One reason for this is that New Zealanders have relatively little direct exposure to the stock market. On average over our sample, direct equity holdings constitute only 14% of financial assets, and 5% of total assets.

Finally, note that total household liabilities have increased more steeply throughout the sample than asset wealth has. Correspondingly, household leverage rates have increased: the ratio of household liabilities to total assets has gradually increased from 14% in 1990 to 22% in 2006.

Table 1 presents summary statistics for the quarterly growth rates in real per-capita consumption, gross housing wealth, gross financial wealth, household debt, and income. In New Zealand, quarter-on-quarter consumption growth is relatively volatile, implying that the gap between the volatility of wealth and consumption is relatively small. Furthermore, notice that financial wealth is somewhat less volatile than housing wealth. In terms of first-order serial correlation, housing wealth growth is more persistent than

financial wealth growth, while there appears to be very little serial correlation in consumption growth.

From the same table, note that consumption growth is strongly correlated with housing wealth and household debt, more so than with financial wealth or labor income. The strongest correlation is that between housing wealth growth and household debt. In the remainder of this paper, we examine which economic relationships these correlations reflect.

Table 1 Summary statistics								
	Mean (%)	Standard deviation	Serial correlation	Correlation matrix				
				Δc	Δh	Δf	Δd	Δy
Consumption growth (Δc)	0.38	0.008	0.06	1	0.43	0.15	0.38	0.12
Housing wealth growth (Δh)	1.39	0.020	0.72		1	0.07	0.68	0.25
Financial wealth growth (Δf)	0.46	0.014	0.24			1	0.25	0.10
Household debt growth (Δd)	1.77	0.008	0.51				1	0.16
Income growth (Δy)	0.23	0.012	-0.15					1

Note: All of the series are defined as quarterly log differences. The degree of serial correlation is computed from an OLS regression of growth in each series against its lag and a constant.

4 Estimates of the net wealth model

4.1 Existence of a long-run relationship

We now turn to the estimation of the net wealth model, which involves the long-run relationship (7) between consumption, net housing wealth, net financial wealth, and

income. For the concept of a long-run relationship to be meaningful, the variables appearing in the equation should be non-stationary. Table A.1 in the appendix presents the results from Augmented Dickey-Fuller (ADF) tests for stationarity using various orders of augmentation in the testing regression. At the orders of augmentation chosen by the Akaike and Schwarz information criteria, we cannot reject the null hypothesis of a unit root for any of the four variables that potentially enter the net wealth model's long-run relationship.

Table A.2 in the appendix presents results from tests for cointegration among these four variables. First, we implement Engle-Granger cointegration tests. This procedure consists of applying ADF tests to the residuals of a static OLS regression of consumption on income, net housing wealth and net financial wealth. At the order of augmentation selected by the Akaike and Schwarz information criteria, we reject the null hypothesis of a unit root in these residuals at the 5% level, which suggests cointegration. 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, captured by equation (11) below. Both the L-max test and the Trace test reject the null hypothesis of no cointegration at the 5 percent level. In sum, we find substantial evidence for the existence of the cointegration relationship hypothesized in equation (7).

4.2 Estimates of the long-run relation and short-run dynamics

We just provided evidence that equation (7) can be interpreted as a cointegration relationship, which is consistent with the implications of the intertemporal budget constraint. We now move to estimating this relationship and to characterizing convergence to the long-run equilibrium relation. On the latter point, note that the log-linearized budget constraint, or the cointegration relation equation (7) which is derived from it, impose few if any restrictions on how equilibrium is restored. Cointegration only implies that if the cointegrating residual η_t is positive, indicating that households are currently running a high consumption-to-wealth ratio, then future wealth gains need

to offset future consumption growth. However, it does not impose any restriction on the relative role of the four cointegrating variables in restoring equilibrium. High levels of consumption relative to contemporaneous wealth and income may tend to be adjusted for by slowdowns in consumer spending. Yet high consumption-wealth ratios can just as well systematically precede higher wealth or income growth, as would be the case if relatively high spending levels tend to reflect expectations of future wealth or income gains which are later confirmed.

To maintain this level of generality, we use a Vector Error Correction Model (VECM) which in principle allows any or all of the cointegrated variables to play a role in correcting for deviations from the long-run equilibrium relation. We estimate the following VECM:

$$\Delta X_t = A_0 + A(L)\Delta X_t + \alpha\beta' X_{t-1} + \varepsilon_t \quad (11)$$

where $X_t = (c_t, y_t, nh_t, nf_t)'$ contains data on consumption, labor income, and net housing and financial wealth, A_0 is a vector of constants, $A(L)$ is a polynomial in the lag operator, β is a vector with estimated coefficients from the cointegration relation, such that $\beta' X_{t-1} = \eta_{t-1}$, α is a vector of coefficients capturing whether and how strongly each variable corrects for any deviation from equilibrium, and ε_t is a vector of residuals.

The results reported in this paper are based on estimating the VECM jointly with the long-run relationship using the Johansen FIML estimator.¹⁰ We estimate a VECM of order one, based on the fact that the Akaike and Schwarz information criteria suggest using two lags in the corresponding VAR in levels.

¹⁰ To check for robustness, we also estimated the long-run relationship using the Dynamic OLS estimator, after which we estimated a VAR model imposing this estimated long-run relationship, as in Lettau and Ludvigson (2004). Results, available upon request, were very similar. We report results from the Johansen estimator because in section 5 we will estimate a model with multiple cointegrating vectors, which is not possible with the Dynamic OLS estimator.

