



DP2010/09

Debt Dynamics and Excess Sensitivity of
Consumption to Transitory Wealth Changes

Emmanuel De Veirman and Ashley Dunstan

October 2010

JEL classification: C22, C32, E21

www.rbnz.govt.nz/research/discusspapers/

Discussion Paper Series

ISSN 1177-7567

DP2010/09

Debt Dynamics and Excess Sensitivity of Consumption to Transitory Wealth Changes*

Emmanuel De Veirman and Ashley Dunstan[†]

Abstract

We analyse the consumption-wealth relationship using a framework that accounts for transitory variation in wealth, and in a setting where transitory variation in household net worth is not dominated by boom and bust cycles in stock markets. We find that transitory variation in consumption depends positively on recent transitory changes in wealth. In addition, we find that gross asset wealth and household debt are positively related. Both findings constitute departures from standard lifecycle/ permanent income hypothesis theory with complete financial markets, but can be explained by the introduction of liquidity constraints.

* The Reserve Bank of New Zealand's discussion paper series is externally refereed. The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Reserve Bank of New Zealand. We thank Phil Briggs, Andrew Coleman, Felix Delbrueck, Karen Dynan, Arthur Grimes, Oleg Korenok, Andreas Lehnert, Thomas Lubik, Pierre-Daniel Sarte, Jirka Slacalek, as well as seminar participants at the Federal Reserve Bank of Richmond, the Banque de France, the College of William and Mary, Virginia Commonwealth University, the University of Alabama, the University of Amsterdam, the University of Leuven, and the Reserve Bank of New Zealand for comments and discussion.

[†] Emmanuel De Veirman (corresponding author), Reserve Bank of New Zealand, 2 The Terrace, Wellington 6140, New Zealand. Email address: emmanuel.deveirman@rbnz.govt.nz. Ashley Dunstan, Reserve Bank of New Zealand. Email address: ashley.dunstan@rbnz.govt.nz.

1 Introduction

Standard permanent income hypothesis/lifecycle theory implies that consumption does not respond to transitory fluctuations in household net worth. However, the relation between consumption and transitory wealth variation is not well documented. To see why, note that the conventional approach to estimating the consumption-wealth relation only measures the response of consumption to permanent wealth changes, and abstracts from transitory wealth changes altogether.

In this paper, we investigate the effect of transitory wealth changes by using an empirical framework that accounts for transitory variation in wealth and income. This framework was first applied to the consumption-wealth relation by Lettau and Ludvigson (2004). However, Lettau and Ludvigson use a measure of US household net worth that is dominated by transitory fluctuations in stock market wealth. We will argue that for this reason, their study does not yield conclusions on the relation between consumption and transitory variation in overall household wealth.

In this paper, we use data for New Zealand, where stock cycles account for a small share of wealth variation. In such a setting, we are able to measure the response of consumption to transitory variation in wealth. Our findings suggest that transitory changes in wealth predict transitory changes in consumption, a finding which is at odds with standard permanent income hypothesis/lifecycle theory.

The implication that permanent income consumers hold spending stable in the face of transitory wealth shocks requires the assumption of frictionless financial markets that allow households to freely borrow more in response to a transitory wealth decrease, and to temporarily accumulate more savings after a transitory wealth increase. This suggests that the permanent income hypothesis implies a negative relationship between transitory changes in the value of gross asset holdings and changes in household debt. This implication is rarely examined. In fact, the macroeconomic literature on consumption wealth effects typically focuses on net worth without accounting separately for assets and liabilities. In this paper, we detect a positive relationship between gross asset wealth and household debt, both in terms of transitory deviations from trend as well as in the long run. Our results suggest that prolonged increases in asset prices tend to imply

prolonged periods of increased borrowing, a dynamic which is missing in studies that focus only on household net worth. This finding is intuitive from a practitioner's point of view, but it is at odds with standard lifecycle theory.

As will become clear, both our finding that consumption responds to transitory wealth changes and that debt relates positively to changes in gross wealth can be explained by relaxing the traditional assumption of frictionless financial markets, and by instead assuming that a non-trivial fraction of households is subject to a binding liquidity constraint. An alternative explanation is to relax the conventional assumption of perfect foresight/certainty equivalence on two counts, i.e. by jointly assuming that households save to insure against income uncertainty, and that they cannot perfectly distinguish between permanent and transitory wealth shocks.

The remainder of the introduction covers two points. First, we compare our approach to existing empirical literature on the consumption wealth effect. Second, we detail our findings and their implications for lifecycle theory.

On the first point, note that the conventional approach to estimating consumption wealth effects, as reviewed in Davis and Palumbo (2001), implicitly assumes that all wealth variation is permanent. Typically, this approach involves estimating a long-run relationship between consumption and wealth, and assuming that any transitory deviation from the long-run relationship can only be corrected for by changes in consumption. Because this approach abstracts from transitory wealth variation, it does not allow for investigating whether transitory wealth changes affect consumption.

Lettau and Ludvigson (2004) depart from the conventional literature by explicitly accounting for transitory wealth variation associated with departures from the long-run relationship. Their findings suggest that most variation in net worth is transitory, while there appears to be little transitory variation in consumption. From this disconnect between consumption and wealth in terms of the fraction of transitory variation, they conclude that a typical change in wealth is unrelated with consumer spending.

In the present paper, we argue that Lettau and Ludvigson's finding is contingent on the fact that transitory variation in US household net worth is dominated by boom and bust cycles in stock markets. In particular, we

argue that their findings speak for the relation between consumption and cyclical changes in stock prices, but do not yield conclusions about how consumption relates to transitory variation in any other type of wealth. To see this, note that Lettau and Ludvigson's finding that there is little transitory variation in consumption follows from their finding that consumption barely reacts to deviations from the long-run equilibrium relation, as captured by a cointegration residual. Lettau and Ludvigson interpret their cointegration residual as being primarily driven by a sequence of bull and bear markets in stock prices. Correspondingly, changes in housing wealth and the value of interest-bearing assets appear to have a negligible role in explaining the extent of deviation from long-run equilibrium. Therefore, the finding that consumption does not significantly respond to deviations from long-run equilibrium does not yield a conclusion about whether consumption reacts to transitory variation in housing wealth and the value of liquid financial assets, asset classes which together comprise two thirds of US household assets.

In this paper, we examine whether there are situations in which consumption does respond to transitory variation in the value of some asset classes. To do so, we use data for New Zealand, where stock cycles account for a small fraction of transitory variation in net worth. One reason for this is that the share of housing wealth in total assets is more than twice as large in New Zealand as it is in the United States. Another reason is that the stock market is small in New Zealand, in the sense that its ratio of total stock market capitalization to GDP is less than one third of the ratio in other Anglo-Saxon economies including the United States.¹

In this setting, we find that there is substantial transitory variation in consumption as well as in wealth. In addition, there turns out to be a systematic relation between transitory variation in wealth and transitory variation in consumption. In particular, consumption responds to mean-reverting changes in wealth with a lag. We also find that there is substantial permanent variation in wealth and consumption, and confirm the conventional finding that permanent changes in consumption relate positively to permanent changes in wealth.

Therefore, our evidence suggests that wealth changes do have a measurable

¹De Veirman and Dunstan's (2008) findings suggest that transitory variation in New Zealand gross wealth reflects mean-reverting fluctuations in housing wealth, while most variation in New Zealand financial wealth is permanent.

effect on consumption. As for the theoretical foundation of this relation, our finding that consumption responds to temporary swings in wealth is at odds with the conventional interpretation of the lifecycle/permanent income hypothesis which implies that households hold consumption stable in the face of predictable transitory fluctuations in wealth. This finding is related to existing findings of ‘excess sensitivity’ of consumption to current income,² but is different in the sense that we explicitly characterize the transitory component of wealth and document its relation to consumer spending.

Our finding that household debt responds with a positive sign to permanent as well as to transitory changes in gross asset wealth constitutes a second departure from the standard model. Since we document two departures from standard lifecycle theory in the same paper, we can be quite precise in terms of providing an explanation. As we discuss in our interpretation section, the hypothesis that a non-trivial fraction of households is subject to a binding liquidity constraint on borrowing can, in and of itself, explain both departures from traditional theory. Alternatively, we can explain our findings by combining the hypothesis that households have a precautionary saving motive induced by income uncertainty with the hypothesis that they cannot distinguish between permanent and transitory wealth shocks.

The remainder of our paper is structured as follows. Section 2 discusses data choices and characterizes the data. Section 3 explains the theoretical framework. Section 4 presents our empirical specification and results. We interpret the results in Section 5. Section 6 concludes.

2 Data

In this section, we first discuss our choices regarding data definitions. We then discuss the data, and document that stock market cycles do not dominate wealth variation in New Zealand. We refer to the Appendix

²Flavin (1981) documents that consumption responds more to current income than warranted by the information content of current income for permanent income. Campbell and Mankiw (1989) provide evidence in favor of a model where a fraction of households consume their current income rather than their permanent income.

for detailed information on data sources and the methodology used for constructing the variables.