The top portion of table 2 shows the estimated long-run relationship. The long-run income elasticity is estimated to be 0.68. A permanent one-percent increase in real per-capita net housing wealth is associated with a 0.09 percent increase in real per-capita consumption in the long run. The corresponding effect for net financial wealth is 0.19 percent. Each of these elasticities is statistically significant at the 5% level.

Note that the long-run elasticity from financial wealth is larger than that for housing wealth,¹¹ even though financial wealth accounts for a substantially smaller share of total net worth. This can only be true if the marginal propensity to consume (MPC) from financial wealth is particularly large relative to the MPC from housing wealth. We compute approximate MPCs from housing and financial wealth by multiplying the estimated elasticities by the average ratio of consumption to the relevant type of wealth. This calculation suggests that a permanent one-dollar increase in net housing wealth translates to a long-run increase in consumption of 3.34 cents, while the MPC from net financial wealth is particularly large at 11.36.

In this paper, we focus on the estimates from the VECM framework since this model allows us to investigate whether wealth responds to disequilibria, and to quantify how important transitory wealth variation is. De Veirman and Dunstan (2008) also apply two techniques which do not hinge on stability in the long-run relationship, but do yield estimates of the long-run MPCs from both wealth types. Both an empirical model inspired by Carroll, Otsuka and Slacalek (2006) and a technique due to Aron and Muellbauer (2006) confirm that the MPC from financial wealth is two to three times larger than the MPC from housing wealth.

¹¹ This difference is not statistically significant at the 5% level.

Table 2 Cointegration and VECM estimates for the net wealth model

Long-run relationship					
	Consumption	Net housing wealth	Net financial wealth	Income	Constant
Estimate	1	0.09** (5.12)	0.19** (4.10)	0.68** (8.78)	0.42

Short-run dynamics				
	Equation			
	Δc_t	Δnh_t	Δnf_t	Δy_t
Constant	0.002* (1.92)	0.001 (0.74)	0.003* (1.69)	0.000 (0.44)
$\hat{\eta}_{t-1}$	-0.27** (-3.07)	0.55** (2.99)	0.01 (0.09)	0.13 (0.92)
Δc_{t-1}	-0.07 (-0.56)	0.47* (1.86)	-0.41* (-1.87)	0.13 (0.92)
Δnh_{t-1}	0.16** (3.54)	0.56** (5.86)	0.02 (0.28)	0.25 (1.32)
Δnf_{t-1}	0.01 (0.08)	0.48** (3.28)	0.23* (1.80)	0.05 (0.64)
Δy_{t-1}	0.01 (0.08)	0.18 (0.91)	0.40** (2.41)	0.10 (0.94)
R^2	0.30	0.63	0.20	0.12

Note: The top schedule of this table reports cointegration estimates for the net wealth model, while the bottom schedule reports the corresponding VECM estimates. t-statistics are in parentheses; * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level.

The lower part of table 2 presents the estimated short-run dynamics of the model. In both the consumption and the housing wealth equations, the adjustment coefficient on the cointegrating residual η_{t-1} is statistically significant at the five percent level. All other things equal, consumption tends to fall and housing wealth tends to rise in response to a positive cointegrating residual. Therefore, both consumption and housing wealth tend to restore equilibrium, i.e. they error-correct. In contrast, neither financial wealth nor labor income respond to the error-correction term in a statistically significant fashion.

As for further dynamics, note that consumption growth depends positively and significantly on lagged housing wealth growth, but is only weakly related with lagged growth in financial wealth and income. All other things equal, this implies that consumption will respond faster to a given change in housing wealth than implied by the estimated error-correction parameter. Finally, note that housing wealth is more strongly associated with its own lag than any of the other variables are, which mirrors our finding in section 3 that housing wealth growth is the most persistent series in terms of first-order autocorrelation.

4.3 Permanent-transitory decomposition

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. Doing so will help us assess whether most changes in wealth are related with consumption changes, a question which Lettau and Ludvigson (2004) have answered negatively for the United States.

We perform the permanent-transitory decomposition using the approach developed by Gonzalo and Ng (2001). We construct a matrix G that transforms the reduced-form VECM residuals ε_t from equation (11) into a set of permanent and transitory shocks u_t :

$$u_t = G \varepsilon_t \tag{12}$$

Gonzalo and Ng (2001) show that G depends on the short-run adjustment parameters α to the cointegrating residual and on the coefficients β of the long-run relation:

$$G = \begin{bmatrix} \alpha_{\perp}' \\ \beta \end{bmatrix} \tag{13}$$

The matrix α_{\perp} is orthogonal to α such that $\alpha'_{\perp}\alpha = 0$. In the net wealth model, ε_t and u_t are each four-by-one vectors. We found evidence for only one cointegration relation, which implies that there is a single transitory shock, appearing as the fourth entry of the vector u_t , while the first three entries of u_t represent three permanent shocks. As recommended by Gonzalo and Ng (2001), we set insignificant adjustment parameters (those for net financial wealth and income) to zero for performing the permanent-transitory decomposition.

For each variable in the model, table 3 presents the decomposition of the h -quarter ahead forecast error variance into the contribution from transitory and permanent shocks, labeled ‘ T ’ and ‘ P ’ respectively. The table reports the results for two cases: orthogonalized and unorthogonalized shocks. The results with unorthogonalized shocks follow directly from decomposing shocks according to equation (12). For every variable, this case includes a column labeled P,T which captures the contribution of the covariances between the permanent and transitory shocks. The results labeled ‘orthogonalized shocks’ are obtained by using equation (12) first, and then applying a Choleski decomposition to the vector u_t , where the Choleski order is such that the transitory shock does not affect any of the permanent shocks contemporaneously.¹²

For consumption, the transitory shock explains a relatively large portion of the forecast error variance. Transitory variation is also quite important in the case of housing wealth. On the other hand, the forecast errors for financial wealth and income have been almost uniquely due to permanent shocks.

¹² The results therefore provide a lower bound for the portion of the variance explained by the transitory shock.

Table 3 Transitory-permanent decomposition for the net wealth model


Orthogonalized shocks												
	$\Delta c_{t+h} - E_t \Delta c_{t+h}$		$\Delta h_{t+h} - E_t \Delta h_{t+h}$			$\Delta f_{t+h} - E_t \Delta f_{t+h}$			$\Delta y_{t+h} - E_t \Delta y_{t+h}$			
Horizon	<i>T</i>	<i>P</i>	<i>T</i>	<i>P</i>		<i>T</i>	<i>P</i>		<i>T</i>	<i>P</i>		
1	26	74	26	74		0	100		0	100		
4	28	72	13	87		1	99		0	100		
8	28	72	11	89		1	99		0	100		
12	28	72	10	90		1	99		0	100		
∞	28	72	10	90		1	99		0	100		

Unorthogonalized shocks												
	$\Delta c_{t+h} - E_t \Delta c_{t+h}$			$\Delta h_{t+h} - E_t \Delta h_{t+h}$			$\Delta f_{t+h} - E_t \Delta f_{t+h}$			$\Delta y_{t+h} - E_t \Delta y_{t+h}$		
Horizon	<i>T</i>	<i>P</i>	<i>P,T</i>	<i>T</i>	<i>P</i>	<i>P,T</i>	<i>T</i>	<i>P</i>	<i>P,T</i>	<i>T</i>	<i>P</i>	<i>P,T</i>
1	155	116	-170	73	187	-160	0	100	0	0	100	0
4	152	107	-158	40	148	-87	5	84	11	0	100	0
8	151	107	-158	36	134	-70	5	84	11	0	100	0
12	151	107	-158	35	131	-66	5	84	11	0	100	0
∞	151	107	-158	35	130	-65	5	84	11	0	100	0

Note: This table reports the percentage of the forecast error variance at horizon h that is attributable to permanent (P) vs. transitory (T) shocks. The top half of the table presents results from applying the Gonzalo-Ng decomposition followed by a Choleski decomposition assuming no contemporaneous effect of the transitory shock on the permanent shocks. The bottom half of the table presents results from applying the Gonzalo-Ng decomposition without imposing any further identification on the shocks. In that case, we also list the contribution of the covariance (P,T) between the permanent and transitory shocks to the total forecast error variance.

Our finding that there is non-trivial transitory variation in consumption and housing wealth is in line with the fact that both variables error-correct. To better understand the role which consumption and housing wealth play at the occasion of a transitory shock, figure 3 plots the VECM impulse responses to a one standard deviation transitory shock.

Figure 3 Responses to one standard deviation transitory shock in net wealth model



The transitory shock implies that on impact, housing wealth rises 6 percent above its steady-state level and consumption falls by around 3 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, but starts to decline after about a year. In contrast, consumption immediately starts to revert back up to its steady-state growth path. In line with the estimated VECM adjustment parameters, this scenario shows that after the shock, both consumption and housing wealth play an important role in bringing the error-correction term back to zero.

Figure 3 omits the income response. It suffices to say that according to a similarly scaled impulse-response graph for income, there is no visible change in income associated with the transitory shock.

Note that according to the impulse-responses, a transitory shock implies particularly large variation in housing wealth, and somewhat smaller changes in consumption. On the other hand, the variance decomposition suggested that transitory shocks matter most for consumption and somewhat less so for housing wealth. This is explained by the fact that the overall error variance of housing wealth is a multiple of the error variance for consumption, implying that transitory changes in housing wealth account for a relatively small share of its total error variance.

4.4 Lessons from the net wealth model

For the United States, Lettau and Ludvigson (2004) found that wealth movements were mainly transitory and unrelated to changes in consumption, which were mainly permanent. Based on these findings, they concluded that their estimated long-run relation greatly overstated the true effects of a typical wealth change.

In the present paper, we find, in an environment where stock market cycles are relatively unimportant factors behind variation in household wealth, that most of the variation in wealth has been permanent. Therefore, our estimated long-run relationship adequately captures the effect of a typical wealth change.

Unlike in the United States, we find that there is quite some transitory variation in consumption, and that consumption error-corrects to short-term discrepancies between consumption and wealth. Therefore, permanent changes in wealth will tend to be followed by changes in consumption. Combined with our finding that most variation in wealth is permanent, this completes the picture that consumption responds to a typical change in wealth.

While permanent shocks account for most of the variation in housing wealth, we do find that there is non-trivial transitory variation in housing wealth. This reflects the fact that housing wealth adjusts to short-run discrepancies between wealth and consumption.

Arguably, endogenous changes in wealth such as the ones which we detect in the data are not the focus of lifecycle theory. Typically, the ‘wealth effect’ is interpreted as the effect of unexpected wealth movements (or changes in expectations regarding future wealth movements) on consumption, while endogenous changes in wealth are already incorporated in the households’ spending plan.

However, our finding that housing wealth responds to short-run deviations from equilibrium does not necessarily imply a rejection of the permanent income hypothesis. Relatively standard lifecycle theory would imply that changes in expectations regarding future income or wealth movements imply a change in permanent income, and therefore affect consumption. Empirically, it is likely that expectations about the future also affect house prices, implying endogenous variation in housing wealth other than that driven by households’ decisions on the quantity of housing services they wish to purchase.