In this paper, we use aggregate New Zealand data on real per-capita total household consumption, household wealth, and after-tax labor income. All data are for the period 1990Q1-2009Q1.

As we will discuss in Section 3, a household's consumption choices are in principle influenced by its total expected lifetime resources. Lifetime resources encompass current asset net worth and human wealth.

In this paper, we split asset net worth into two separate variables: household debt and the gross value of current asset holdings. Our measure of gross asset wealth combines the aggregate value of the housing stock with the total value of financial assets.

As for human wealth, current income constitutes its observable part. Our income measure is after-tax labor income, which differs from disposable income in that it excludes property income. Labor income is the standard income measure in studies on the consumption wealth effect.³ The reasoning behind using labor income is that property income is already accounted for by changes in the stock of wealth.

We use a measure of household consumption that includes spending on durable goods as well as on non-durable goods and services. This is consistent with the fact that our asset wealth measure excludes the stock of durable goods. It means that we treat the purchase of durable goods as consumption rather than as wealth accumulation.

As emphasized by Rudd and Whelan (2006) in a critique of Lettau and Ludvigson (2004), for estimating a long-run relation between consumption and wealth it is crucial to deflate all variables by the same price index. We use the Consumer Price Index (CPI) for that purpose.

We are now ready to discuss the patterns in the data. Figure 1 plots consumption and after-tax labor income. The recession of the early 1990s reflects a disinflation associated with the first years of inflation targeting by the Reserve Bank of New Zealand. In the mid- and late 1990s, consumption and income grew at similar rates. Consumption grew faster than labor income in the 2000s.

³See for instance Lettau and Ludvigson (2004) and Rudd and Whelan (2006).

Figure 2 shows household net worth and its components. The figure suggests that patterns in total net worth largely reflect changes in the value of the housing stock. This is consistent with the large share of housing in total assets. On average over our sample period, housing accounts for 68% of total household assets, which is more than double the share in the United States (31% over our sample period⁴). The share of housing wealth in New Zealand household assets increases from 65% in 1990 to 74% in 2009, an increase which largely reflects a housing boom between 2001 and 2007. The fact that housing wealth exerts such a large influence on developments in New Zealand household net worth implies that stock market cycles have only a minor influence on net worth.

A final point to note from Figure 2 is that the balance of outstanding household liabilities has tended to increase at a faster pace than asset worth has.

To better understand the role of stock wealth, we now turn to Figure 3. That figure decomposes household financial assets into stock wealth (labeled ‘equity’) and other financial assets. We define stock wealth as the sum of households’ direct holdings of stocks and indirect stock holdings through financial intermediaries such as mutual funds and pension schemes.⁵ The striking feature of Figure 3 is the comparatively small share of stock wealth in financial wealth. On average over our sample period, direct and indirect equity combined account for less than one third of financial wealth, and 9 percent of total household assets. This corresponds to the fact that the New Zealand stock market is small compared to those in other OECD economies, in the sense that a comparatively small fraction of New Zealand companies are listed on the stock market. Relative to GDP, New Zealand stock market capitalization is only about one-third to one-fourth as large as in other Anglo-Saxon economies.⁶

⁴Both the New Zealand and the US ratios are as a share of household assets excluding durable goods.

⁵The stock wealth series accounts for direct and indirect holdings of domestic as well as foreign stocks. It excludes assets held by financial intermediaries other than stocks.

⁶This comparison holds between New Zealand on the one hand, and the United States, the United Kingdom, and Australia on the other hand.

3 Analytical Framework

In this section, we discuss the theoretical underpinnings of our framework for estimating the long-run relationships between consumption, gross wealth, debt and income.

There are two major ways to motivate a long-run consumption relation. The first approach implies deriving a log-linear consumption equation from the traditional consumption function according to which aggregate consumption C_t in any year t is a constant multiple m of total lifetime wealth, consisting of asset wealth A_t and human wealth H_t :

$$C_t = m (A_t + H_t) \tag{1}$$

This relationship reflects the conventional interpretation of the ‘permanent income hypothesis’. According to equation (1), households consume their permanent income, which is the annuity value of their lifetime resources. This equation follows from a lifecycle/permanent income hypothesis model under particular assumptions. For instance, it requires that consumer choices are not influenced by uncertainty, either because households have perfect foresight about the future or because household preferences take a particular form (quadratic utility) that implies certainty equivalence. It also requires that households can freely lend and borrow in frictionless financial markets in order to ensure that consumption stays stable in the face of transitory wealth shocks. In this paper, we refer to the perfect foresight/certainty equivalence model with frictionless financial markets as the ‘standard’ or ‘conventional’ permanent income hypothesis/lifecycle model.⁷

A second way to motivate a long-run consumption relation is to derive it from a single equation: the intertemporal budget constraint. This is the approach we adopt in this paper. Unlike the first approach, this second

⁷The permanent income hypothesis is due to Friedman (1957). Modigliani and Brumberg (1954) developed the original lifecycle model, with subsequent empirical evidence by Ando and Modigliani (1963). Quadratic preferences and certainty equivalence have been a common assumption since Hall (1978). Attanasio (1999) provides a clear overview of the various incarnations of the lifecycle model. In this regard, see also Deaton (1992), Muellbauer (1994), and Davis and Palumbo (2001).

approach does not impose that all households are permanent income consumers, nor does it impose the above-mentioned assumptions required to derive equation (1). Therefore, it allows for the possibility that households take uncertainty into account, for instance by saving out of precaution against adverse shocks. Similarly, the existence of an intertemporal solvency constraint neither requires nor rules out the possibility that households are liquidity constrained in terms of the amount they can borrow in any given period.

Given this level of generality, we use the log-linearized budget constraint as part of our framework for testing two implications of the conventional permanent income hypothesis. First, equation (1) implies that consumption only responds to wealth changes that permanently alter the level of resources available over the lifetime. In this paper, we examine whether consumption also responds to transitory variation in wealth. Second, the desire by households to smooth consumption in the face of mean-reverting wealth shocks implies that they temporarily accumulate savings after a transitory wealth increase, and temporarily borrow more in response to a transitory wealth decrease. This suggests a negative relationship between transitory wealth changes and household debt. To test this second implication, we extend the existing budget constraint framework in such a way that debt enters as a variable separate from gross asset wealth, and investigate the sign of the relationship between changes in assets and liabilities.

We now explain how to derive a long-run equation involving consumption, gross wealth, debt and income from the intertemporal budget constraint. In sequence, we discuss the steps carried out by Campbell and Mankiw (1989), Lettau and Ludvigson (2001), and our extension to allow for explicit debt.

Campbell and Mankiw (1989) start from an aggregate intertemporal budget constraint of the following form:

$$W_{t+1} = (1 + R_{w,t+1}) (W_t - C_t) \tag{2}$$

where W_{t+1} represents total lifetime resources in period $t+1$, C_t represents consumption, and $R_{w,t+1}$ is the net return on wealth. In this specification, W_{t+1} includes asset wealth and human wealth. This reflects the assump-

tion that all wealth, including expected future labor income, in principle has a market value which can be spent on consumption. That is to say, equation (2) by itself allows for the possibility that households borrow against future labor income and spend the borrowed funds today. This is because the budget constraint is purely an intertemporal solvency condition that prevents households from borrowing more than they will ever be able to repay. However, the budget constraint can in principle co-exist with additional, more stringent constraints on household borrowing, for instance a liquidity constraint that restricts borrowing to a fraction of current assets.

Dividing equation (2) by W_t , taking logarithms, and approximating the term $\log(1 - \exp(c_t - w_t))$ around the steady-state log consumption-wealth ratio $c - w$, Campbell and Mankiw (1989) derive an accumulation equation for wealth, which decomposes wealth growth into a contribution of wealth returns and a contribution due to the fraction of wealth that is consumed in any given period:

$$\Delta w_{t+1} = r_{w,t+1} + \left(1 - \frac{1}{\rho}\right)(c_t - w_t) \quad (3)$$

Lower caps denote variables in logarithms. We omit constants throughout the theoretical section. Wealth growth Δw_{t+1} depends on the logarithm of the rate of return $r_{w,t+1} = \log(R_{w,t+1}) \approx \log(1 + R_{w,t+1})$, and it depends negatively on the log ratio of current consumption to wealth $c_t - w_t$. The parameter ρ equals the steady-state ratio of invested wealth to total wealth. In particular, $\rho = 1 - \exp(c - w) = [W - C]/W$. Since the derivation assumes that there are steady-state wealth shares, it requires that wealth shares be stationary.

Re-arranging the wealth accumulation equation, solving forward, and imposing the transversality condition $\lim_{i \rightarrow \infty} \rho^i (c_{t+i} - w_{t+i}) = 0$, Campbell and Mankiw (1989) arrive at the equation:

$$c_t - w_t = \sum_{i=1}^{\infty} \rho^i (r_{w,t+i} - \Delta c_{t+i}) \quad (4)$$

This equation implies that a high current ratio of consumption to wealth

must be followed by high future rates of return on wealth or by low future consumption growth, else it would not be sustainable in an intertemporal sense.

Lettau and Ludvigson (2001) build on this framework. They develop equation (4) further using three more equations, which we discuss one by one.