Our VECM evidence suggests that housing wealth growth is particularly persistent, while there is virtually no intrinsic persistence in consumption growth. To the extent that this difference is a structural feature of the economy (rather than being due to a difference in the persistence of shocks affecting these variables), changes in expectations about the future will tend to affect consumption relatively quickly, but imply a delayed response on the part of housing wealth. Therefore, such changes in consumption will tend to predate changes in housing wealth, where the latter appear to correct for short-run discrepancies between consumption and wealth.

5 Estimates of the gross wealth model

5.1 Existence of two long-run relationships

In this section, we explicitly model gross housing wealth, gross financial wealth and household debt as three separate variables, rather than using net housing and financial

wealth as in the previous section. The error-correction system is as in equation (11), except for the fact that the data vector X_t , now contains five entries: consumption, income, gross housing wealth, gross financial wealth, and household debt.

Table A.3 in the appendix documents the results from Johansen L-max and Trace tests for the existence and number of cointegrating vectors in the gross wealth model. We find evidence for two, but not more than two, cointegrating vectors at the 5 percent significance level. We will discuss our interpretation of the two long-run relations below.

5.2 Estimates of the long-run relations and short-run dynamics

Recall that there was a single cointegrating vector in the net wealth model. It required one normalization for identification: we normalized the vector such that the coefficient on log consumption equaled one. Since there are two cointegrating vectors in the gross wealth model, each long-run equation now needs two normalizations. To be concrete, we repeat the gross wealth model's first cointegrating relation below:

$$c_t - \tau_t - \omega_y y_t - \omega_h h_t - \omega_f f_t - \omega_d d_t = \eta_{1t} \quad (9')$$

This equation normalizes the consumption coefficient to one. We refer to it as the consumption equation. In addition, we require that the elasticities from income and the components of wealth sum to one, i.e. $\omega_y + \omega_h + \omega_f + \omega_d = 1$. Recall from section 2 that in theory, the coefficients are wealth shares (or rather wealth contributions since $\omega_d \leq 0$), which does imply their summing to one. In addition, the estimated elasticities on wealth and income did approximately sum to one in the net wealth model.

We only find evidence for a second cointegrating relation once we separate debt from gross asset wealth. This suggests that the second relation captures a previously implicit

relation between precisely those variables. We normalize the second cointegrating vector accordingly: we require the coefficient on debt to equal one, and the coefficient on consumption to equal zero. We thereby obtain equation (10), which we restate below:

$$d_t - \tau_2 - \theta_y y_t - \theta_h h_t - \theta_f f_t = \eta_{2t} \quad (10')$$

That is, the second long-run relation expresses debt in terms of gross wealth and income.

The top part of table 4 shows the estimates of the two cointegrating relations. The long-run elasticities in the consumption equation are qualitatively similar to their analogues in the net wealth model. The same is true for the implied approximate MPCs: a one-dollar increase in gross financial wealth is associated with a 12.5 cents rise in consumption, while a one-dollar increase in gross housing wealth translates into a 3.8 cent consumption increase.

In section 2, we discussed that in theory, debt should enter with a negative coefficient in the consumption equation, reflecting the fact that higher long-run debt levels imply higher costs of debt servicing. The estimated sign is indeed negative. In terms of the implied dollar-for-dollar effect, a permanent one-dollar increase in household debt implies a 6 cent decrease in consumption in the long run.

The results of the long-run debt equation suggest that permanent increases in wealth and income imply statistically significant increases in long-run debt levels. The percentage effects from housing and financial wealth are similar: a one percent increase in either housing or financial wealth leads to an increase in debt by about two third of a percent. A change in income is associated with a roughly proportional change in debt.

The results for the debt equation suggest factors such as liquidity constraints, a precautionary saving motive or passive equity withdrawal are present in the data.

Table 4 Cointegration and VECM estimates for the gross wealth model

Long-run relationships						
	Consumption	Housing	Financial	Debt	Income	Constant
Consumption equation	1	0.12** (3.76)	0.21** (3.46)	-0.05** (2.07)	0.72** (12.12)	0.57
Debt equation	0	0.69** (6.24)	0.62** (2.56)	1	1.14** (2.71)	3.85

Short-run dynamics					
	Equation				
	Δc_t	Δh_t	Δf_t	Δd_t	Δy_t
Constant	0.000 (0.25)	0.000 (0.15)	0.002 (0.34)	0.014** (6.55)	-0.004 (-1.13)
$\hat{\eta}_{1t-1}$	-0.28** (-3.30)	0.39** (2.75)	0.01 (0.06)	0.11 (1.55)	0.08 (0.69)
$\hat{\eta}_{2t-1}$	0.04** (2.15)	0.04 (1.47)	0.05 (1.53)	-0.03** (-2.15)	0.08** (2.97)
Δc_{t-1}	-0.12 (-0.98)	0.24* (1.17)	-0.47** (-2.2)	-0.08 (-0.8)	0.10 (0.57)
Δh_{t-1}	0.21** (3.07)	0.59** (5.18)	0.03 (0.23)	0.25** (4.48)	0.02 (0.22)
Δf_{t-1}	0.01 (0.16)	0.43** (3.45)	0.21* (1.6)	0.21** (3.43)	0.09 (0.77)
Δd_{t-1}	0.04 (0.24)	0.12 (0.42)	0.14 (0.45)	-0.04 (-0.26)	0.35 (1.32)
Δy_{t-1}	-0.02 (-0.23)	0.10 (0.62)	0.36** (2.21)	0.06 (0.75)	-0.21 (-1.53)
R^2	0.35	0.68	0.23	0.55	0.25

Notes: This table reports VECM estimates. t-statistics are shown in parentheses; * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level.