First, they log-linearize the wealth identity $W_t = A_t + H_t$ to arrive at $w_t \approx \omega a_t + (1 - \omega) h_t$. In this equation, a_t and h_t are log asset wealth and log human wealth, respectively, and $\omega = A/W$ is the steady-state share of asset wealth in total lifetime resources.

Second, they assume that the nonstationary component of log human wealth h_t is fully captured by log labor income y_t , which implies the assumption that discounted future labor income can be captured by a stationary term z_t . In particular, they assume that $h_t \approx y_t + z_t$, again omitting a constant.

Third, they assume that the log rate of return on total wealth depends linearly on the log returns on asset wealth $r_{a,t}$ and on human wealth $r_{h,t}$, with weights depending on the steady-state share ω of asset wealth in total resources. In particular, they assume that $r_{w,t} \approx \omega r_{a,t} + (1 - \omega) r_{h,t}$. This equation is a linear approximation of the logarithm of an accounting identity for the gross rate of return on total wealth $1 + R_{w,t+1}$ that was first performed by Campbell (1996).

Substituting these three equations into equation (4), and taking expectations conditional on information available at time t , Lettau and Ludvigson (2001) arrive at the following equation:

$$\begin{aligned} & c_t - \omega a_t - (1 - \omega) y_t \\ = & (1 - \omega) z_t + E_t \sum_{i=1}^{\infty} \rho^i (\omega r_{a,t+i} + (1 - \omega) r_{h,t+i} - \Delta c_{t+i}) \quad (5) \end{aligned}$$

The left-hand side of equation (5) is a log-linear combination of consumption, asset net worth, and current labor income. The right-hand side in

theory involves only stationary terms: the stationary component of human wealth z_t and the expected value of a term involving future returns on asset wealth $r_{a,t+i}$, future returns on human wealth $r_{h,t+i}$, and expected future consumption growth Δc_{t+i} .

In this paper, we extend equation (5) so as to explicitly account for the difference between debt and gross asset wealth. To do so, we log-linearize the identity $A_t = GA_t - D_t$, which states that asset net worth A_t equals gross asset wealth GA_t minus the stock of outstanding debt D_t . The log-linear approximation reads:

$$a_t \approx \psi ga_t - (\psi - 1) d_t \quad (6)$$

where ga_t is log gross asset wealth and d_t is log debt. ψ indicates the steady-state ratio GA/A of gross asset wealth to asset net worth. ψ is larger than 1 as long as steady-state debt D is positive. This implies that the coefficient $-(\psi - 1)$ assigned to debt is negative.

Substituting equation (6) into Lettau and Ludvigson's (2001) log-linearized wealth identity $w_t \approx \omega a_t + (1 - \omega) h_t$, we obtain:

$$w_t = \omega \psi ga_t - \omega (\psi - 1) d_t + (1 - \omega) h_t \quad (7)$$

The compound coefficients can be interpreted as follows. $\omega \psi$ is the steady-state share GA/W of gross asset wealth into total resources, while $-\omega (\psi - 1)$ is the steady-state contribution D/W of debt to total lifetime resources. This contribution is negative.

Substituting equation (7) into equation (4) along with the previously discussed assumption that current income captures the nonstationary component of human wealth, and the log-linearized equation for wealth returns, we obtain, after taking expectations:

$$\begin{aligned} & c_t - (1 - \omega) y_t - \omega \psi ga_t + \omega (\psi - 1) d_t \\ = & (1 - \omega) z_t + E_t \sum_{i=1}^{\infty} \rho^i (\omega r_{a,t+i} + (1 - \omega) r_{h,t+i} - \Delta c_{t+i}) \quad (8) \end{aligned}$$

On the left-hand side, consumption and debt enter with a positive sign, while income and gross asset wealth enter with a negative sign. Therefore, high rates of debt affect intertemporal sustainability beyond the effect of high levels of consumption. This is because all other things equal, high debt means low net worth. As will become clear in Section 4, the difference of our empirical framework from earlier work is that it allows for the possibility that high debt levels have different implications than low levels of gross asset wealth, even when they translate into the same net worth position.

In theory, the right-hand side of equation (8) is stationary. In the next section, we will test for cointegration to examine whether there indeed exist stationary log-linear combinations of consumption, debt, gross wealth and income. If the right-hand side indeed proves to be stationary empirically, then this is in line with the assumptions that current income fully captures the nonstationary component of human wealth such that z_t is stationary, and that wealth returns are stationary. Most importantly, empirical stationarity would be in line with the hypothesis that, according to rational expectations, households believe to live in accordance with the intertemporal budget constraint. If households were to violate the intertemporal budget constraint in expectation, then the transversality condition would not hold, such that a term $E_t \lim_{i \rightarrow \infty} \rho^i (c_{t+i} - w_{t+i}) \neq 0$ would enter the right-hand side of equation (8). If this term was non-zero, that would imply that households could only expect to be intertemporally solvent if they expect wealth (or consumption) to be infinite in the distant future.

4 Findings

In this section, we first specify and estimate the empirical long-run relations between consumption, gross wealth, debt, and income corresponding to equation (8). Next, we specify and estimate our empirical model for short-run transitions towards the equilibrium relationships. Finally, we decompose variation into transitory and permanent components, and document that transitory changes in wealth predict transitory variation in consumption and borrowing. We will interpret our findings in Section 5.

4.1 Long-Run Specification and Results

Table 1 presents the results from Augmented Dickey-Fuller (ADF) tests for stationarity of consumption, gross wealth, debt, and income. We enter a deterministic trend in every ADF test regression, implying that we test for the existence of a stochastic trend. For the orders of augmentation chosen by the Akaike and Schwarz information criteria, we cannot reject the null hypothesis of a unit root in either of these four variables.

Since the variables are non-stationary, we perform Johansen Trace and L-max tests for the existence of cointegration between those four variables. Table 2 documents the cointegration test results. For the purpose of the cointegration tests, we control for deterministic trends in the long-run equations. This implies that we test for the existence of common stochastic trends, but in principle allow for the variables to have different deterministic trends. The relevance of allowing for different deterministic trends will become clear later in this subsection.

The test results constitute strong evidence in favor of the existence of stochastic cointegration between consumption, gross wealth, debt and income. This suggests that households make their consumption decisions in such a way that, in expectation, they do not violate the intertemporal budget constraint.

In addition, we find that there exist two cointegrating relations in the four-variable system. This suggests that the right-hand side of (8) is best interpreted as reflecting two cointegrating residuals. Therefore, we capture the long-run links by a system of two empirical long-run equations. Identifying the two equations requires us to normalize every equation on two counts.

In the normalized system, the first equation is such that the consumption coefficient equals one, and the debt coefficient is zero. With these normalizations, we can interpret the first cointegrating relation as a consumption equation:

$$c_t - \kappa_1 - \beta_{y,1} y_t - \beta_{ga,1} ga_t = \eta_{1t} \quad (9)$$

In our notation, κ_1 is a constant and η_{1t} is the cointegrating residual to the

consumption equation. This equation differs from the usual specification in that the wealth measure ga_t is not net of liabilities. As we will see in a minute however, debt in principle relates to income and gross assets through a second cointegrating relation, and therefore plays an implicit role in the consumption equation. In the reported results, we do not enter a deterministic trend in equation (9), but the results are very similar when we do.

We now turn to the second long-run relation. In a three-variable system with consumption, net worth, and income, (unreported) L-max and Trace tests reveal the existence of a single cointegrating relationship. That is, we detect two cointegrating relations only when we decompose net worth into gross asset wealth and debt. This suggests that the second cointegrating equation captures a relationship between gross wealth and debt which is implicit in the system with net worth. Accordingly, the normalized second long-run relation has a coefficient of one on debt, and has a zero coefficient on consumption. We interpret the second cointegrating relation as a debt equation:

$$d_t - \kappa_2 - \beta_{y,2} y_t - \beta_{ga,2} ga_t - \gamma_{trend} t = \eta_{2t} \quad (10)$$

Where κ_2 is a constant, t is a deterministic trend, and η_{2t} is the cointegrating residual for the debt equation.

Based on a likelihood ratio test, we reject the null hypothesis of no deterministic trend in equation (10) at the 1 percent significance level. This implies that debt, assets and income differ in terms of their deterministic trends in a finite sample. As we will discuss shortly, this is in line with a process of financial liberalization over the course of our sample.⁸

We estimate equations (9) and (10) using Dynamic OLS.⁹

⁸In an open-economy model of the current account, Corsetti and Konstantinou (2009) analogously motivate including trends in an intertemporal current account equation by arguing that financial market liberalization may have induced trends in assets and liabilities, where the in-sample trend in assets may differ from the in-sample trend in liabilities.

⁹We entered three leads and lags of the differenced log variables in the DOLS regression. We adjust standard errors for serial correlation using the procedure explained in Hayashi (2000, p. 654 ff.)