The bottom half of table 4 shows the estimated short-run dynamics. The estimates are qualitatively similar to those in the net wealth model, except of course for the

coefficients governing adjustment to the residual from the long-run debt relation, and for the estimated short-run equation for debt growth.

The adjustment parameters suggest that debt and income error-correct to the second cointegration residual. Consumption also responds significantly to discrepancies in the debt-income relation, although it does not directly affect error-correction since it does not appear in that long-run relation. These findings suggest that when debt is high relative to contemporaneous income, households tend to gradually build down debt in order to conform to the debt levels that they can run in the long run given wealth and income. It also suggests that periods of relatively high debt tend to predate income (and consumption) increases, consistent with the hypothesis that households borrow in anticipation of future income and associated consumer spending.

The only short-run equation for which the constant is significantly different from zero is that for debt growth. This suggests that (in a linear model) the trend increase in debt over our sample cannot fully be captured by gross wealth and income changes. This is consistent with the fact that leverage ratios have gradually increased over our sample.¹³

5.3 Permanent-transitory decomposition

As with the net wealth model, we use the VECM estimates to rotate the reduced form residuals into a set of permanent and transitory shocks. Given that we now have five variables and two cointegrating vectors, there are three permanent shocks and two transitory shocks.

Table 5 shows the results of the decomposition into orthogonalized permanent and transitory shocks.¹⁴ It remains true that permanent shocks account for the majority of the

¹³ As a robustness test, we entered a linear trend in the long-run debt equation. The trend turned out to be statistically insignificant.

¹⁴ To perform the permanent-transitory decomposition, we have ordered the transitory shock corresponding to the debt equation last, and the transitory shock corresponding to the consumption equation second-to-last. Thus, we assume that the transitory debt shock does not affect the transitory

variation in every variable. Yet the fraction of transitory variation has increased relative to the net wealth model, in particular for income and consumption. The increase in transitory variation for income corresponds to our finding that income adjusts in response to the second cointegrating residual. Furthermore, we find non-trivial transitory variation in household debt, consistent with the finding that debt error-corrects to the second cointegrating residual.

Table 5 Variance decomposition by persistence of shocks (gross wealth model)

Orthogonalized shocks										
	$\Delta c_{t+h} - E_t \Delta c_{t+h}$		$\Delta h_{t+h} - E_t \Delta h_{t+h}$		$\Delta f_{t+h} - E_t \Delta f_{t+h}$		$\Delta d_{t+h} - E_t \Delta d_{t+h}$		$\Delta y_{t+h} - E_t \Delta y_{t+h}$	
Horizon	<i>T</i>	<i>P</i>	<i>T</i>	<i>P</i>	<i>T</i>	<i>P</i>	<i>T</i>	<i>P</i>	<i>T</i>	<i>P</i>
1	42	58	30	70	7	93	13	87	25	75
4	45	55	16	84	10	90	19	81	26	74
8	45	55	14	86	10	90	18	82	26	74
12	45	55	14	86	10	90	17	83	26	74
∞	44	56	14	86	10	90	17	83	26	74

Note: This table reports the percentage of the forecast error variance at horizon h that is attributable to orthogonalized permanent and transitory shocks.

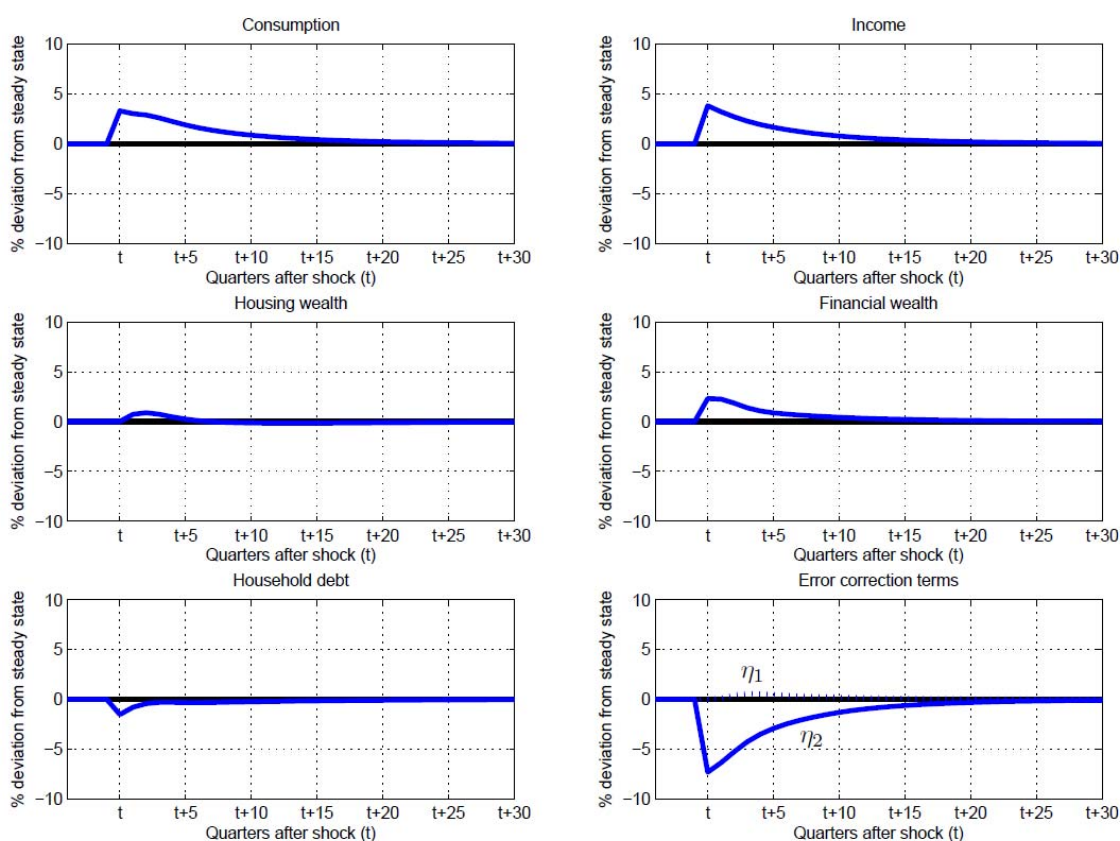
Figure 5 illustrates the extra dynamics introduced by the equation for equilibrium debt by plotting the impulse-responses to a one standard deviation surprise in the transitory shock corresponding to the second cointegrating vector, i.e. that of the debt equation.