Table 3 contains the estimated long-run coefficients. According to the consumption equation, a permanent one percent increase in gross asset wealth is associated with a permanent 0.08 percent quarterly increase in consumption. A one percent increase in income is associated with a long-run increase in consumption of 0.85 percent.¹⁰

From the debt equation, we see that permanent changes in income and gross asset wealth are positively and significantly associated with permanent changes in debt. We also find that the coefficient on the deterministic trend is positive.

As we will discuss in Section 5, the positive long-run relationship between debt and gross wealth points to the importance of financial market imperfections. Assuming that financial markets are incomplete, the presence of a positive trend in equation (10) is consistent with the interpretation that the effect of financial frictions has gradually been reduced by a trend increase in financial intermediaries' willingness to lend for any given level of current assets and income. In New Zealand, banks have indeed relaxed lending standards over the course of our sample, a phenomenon that we attribute to financial liberalization and increased international financial integration.

As we are about to see in equation (11), the two long-run relations form part of a Vector Error Correction Model (VECM) for short-run dynamics. We test for stability in the long-run component of this model by computing recursive test statistics for a structural break as described in Hansen and Johansen (1999). Figure 4 documents that the test statistics never exceed the 10 percent critical value, indicating stability of the two long-run equations as well as the VECM error-correction coefficients α that govern adjustment to deviations from the long-run relationships.

4.2 Short-Run Specification and Results

Equations (9) and (10) jointly tell us how the four variables in the system relate in steady state. In the short run, the variables will rarely

¹⁰These elasticities are within the range of studies for earlier countries. For cross-country studies in this respect, see Ludwig and Slok (2004), Case, Quigley, and Shiller (2005), and Slacalek (2009).

exactly satisfy the long-run relations, implying that most observations for the residuals η_{1t} and η_{2t} are non-zero. Cointegration implies that such deviations from steady state are temporary, in the sense that they tend to be corrected for by equilibrium-restoring changes in at least one of the variables in the system. In analogy with Lettau and Ludvigson (2004), we in principle allow any of the four variables of the system to correct for deviations from the long-run relations.

To do so, we estimate the following Vector Error Correction Model:

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

where X_t is a four-by-one vector containing data on consumption, debt, income, and gross asset wealth, A_0 is a vector of constants, and $A(L)$ is a matrix polynomial in the lag operator.

The five-by-one vector \mathcal{X}_{t-1} stacks the data vector with a deterministic trend:

$$\mathcal{X}_{t-1} = \begin{bmatrix} X_{t-1} \\ t \end{bmatrix} \quad (12)$$

The two-by-five matrix β' contains estimated coefficients for the two cointegrating relationships, such that:

$$\beta'\mathcal{X}_{t-1} = \begin{bmatrix} \eta_{1,t-1} \\ \eta_{2,t-1} \end{bmatrix} \quad (13)$$

The four-by-two matrix α captures every variable's response to each of the residuals $\eta_{1,t-1}$ and $\eta_{2,t-1}$. Based on the estimated adjustment parameters in α , we will be able to tell whether and how strongly any variable error-corrects. ε_t is a vector of short-run residuals.

To estimate this system, we first estimate the two long-run relationships $\beta'\mathcal{X}_{t-1}$, storing the two cointegrating residuals, and then enter $\beta'\mathcal{X}_{t-1}$ as a regressor as stated in equation (11). Both the Akaike Information Criterion and the Bayesian Information Criterion select a lag length of order

two in a levels VAR, corresponding to a lag length of order one in the VECM. Therefore, in our case $A(L)\Delta X_t = A_1 \Delta X_{t-1}$.

Table 4 provides results. The estimated adjustment parameters on the residual $\eta_{1,t-1}$ imply that income and gross asset wealth tend to restore their long-run relation with consumption at the 5% significance level or better, while consumption error-corrects at the 10% level. For instance, if consumption is high relative to current income and asset wealth, implying a positive value for $\eta_{1,t-1}$, that tends to predict rising income and asset wealth, and declining consumption. In Section 5, we will provide more economic intuition behind the present section's results.

Our results also show that deviations from equilibrium in the debt equation, as captured by $\eta_{2,t-1}$, tend to be corrected for by changes in debt and income, at the 5% level. Therefore, our model predicts that if debt is high relative to gross wealth and income, the household balance sheet position will tend to be restored through reduced borrowing and increases in income.

The coefficients on the lagged growth rates suggest that growth in asset wealth plays a key role in terms of the short-run dynamics. Growth in gross asset wealth affects all other variables significantly with a comparatively large coefficient. In addition, asset wealth growth is the only truly persistent variable in the system. Because of these two factors, the persistence in the system is primarily determined by persistence in asset wealth growth. In other words, the speed of error-correction in the system depends crucially on the persistence in the error-correction of asset wealth growth. The implications of this will become clear when we discuss impulse-responses to transitory shocks in the next subsection.

4.3 Permanent-Transitory Decomposition

Cointegration implies that any deviation from the two long-run relationships is temporary, but does not imply that all variation in the individual variables is transitory. In fact, every individual variable is nonstationary and therefore contains some permanent variation. In this subsection, we disentangle the fraction of the variability in the short-run shocks ε_t that implies permanent changes in the individual variables from the fraction

that implies transitory changes. We also document that transitory wealth variation predicts transitory variation in consumption and borrowing.

To decompose variation in the reduced-form VECM residuals ε_t from equation (11) into permanent and transitory components, we use a technique developed by Gonzalo and Ng (2001). Gonzalo and Ng (2001) show that, by multiplying the vector ε_t by a matrix that depends on the estimated short-run error-correction parameters α as well as on the estimated long-run coefficients β , one obtains a set of permanent and transitory shocks u_t . Explicitly writing ε_t and u_t in terms of their entries, this means in our case:

$$\begin{bmatrix} u_{P1,t} \\ u_{P2,t} \\ u_{T1,t} \\ u_{T2,t} \end{bmatrix} = \begin{bmatrix} \alpha'_{\perp} \\ \beta \end{bmatrix} \begin{bmatrix} \varepsilon_{c,t} \\ \varepsilon_{d,t} \\ \varepsilon_{y,t} \\ \varepsilon_{ga,t} \end{bmatrix} \quad (14)$$

Where $\varepsilon_{j,t}$ is the reduced-form residual to the VECM equation for the variable j , for $j = c, d, y, ga$. The matrix α_{\perp} is orthogonal to α such that $\alpha'_{\perp}\alpha = 0$. In our application, the permanent-transitory decomposition yields two permanent shocks $u_{P1,t}$ and $u_{P2,t}$ and two transitory shocks $u_{T1,t}$ and $u_{T2,t}$. The permanent shocks encompass any shocks that induce nonstationary variation in the variables, such as productivity shocks or permanent changes in consumer preferences. The transitory shocks explain mean-reverting fluctuations around the two long-run relationships.

The number of permanent and transitory shocks follows from the number of variables in our system as well as from the number of cointegrating relationships. We find evidence for two common stochastic trends in a four-variable system, which leaves two separate stochastic trends, each of which corresponds to one of the permanent shocks. Conversely, each of the two cointegrating relationships corresponds to one transitory shock.

The set of permanent and transitory shocks are in principle correlated. To identify the shocks, we orthogonalize the permanent and transitory shocks by applying a Cholesky decomposition to the vector u_t . Consistent with Lettau and Ludvigson (2004), we order the permanent shocks first, implying the assumption that transitory shocks do not affect permanent shocks contemporaneously.

Table 5 provides results from decomposing the variance of the forecast

errors associated with the growth rates of each of the four variables into variation induced by permanent and transitory shocks.¹¹ The upper part of the table shows the fraction of transitory variation, while the lower part displays the share of permanent variation. We find that transitory variation accounts for a substantial fraction of the variation in either variable, but that overall, most variation is permanent. A crucial difference with Lettau and Ludvigson (2004) is that in our setting, there is substantial transitory variation in consumption as well as in wealth, on the order of 20 to 30 percent of total variation. This absence of a disconnect between wealth and consumption in terms of the fraction of transitory variation in principle allows for the possibility that permanent as well as transitory changes in wealth relate to consumption.¹²

In Subsection 4.1, we found that permanent changes in wealth are positively related to permanent changes in consumption and debt. We now investigate whether transitory variation in wealth, borrowing and spending are related. For this purpose, we produce impulse-responses from applying one-standard deviation transitory shocks to the estimated VECM from equation (11). We first consider the effect of a transitory shock associated with the long-run consumption equation before turning to the debt equation shock.

Figure 5 graphs the impulse-responses for a time- t shock to the residual of the consumption equation. The bold line reflects a Cholesky decomposition which orders the consumption equation shock before the debt equation shock. The dotted line reverses the ordering of the two transitory shocks.¹³

In the case where we order the consumption equation shock before the debt equation shock, the shock to the consumption equation contemporaneously affects the debt residual. Therefore, in this case we effectively

¹¹In the reported results, we do not set insignificant adjustment parameters to zero for performing the permanent-transitory decomposition. Results are very similar when we do.

¹²As a robustness check, we also consider the case of unorthogonalized shocks. In this case, we derive correlated transitory and permanent shocks without imposing any Cholesky decomposition. Our (unreported) results suggest that our conclusions are robust to this alternative procedure. In particular, there continues to be substantial transitory as well as permanent variation in all four variables.