On impact, this shock implies a decrease in debt to 1.5 percent below its steady-state path, while consumption, income and financial wealth increase above their steady-state levels. After the shock, debt immediately starts to increase. Consumption, income and

consumption shock contemporaneously. Note that table 5 lists the proportion of the variance explained by these two transitory shocks combined, such that this assumption does not have much of an impact on the reported results.

financial wealth start decreasing immediately but take three to four years to return to their steady-state growth paths. The increase in debt and the decreases in income and wealth tend to restore equilibrium. However, error-correction takes much longer to achieve than in the case we analyzed in the net wealth model. In this case, debt is still out of line with wealth and income by about 1% three years after the shock.

Figure 5 Impulse-responses from a one standard deviation transitory shock associated with long-run debt equation



6 Robustness

In this section, we find evidence for stability in the long-run coefficients over our sample. We also implement an alternative framework for testing the implications of the aggregate budget constraint which does not rely on long-run coefficient stability.

6.1 Testing for long-run coefficient stability

Any application of the cointegration framework assumes that there exists a stable long-run relationship between the cointegrated variables. In truth, there may be structural factors which cause instability in particular over longer horizons. Theoretically, any change in steady-state wealth shares would imply structural change in the long-run coefficients. The question we turn to now is whether we can detect any instability in our sample and if so, whether it invalidates our main conclusions.

We test for a structural break in the long-run relationships using the Sup-LM, Ave-LM and Exp-LM tests with unknown breakpoint developed in Seo (1998).¹⁵ The tests in Seo (1998) have the same asymptotic distribution as the well-known break tests in Hansen (1992), but the Seo tests apply to the maximum likelihood estimation technique which we use in this paper.

Table 6 reveals that both in the net and gross wealth models, neither of the three tests rejects the null hypothesis of stability in the long-run coefficients at the 5 percent level. Note that in the gross wealth model, the test evaluates stability for the two cointegrating relations jointly.

¹⁵ The Sup-LM test tests against the alternative of a single structural break. The Ave-LM and Exp-LM tests concern the alternative that the coefficients follow a martingale, without yielding an estimate for a particular break date. The Exp-LM test has optimal power against alternatives distant from the null hypothesis, while the Ave-LM tests focuses on alternatives near the null hypotheses.

Table 6 Tests for stability in long-run relationships

	Net wealth model		Gross wealth model	
	Test statistic	5 % critical value	Test statistic	5 % critical value
Sup-LM	6.4	14.8	13.9	20.8
Ave-LM	1.9	6.0	5.4	10.2
Exp-LM	1.5	4.2	4.3	6.9

Note: This table reports results from tests for structural change in the estimated cointegrating vectors in the net wealth and gross wealth models, using the Seo (1998) structural break test with unknown breakpoint.

6.2 Testing the implications of the budget constraint without cointegration

Whelan (2008) provides an alternative framework to Lettau and Ludvigson (2004) for testing whether households behave in accordance with the intertemporal budget constraint. One advantage of this framework is that one does not need to estimate any cointegrating coefficients, eliminating the need for assuming long-run coefficient stability. Another advantage is that it does not require any assumptions on how human wealth relates to observable variables. On the other hand, this method does not provide estimates of either the short-run interactions or long-run relations between the variables in the system.

In section 2, we showed how a long-run relation between consumption, wealth and income can be derived from the budget constraint equation (1), which accounted for both human and non-human wealth. The derivation assumed that the non-stationary component of human wealth could be captured by current income. Whelan (2008) instead uses a more standard version of the budget constraint which does not track human wealth:

$$A_{t+1} = R_{a,t}(A_t + Y_t - C_t) \quad (14)$$

Where A_t stands for household net worth including housing and financial wealth, $R_{a,t}$ is the gross time-varying return on assets, Y_t is income and C_t consumption.

In analogy with Campbell and Mankiw (1989), Whelan derives the following equation summarizing the dynamic implications of the budget constraint:

$$xc_t - a_t = E_t \sum_{k=1}^{\infty} \rho_a^k (r_{a,t+k} - \Delta xc_{t+k}) \quad (15)$$

Where xc_t is the log of excess consumption, and excess consumption is defined as the difference between consumption and current labor income. The implications of this equation are much the same as those of equation (2), except that there is no role for human wealth other than current income. Intuitively, the ratio of excess consumption to assets, the left-hand side of equation (15), reflects how much households are eating into their existing assets in any given period. A household can only afford to run a positive ratio of excess consumption to assets if expected future wealth returns more than offset its expected future excess consumption growth.