¹³In both cases, we continue to order the permanent shocks before the transitory shocks.

consider a shock that implies simultaneous deviations from both long-run relationships. This transitory shock is composed of time- t increases in asset wealth and income, and a time- t decrease in consumption. Because of this, the shock implies that consumption is below its long-run level as implied by current assets and income for reasons that are purely transitory.

After the shock, wealth continues to increase for two quarters, but then starts to error-correct. It does so gradually, and only reaches its pre-shock level about thirty quarters after the shock. Consumption increases in the first three quarters after the shock, but then starts to decline gradually. The dynamics of consumption track those of asset wealth with a one-quarter lag. In this respect, recall that our VECM estimates -in Table 4- imply that consumption growth strongly depends on lagged wealth growth. The impulse-responses suggest that this dependence on lagged wealth growth is a key driver of transitory consumption changes. Our finding that consumption temporarily increases in response to transitory wealth changes is at odds with standard lifecycle theory, an observation which we elaborate on in Section 5.

In Figure 5, debt also tracks the dynamics of asset wealth with a lag. This finding is consistent with the estimated strong dependence of debt growth on lagged wealth growth in our VECM. It suggests that debt temporarily increases in response to transitory increases in wealth, which is equally at odds with standard lifecycle theory.

The dotted line in Figure 5 reflects a Cholesky decomposition which orders the debt equation shock before the consumption equation shock. The results are similar to the case we just discussed.

Next, consider Figure 6, which graphs the impulse-responses for a one standard deviation shock to the residual of the debt equation. As before, we consider two slightly different scenarios, differing by the ordering of the transitory shocks. The interesting feature is that wealth dynamics are different across these two experiments, which allows us to further document that transitory wealth variation is an important determinant of transitory variation in consumption and debt.

When we order the consumption equation shock before the debt equation shock (the bold line), wealth does not jump on impact, but gradually declines with respect to its steady-state growth path. In that case, wealth

does not form part of the transitory shock. On the other hand, when we order the debt equation shock before the consumption shock (dotted line), wealth increases on impact.¹⁴

In the case where wealth forms part of the transitory shock, consumption and debt increase temporarily and in so doing mirror wealth dynamics with a lag, much like what we described above for Figure 5. When wealth does not jump up on impact, consumption declines more swiftly and debt rises more slowly. The reason appears to be that, with wealth declining relative to steady state, neither consumption nor debt are pushed upward by a transitory increase in wealth.

5 Interpretation

In this section, we first examine to which extent our evidence supports the conventional lifecycle model. Second, we evaluate candidate explanations for two empirical departures from standard lifecycle theory.

5.1 Comparing Evidence to Conventional Theory

We find that consumption relates to permanent as well as to transitory wealth changes. Therefore, our results suggest that a typical wealth change matters for consumption, as long as one considers a wealth measure not dominated by stock cycles. This greatly nuances Lettau and Ludvigson's (2004) findings.

However, the existence of an empirical relation between consumption and wealth does not necessarily mean that the relation is consistent with standard lifecycle theory. We now assess to which extent our results are in line with the perfect foresight/certainty equivalence lifecycle model with complete financial markets, as summarized in the consumption function (1) according to which households consume their permanent income.

¹⁴The difference reflects the fact that in the VECM results, gross wealth barely responds to the debt equation's error-correction term. When we shock the debt equation in isolation (as in the bold line in Figure 6), this implies that wealth is a negligible part of the transitory shock.

Our finding that there is a positive long-run relationship between wealth and consumption is in line with the permanent income hypothesis. This result is consistent with the interpretation that wealth changes cause consumption changes, even though we find that gross asset wealth and income error-correct. To see this, note that consumption can be high today relative to its long-run relation with income and wealth because households expect high future income or high future wealth returns. In this interpretation, our finding that income and wealth error-correct suggests that expectations about permanent changes in wealth or income tend to be confirmed by actual future wealth and income changes that empirically restore the long-run equilibrium relation.

However, our results are at odds with the conventional permanent income hypothesis on two counts. First, the permanent income hypothesis implies that consumption only responds to permanent wealth changes. This implication is at odds with our finding that transitory asset wealth fluctuations, which leave permanent income unaltered, predict transitory variation in consumption. This suggests that, while consumption plausibly depends on permanent income in the long run, short-run fluctuations in wealth cause substantial fluctuations in consumption around trend which cannot be explained by a conventional permanent income hypothesis/lifecycle model. This finding is akin to findings of excess sensitivity in consumption to current income by authors including Flavin (1981) and Campbell and Mankiw (1989).

Second, we find a positive relation between household debt and gross asset wealth, a finding that the standard permanent income hypothesis model with complete financial markets cannot explain. To see this, note that permanent income hypothesis households who operate in complete financial markets use their ability to freely save and borrow in order to keep consumption stable in the face of transitory shocks. To achieve this, such households temporarily increase savings in response to a transitory wealth gain, and borrow more in the event of a transitory wealth reduction. This suggests that the conventional permanent income hypothesis model implies a negative relationship between transitory wealth variation and household debt. However, this theoretical implication is precisely the opposite of our empirical finding that wealth increases are associated with increases in debt.

5.2 Reconciling Evidence with Theory

We now evaluate three candidate explanations for the two departures from standard lifecycle theory that we just documented.

A first candidate explanation lies in the possibility that households face a limit on current borrowing, implying that financial markets are not complete. In theories of consumer behavior with a credit constraint, the constraint is typically such that the value of outstanding debt cannot exceed a fraction of the current value of the housing stock.¹⁵ The intuition behind this is that housing is a common type of collateral against which lenders, uncertain about borrower revenue, are willing to supply funds.

Models with liquidity constraints are consistent with the transitory responses of consumption and debt to transitory wealth changes, as well as with the positive long-run relation between debt and wealth. This is the case as long as a non-trivial fraction of households is sufficiently impatient for the liquidity constraint to be binding. Since liquidity-constrained households are borrowing less than they would with complete financial markets, they choose to borrow more when a rise in asset values relaxes the constraint. If the wealth increase is permanent, all other things equal aggregate debt will increase permanently. If the wealth increase is transitory, then aggregate debt and consumption will temporarily increase as long as credit constraints are a function of current asset values.

The assumption that the ability to borrow depends on current asset values is not only consistent with the above-mentioned theories with household liquidity constraints, but it is also very much in line with reality. Empirically, banks specify maximum loan amounts as a fraction of the current value of collateral, without making lending contingent on the fraction of changes in asset values that is expected to be permanent. In sum, relaxing the assumption of frictionless financial markets by introducing liquid-

¹⁵Some of these models account for household decisions over the lifetime, while others are business cycle models that restrict heterogeneity to two types of households: impatient agents who will choose to borrow, and lenders. For lifecycle models with collateral constraints, see Japelli and Pagano (1994), Kiyotaki, Michaelides and Nikolov (2006), Ortalo-Magne and Rady (2006) and Coleman (2007). For general equilibrium models with household collateral constraints, see Aoki, Proudman and Vlieghe (2004), Iacoviello (2005), Davis and Heathcote (2005), Campbell and Hercowitz (2005), Calza, Monacelli and Stracca (2007), Monacelli (2009), and Iacoviello and Neri (2010).

ity constraints on borrowing is sufficient to explain both departures from standard theory.

Our second and third candidate explanation each involve relaxing the assumption of perfect foresight/certainty equivalence. The second candidate explanation consists of assuming that households cannot perfectly distinguish transitory from permanent shocks. The finding that consumption responds to transitory wealth fluctuations is in line with the hypothesis that households do not fully anticipate the wealth increase to be transitory, and increase consumption because in expectation, some of the wealth increase is permanent. However, this hypothesis in itself cannot explain the positive relation between debt and wealth. Therefore, allowing for misperceptions about transitory and permanent shocks is in itself not sufficient to reconcile evidence with theory.

A third possible explanation for our findings is that uncertainty about future income may inspire risk averse households to save out of precaution, thereby restricting themselves from borrowing or dissaving as much as they would in a certainty equivalence model. The precautionary saving motive is at the heart of buffer stock saving models.¹⁶ In these models, households make consumption decisions in line with their desire to reach a target ratio of current asset net worth to trend income at which consumer impatience exactly offsets the precautionary motive to save.

If households desire to achieve a particular net worth ratio, a permanent increase in the value of gross assets cannot leave a long-lasting effect on net worth. Therefore, buffer stock models imply that any increase in wealth must be offset by subsequent decreases in gross asset wealth or income, or by increases in debt. This is consistent with our empirical finding that debt cointegrates with gross wealth and income.

In isolation, the hypothesis that households are buffer stock savers appears to be an unconvincing explanation of the observed comovement between wealth, debt and consumption in terms of transitory variation. To see why, note that if households with a target net worth ratio anticipate an increase in gross wealth to be transitory, they may choose to alter consumption very little because they expect that a subsequent decline in gross wealth will

¹⁶For buffer stock saving models, see Zeldes (1989), Kimball (1990), Deaton (1991), Carroll (1997), Attanasio, Banks, Meghir and Weber (1999) and Gourinchas and Parker (2002). In a general equilibrium setting, see Krusell and Smith (1999).

restore their net worth ratio anyway.

However, we can explain both departures from standard theory by combining the second and third candidate explanations, meaning that households have a precautionary saving motive and cannot perfectly distinguish transitory from permanent shocks. Buffer stock savers would increase consumption in response to a transitory wealth increase which they expect to be permanent, because they believe that doing so is necessary to restore their target net worth ratio.

6 Conclusion

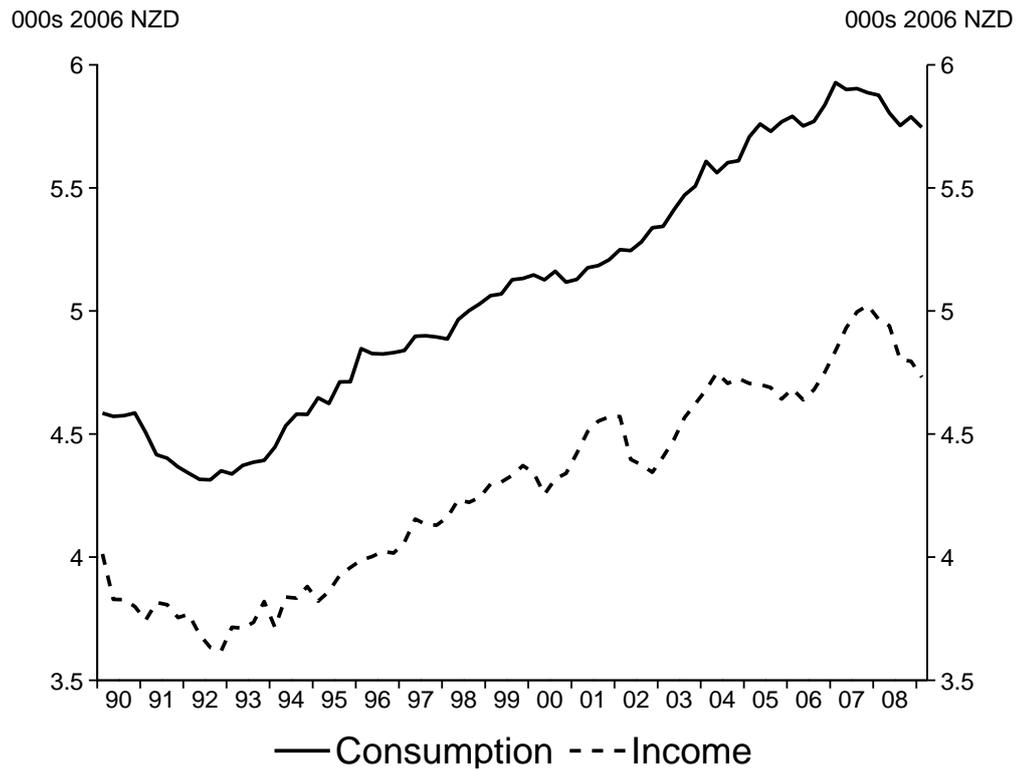
In this paper, we have documented two empirical grounds for rejecting the implication from perfect foresight/certainty equivalence lifecycle/permanent income hypothesis theory with complete financial markets that all households consume their permanent income. First, we find that consumption responds to transitory variation in wealth. Second, we find a positive relation between debt and wealth, both in the long run and in terms of transitory variation.

Our findings indicate that, to explain recent consumption and debt dynamics, it is important to relax the traditional assumption of frictionless financial markets. In particular, we find that the introduction of liquidity constraints by itself is sufficient to explain transitory as well as permanent comovement between wealth, consumption and debt. An alternative explanation is to relax the conventional assumption of perfect foresight/certainty equivalence on two counts, i.e. by jointly assuming that households save to insure against income uncertainty, and that they cannot perfectly distinguish between permanent and transitory wealth shocks.

In this paper, we are able to document the relation between transitory wealth changes and transitory consumption changes because we not only use an empirical procedure that accounts for transitory wealth variation, but also use a measure for household net worth that is not dominated by transitory variation in stock market wealth. A natural extension of our analysis to a broader set of economies would be to examine consumption-wealth comovement after filtering the cyclical component of variation in stock market wealth from net worth.

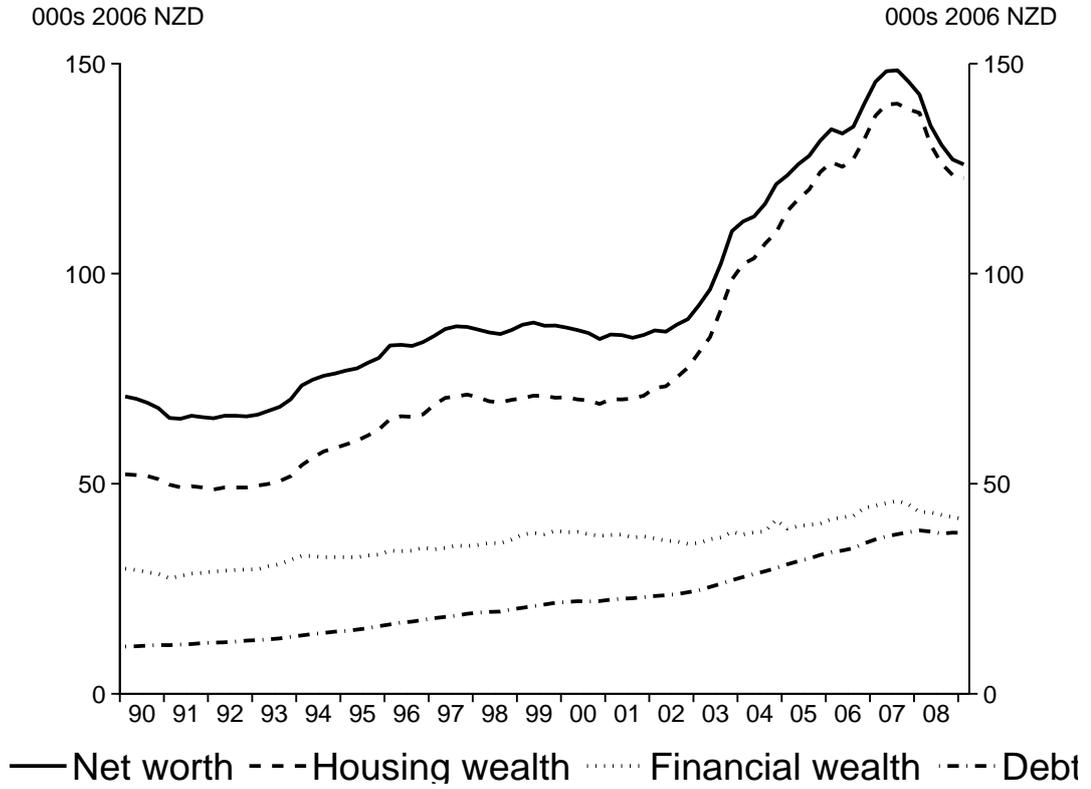
Tables and Figures

Figure 1: Quarterly Real Per Capita Consumption and After-Tax Labor Income



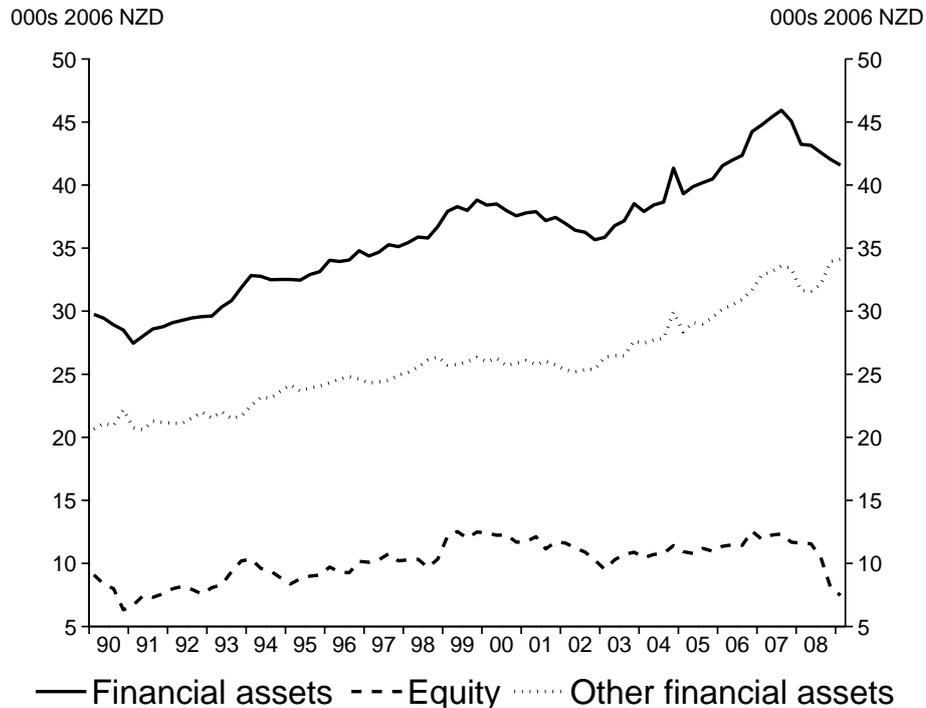
Note: This figure graphs quarterly consumption and after-tax labor income for New Zealand households, in real per capita terms. We obtained real data by deflating the nominal variables by the Consumer Price Index (CPI) with base year 2006. Further details on data construction are available in the Appendix.