Equation (15) involves only observable variables. Therefore, its implications can be tested directly without requiring any further assumptions. If equation (15) holds, then the current ratio of excess consumption to net worth $xc_t - a_t$ should on average predict future realizations of the gap between wealth returns and excess consumption growth. We test for that implication by regressing cumulative sums of the gap $r_{a,t+k} - \Delta xc_{t+k}$ on the ratio $xc_t - a_t$ over various forecasting horizons. To construct a measure of the rate of return on assets $R_{a,t}$, we follow Whelan in defining the return on assets implicitly from equation (14):

$$R_{a,t} = \frac{A_{t+1}}{(A_t + Y_t - C_t)} \quad (16)$$

Table 8 presents the results. 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 tend to undo any short-term deviations of that ratio from its sustainable level.

To investigate the relative role of wealth returns and excess consumption growth in restoring the excess consumption ratio, we perform two sets of additional regressions, one in which we investigate the predictive power of the excess consumption ratio for future net worth returns $r_{a,t+k}$, and one where we assess forecasting ability for the negative of excess consumption growth $-\Delta x_{c,t+k}$. Taking 5% significance as the threshold, the results suggest that excess consumption growth plays the most important role in maintaining consistency with the budget constraint. This contrasts with the findings of Whelan (2008) for the United States, which suggested that the ratio of excess consumption to net worth mainly forecasted the returns on the stock of net worth. This mirrors our finding that in the cointegration framework, error-correction by household wealth plays a less important role in New Zealand than in the United States, and that consumption error-corrects unlike in the United States.

Table 8 Long-horizon regressions based on the Whelan framework

	Forecast Horizon, N (quarters)				
	1	4	8	12	20
$Z_{t+k} = (r_{t+k}^a - \Delta xc_{t+k})$					
λ	0.59** (4.05)	1.04** (8.27)	1.06** (9.09)	1.11** (8.34)	1.15** (5.80)
\bar{R}^2	0.31	0.54	0.56	0.58	0.58
$Z = r_{t+k}^a$					
λ	0.04* (1.59)	0.23* (1.89)	0.34* (1.62)	0.29 (1.01)	-0.01 (-0.06)
\bar{R}^2	0.07	0.17	0.12	0.06	-0.02
$Z = -\Delta xc_{t+k}$					
λ	0.54** (3.41)	0.82** (4.23)	0.72** (2.54)	0.82** (2.15)	1.16** (5.80)
\bar{R}^2	0.26	0.32	0.21	0.22	0.43
<i>Notes:</i> This table presents long-horizon regressions of the form:					
$\sum_{k=1}^N Z_{t+k} = \tau + \lambda(xc_t - a_t) + e_t$					
Where N is the forecast horizon.					
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.					

6 Conclusion

Our findings suggest that, in an environment where variation in household wealth is not dominated by stock market cycles, a typical change in wealth is permanent, and that permanent wealth changes imply a lagged consumption response. Therefore, our evidence suggests that a typical change in wealth does matter for consumption.

In appearance, our findings contrast with earlier studies applying similar techniques to other economies, including the United States. In the United States, most wealth variation is transitory and unrelated with consumer spending. We interpret that difference as suggesting that transitory variation in stock market wealth, which is plausibly unrelated with consumption, dominates the observed volatility in household wealth in the United States, and therefore makes it hard to find out whether any other type of wealth change is actually related to consumption.

In future research, it may be worthwhile to filter the wealth data so as to extract variation which is thought to reflect stock market cycles, or for that matter, any type of wealth change which does not appear to be related to substantial consumption changes. In that way, one could obtain a more comprehensive view about which types of wealth changes -e.g. decomposed by frequency of variation- actually do matter for consumption, and which do not.

Furthermore, we show that a model which allows for an explicit role of household debt suggests that debt depends positively on wealth and income. This suggests the presence of factors such as liquidity constraints or precautionary saving behavior, which tend to increase the marginal propensity to consume out of wealth. Finally, the model with debt as a separate variable confirms that even when accounting for two cointegrating relations, we find that most variation in wealth and consumption is permanent.

Appendix: results from tests for stationarity and cointegration

Table A.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**	
Net Housing wealth	0.12	-1.93 ^{A,S}	-2.02	-2.43	-2.26	-2.97	-1.90	
Net 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 values -3.45, -3.15

Note: For each of the four series, the corresponding row reports tests of the null hypothesis that the series is I(1). All of the tests include both a deterministic trend and a constant. 'Lag length' refers to the order of augmentation of the ADF testing regression. The superscripts A and 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. According to the ADF tests, we cannot reject the null hypothesis of I(1) behavior for any of the four series.

Table A.2 Tests for cointegration with the net wealth model

	Lag length						
	0	1	2	3	4	8	12
Residual based ADF test	-5.0** ^{A,S}	-3.5	-3.6	-3.7	-4.3**	-2.6	-1.6
5, 10 percent critical value	-4.16, -3.84						
L-max test for cointegration	39.6**	31.7** ^{AS}	34.2**	31.8**	42.1*	55.6**	97.9**
5 percent critical value	27.58						
Trace test for cointegration	72.8**	57.1** ^{AS}	56.4**	55.1**	77.0**	89.9**	174.9**
5 percent critical value	47.85						

Note: The table reports tests of the null hypothesis that there is no cointegration between consumption, income, net housing wealth and net financial wealth. ‘Lag length’ refers to the order of augmentation for the ADF tests, and the number of lagged differences included in the VECM for the Johansen L-max and trace tests. The superscripts A and S refer 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 (2000). The critical values for the Johansen tests are from MacKinnon, Haug, and Michelis (1999). * indicates significance at the 10 percent level; ** indicates significance at the 5 percent level. At the lag lengths selected by the information criteria, each of the three tests suggests that the variables in the net wealth model are cointegrated.