Figure 2: Real Per Capita Household Net Worth and its Components



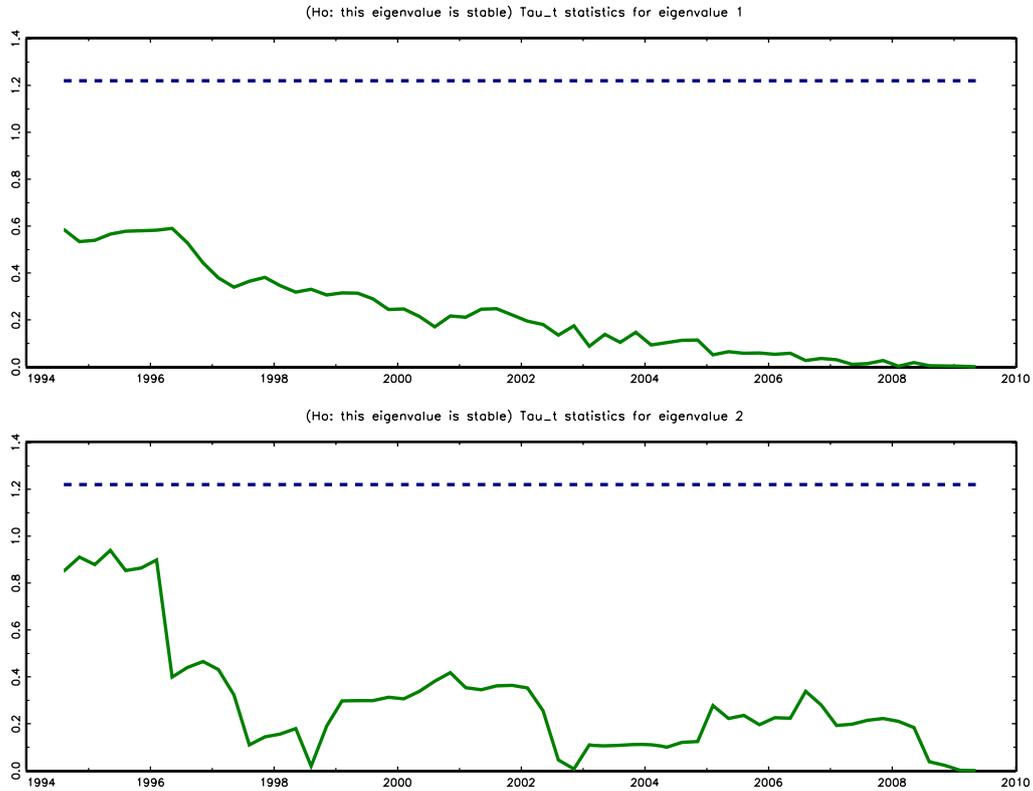
Note: This figure graphs quarterly series, in real per capita terms, for New Zealand household net worth and its subcomponents. The subcomponents are gross housing wealth, gross financial wealth, and household debt. This graph shows that fluctuations in financial wealth play a minor role in explaining fluctuations in net worth, as compared to the role played by housing wealth. We obtained real data by deflating the nominal variables by the Consumer Price Index (CPI) with base year 2006. Further details on data construction are available in the Appendix.

Figure 3: Real Per Capita Household Financial Wealth and its Components



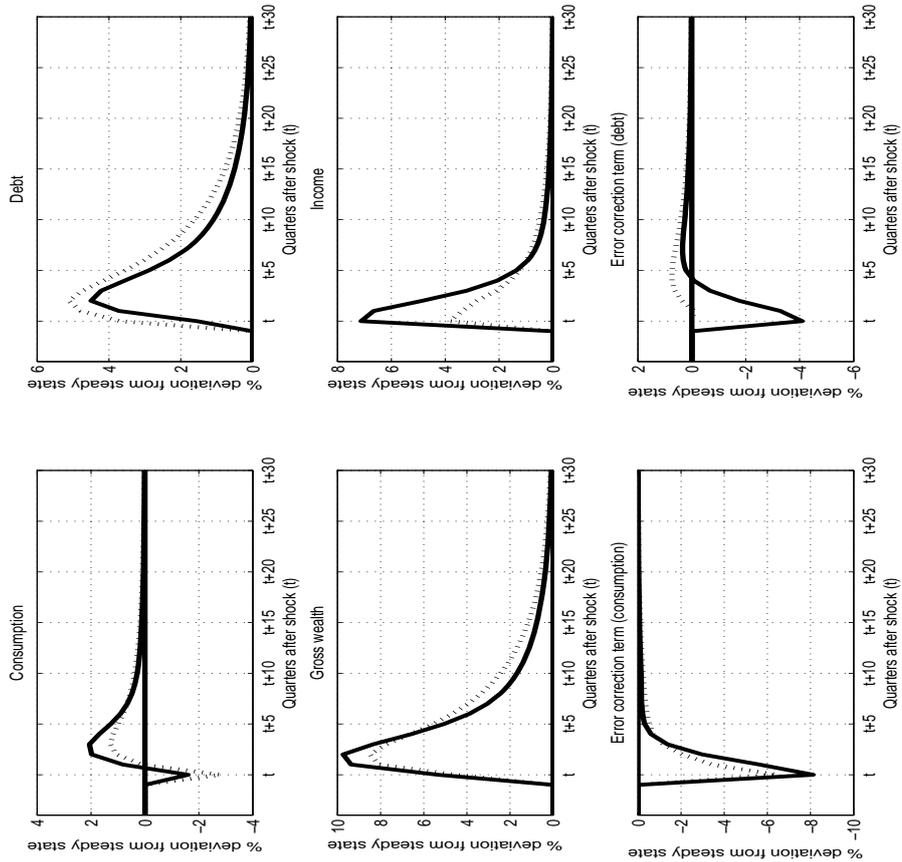
Note: This figure graphs quarterly series, in real per capita terms, for the financial wealth of New Zealand households and its components. Equity combines households' direct holdings of stocks with indirect investments in stocks through financial intermediaries including mutual funds and pension funds. On average over the sample, the sum of direct and indirect equity is only about one third of New Zealand financial assets. We obtained real data by deflating the nominal variables by the Consumer Price Index (CPI) with base year 2006. Further details on data construction are available in the Appendix.

Figure 4: Long-Run Stability Tests



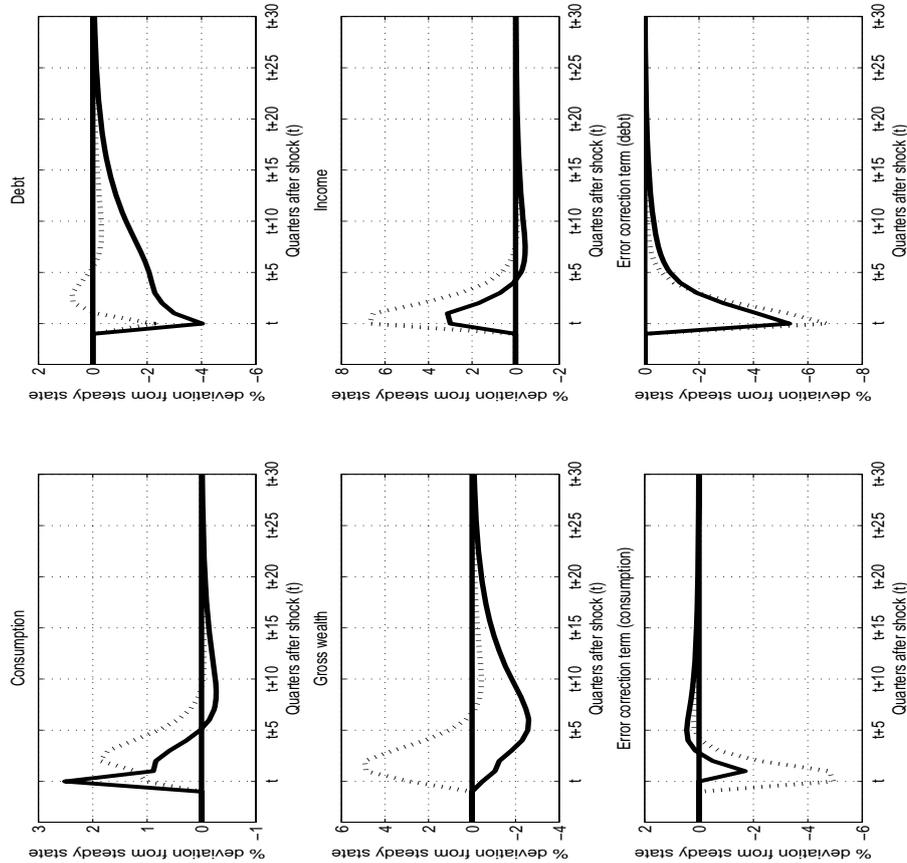
Note: This figure reports results from tests for a structural break in the long-run component of the VECM model (11) using the approach of Hansen and Johansen (1999). The top panel shows recursively computed test statistics for joint stability of the long-run consumption equation (9) and the VECM adjustment parameters to the consumption equation's error-correction term. The bottom panel shows recursive tests for the joint stability of the long-run debt equation (10) and the associated VECM adjustment parameters. Neither of the test statistics ever exceeds the 10 percent critical value. Therefore, we cannot reject the stability of the long-run component of the VECM at the 10 percent level.