Table A.3 Johansen L-max and Trace tests for the gross wealth model

L-max test		Trace test	
Alternative hypothesis	p-value	Alternative hypothesis	p-value
$r = 1$	0.06	$r \geq 1$	0.00
$r = 2$	0.01	$r \geq 2$	0.00
$r = 3$	0.19	$r \geq 3$	0.06
$r = 4$	0.16	$r \geq 4$	0.14
$r = 5$	0.24	$r \geq 5$	0.24

Note: The table reports p-values for Johansen L-max and Trace tests in a system containing consumption, income, gross housing wealth, gross financial wealth, and household debt. We report results only for a VECM of lag order 1, which is the order selected by the Akaike and Schwarz information criteria. Taken together, the tests suggest the existence of two, but not more than two, cointegration relations.

References

Ando, Albert and Franco Modigliani (1963): “The Life Cycle Hypothesis of Saving: Aggregate Implications and Tests”, *American Economic Review* 53, 55-84.

Aron, Janine and John Muellbauer (2006): “Housing Wealth, Credit Conditions and Consumption”, Centre for the Study of African Economies Working Paper Series paper 252.

Campbell, John Y. and N. Gregory Mankiw (1989): “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence”, in *NBER Macroeconomics Annual*, University of Chicago Press.

Carroll, Christopher D. (1994): “How Does Future Income Affect Current Consumption?”, *Quarterly Journal of Economics* 109, 111-147.

Carroll, Christopher D. (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis”, *Quarterly Journal of Economics* 112, 1-55.

Carroll, Christopher D. (2001): “A Theory of the Consumption Function, With and Without Liquidity Constraints”, *Journal of Economic Perspectives* 15, 23-45.

Carroll, Christopher D., Misuzu Otsuka and Jirka Slacalek (2006): “How Large is the Housing Wealth Effect? A New Approach”, NBER Working Paper W12746.

Case, Karl E., John M. Quigley and Robert J. Shiller (2005): “Comparing Wealth Effects: the Stock Market versus the Housing Market”, *Advances in Macroeconomics* 5, article 1.

Chen, Jie (2006): “Re-evaluating the Association between Housing Wealth and Aggregate Consumption: New Evidence from Sweden”, *Journal of Housing Economics* 15, 321-348.

Davis, Morris A. and Michael G. Palumbo (2001): “A Primer on the Economics and Time Series Econometrics of Wealth Effects”, Finance and Economics Discussion Series 2001-09, Board of Governors of the Federal Reserve System.

De Veirman, Emmanuel and Ashley Dunstan (2008): “How Do Housing Wealth, Financial Wealth and Consumption Interact? Evidence from New Zealand”, Reserve Bank of New Zealand Discussion Paper 2008/05.

Fernandez-Corugedo, Emilio, Simon Price and Andrew Blake (2003): “The Dynamics of Consumers’ Expenditure: The UK Consumption ECM Redux”, Bank of England Working Paper 204.

Fisher, Lance A., Glenn Otto, and Graham M. Voss (2005): “Housing Booms, Non-Financial Wealth and Consumption: Lessons from the Australian Experience”, mimeo.

Flavin, Marjorie A. (1981): “The Adjustment of Consumption to Changing Expectations about Future Income”, *Journal of Political Economy* 89, 974-1009.

Friedman, Milton (1957): “A Theory of the Consumption Function”, Princeton University Press.

Gonzalo, Jesus and Serena Ng (2001): “A Systematic Framework for Analyzing the Dynamic Effects of Permanent and Transitory Shocks”, *Journal of Economic Dynamics and Control* 25, 1527–46.

Hall, Robert E. (1978): “Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence”, *Journal of Political Economy* 86, 971-987.

Hayashi, Fumio (2000): “Econometrics”, Princeton University Press.

Hamburg, Britta, Mathias Hoffmann and Joachim Keller (2007): “Consumption, Wealth and Business Cycles in Germany”, *Empirical Economics* 34(3), 451-476.

Hansen, Bruce E. (1992): “Tests for Parameter Instability in Regressions with I(1) Processes”, *Journal of Business and Economic Statistics* 10, 45-59.

Lettau, Martin and Sydney Ludvigson (2001): “Consumption, Aggregate Wealth, and Expected Stock Returns”, *Journal of Finance* 56, 815–849.

Lettau, Martin and Sydney Ludvigson (2004): “Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption”, *American Economic Review* 94, 276-299.

MacKinnon, James G., Alfred A. Haug and Leo Michelis (1999): “Numerical Distribution Functions of Likelihood Ratio Tests for Cointegration”, *Journal of Applied Econometrics* 14, 563-577.

Rudd, Jeremy and Karl Whelan (2006): “Empirical Proxies for the Consumption-Wealth Ratio”, *Review of Economic Dynamics* 9, 34–51.

Seo, Byeongseon (1998): “Tests for Structural Change in Cointegrated Systems”, *Econometric Theory* 14, 222-259.

Whelan, Karl (2008): “Consumption and Expected Asset Returns Without Assumptions about Unobservables”, *Journal of Monetary Economics* 55, 1209-1221.