Figure 5: Impulse-Responses: Shock to Consumption Equation



Note: This figure provides the VECM impulse-responses for a one-standard deviation shock to the error-correction term corresponding to the consumption equation. We graph the variables in levels, in terms of percentage deviations from their steady-state growth path. The bold line reflects a Cholesky decomposition which orders the consumption equation shock before the debt equation shock, while the dotted line reverses the ordering of the two transitory shocks. In both cases, we find that consumption and debt increase temporarily in response to a transitory increase in wealth.

Figure 6: Impulse-Responses: Shock to Debt Equation



Note: This figure provides the VECM impulse-responses for a one-standard deviation shock to the error-correction term corresponding to the debt equation. We graph the variables in levels, in terms of percentage deviations from their steady-state growth path. The bold line reflects a Cholesky decomposition which orders the consumption equation shock before the debt equation shock, while the dotted line reverses the ordering of the two transitory shocks. In the former scenario, gross wealth gradually decreases with respect to steady state, while in the latter case, the transitory shock implies an upward jump in gross wealth. Consumption decreases more swiftly, and debt decreases more slowly, in the case without a wealth increase. This suggests that one of the reasons for the hump-shaped response in consumption and debt in the other case is their responding to hump-shaped wealth dynamics.

Table 1: Augmented Dickey-Fuller Tests for Stationarity

	Lag length				
	0	1	2	3	4
Consumption	-2.31 ^S	-2.34	-2.28	-2.45 ^A	-2.62
Gross wealth	-1.64	-2.97 ^{A,S}	-2.77	-3.08	-3.11
Debt	-1.08	-1.97	-2.61 ^S	-2.98 ^A	-2.98
Income	-2.42	-2.49 ^{A,S}	-3.36*	-3.33*	-2.93
5% critical value	-3.46				

Note: Every variable is in real per capita terms. For every variable, we report the ADF test statistic for the null hypothesis that the series has a unit root. ‘Lag length’ refers to the order of augmentation of the testing regression. Superscripts A,S refer to the lag length selected by the Akaike and Schwarz information criteria, respectively. In every test regression, we include a constant as well as a deterministic trend. We cannot reject the null hypothesis of a unit root at the 5 percent level for any variable at any of the reported orders of augmentation. *** indicates significance at 1% level, ** at 5% level, and * at 10% level.

Table 2: Johansen L-Max and Trace Tests for Cointegration

Test				Lag length			
	H ₀	H _A	5% crit.	1	2	3	4
L-max test	r=0	r=1	32.11	44.15***	52.8***	51.4***	51.1***
	r=1	r=2	25.8	28.2**	36.8***	39.5***	23.2*
	r=2	r=3	19.4	9.5	10.5	13.6	18.7*
Trace test	r=0	r ≥ 1	63.7	101.1***	98.0***	107.6***	98.2***
	r=1	r ≥ 2	42.8	39.8*	53.9***	57.3***	47.2**
	r=2	r ≥ 3	25.7	16.6	18.4	15.7	20.0

Note: This table presents Johansen L-max and Trace test statistics for the number of cointegrating vectors in a four-variable system containing consumption, gross wealth, debt, and income. ‘5% crit.’ refers to the 5 percent critical value. ‘Lag length’ refers to the number of lagged differences included in the VECM for the purpose of conducting the L-max and Trace tests. The Akaike information criterion suggests to use three lags, while the Schwarz information criterion suggests one lag. All tests include trends in the long-run relationships. For most lag specifications, we can reject the null hypotheses of no cointegration and of a single cointegrating vector, but cannot reject the null hypothesis of two cointegrating vectors, at the 5 percent level. *** indicates significance at 1% level, ** at 5% level, and * at 10% level.

Table 3: Estimated Long-Run Relationships

	consumption equation	debt equation
constant		
income	0.85*** (9.94)	0.50*** (2.90)
gross asset wealth	0.08** (2.46)	0.40*** (8.19)
trend	— —	0.84*** (10.17)

Note: The middle column presents estimates of the long-run consumption equation (9), while the rightmost column contains the estimated coefficients of the long-run debt equation (10). We estimated both equations using Dynamic Ordinary Least Squares (DOLS) with three leads and lags of the differenced log variables. We adjusted t-statistics for serial correlation using the procedure in Hayashi (2000, p. 654 ff.). The results suggest that permanent changes in consumption and debt are positively associated with trend changes in wealth and income. *** indicates significance at 1% level, ** at 5% level, and * at 10% level.

Table 4: Estimated Vector Error Correction Model

	Δc_t	Δd_t	Δy_t	Δga_t
constant	0.03 (0.98)	-0.06* (-1.69)	0.25*** (4.27)	0.11* (1.98)
η_{1t-1}	-0.11* (-1.78)	0.13* (1.95)	0.29*** (2.71)	0.26** (2.48)
η_{2t-1}	0.11** (2.47)	-0.18*** (-3.77)	0.16** (2.10)	-0.004 (-0.060)
Δc_{t-1}	-0.24* (-1.97)	-0.14 (-1.10)	-0.17 (-0.83)	-0.44** (-2.19)
Δd_{t-1}	0.04 (0.37)	-0.10 (-0.89)	-0.32* (-1.72)	-0.07 (-0.37)
Δy_{t-1}	-0.04 (-0.56)	0.06 (0.82)	0.10 (0.89)	0.20* (1.74)
Δga_{t-1}	0.34*** (4.64)	0.37*** (4.78)	0.38*** (3.06)	0.77*** (6.35)
R^2	0.34	0.53	0.32	0.59
\overline{R}^2	0.28	0.49	0.26	0.55

Note: This table reports estimates for the VECM in equation (11). t-statistics are in parentheses. η_{1t-1} is the error-correction term for the long-run consumption equation, while η_{2t-1} is the error-correction term for the long-run debt equation. The results suggest that income and gross asset wealth error-correct to the consumption equation residual at the 5 percent level or better, and that debt and income error-correct to the debt equation residual. The results also imply that lagged growth in gross wealth significantly affects growth in all four variables in the system. *** indicates significance at 1% level, ** at 5% level, and * at 10% level.

Table 5: Permanent-Transitory Decomposition

	Δc	Δd	Δy	Δga
Transitory				
$h = 1$	35.90	42.30	49.07	29.94
$h = 4$	35.93	34.75	49.76	23.68
$h = 8$	35.94	32.98	49.69	23.48
$h = 20$	35.89	32.46	49.63	23.29
$h = \infty$	35.88	32.45	49.63	23.28
Permanent				
$h = 1$	64.10	57.70	50.93	70.06
$h = 4$	64.07	65.25	50.24	76.32
$h = 8$	64.06	67.02	50.31	76.52
$h = 20$	64.11	67.54	50.37	76.71
$h = \infty$	64.12	67.55	50.37	76.72

Note: For consumption, debt, income, and gross asset wealth growth, respectively, this table provides the fraction of the forecast error variance, in percent, explained by transitory and permanent shocks. The column labels are shorthand for the forecast error variance applicable to h -quarter ahead growth forecasts. For instance, Δc stands for $\Delta c_{t+h} - E_t \Delta c_{t+h}$. To compute transitory and permanent shocks, we first apply the Gonzalo-Ng (2001) decomposition as stated in equation (14) to the reduced-form VECM residuals from equation (11), without setting insignificant adjustment coefficients to zero. Next, we orthogonalize the permanent and transitory shocks by applying a Cholesky decomposition to the vector of transitory and permanent shocks, assuming that transitory shocks do not contemporaneously affect permanent shocks. Unlike Lettau and Ludvigson (2004), we find that both permanent and transitory shocks account for a non-trivial part of the variation in every of the four variables.

Appendix: Data Construction

This appendix discusses the data sources and the methodology used to derive the variables in our dataset.

Our series for household consumption is from the quarterly national accounts published by Statistics New Zealand. Household consumption excludes the consumption of private non-profit organisations. We include consumption on durable and non-durable goods and services.

We compute after-tax labor income as follows: pre-tax labor income minus tax payments plus transfer income. In its national accounts releases, Statistics New Zealand (SNZ) produces an annual series for labor income. However, there is no directly available measure for quarterly labor income in New Zealand. We construct quarterly labor income by multiplying average hourly earnings (including overtime payments) from the Quarterly Employment Survey (QES) by hours worked from the Household Labour Force Survey (HLFS). The latter survey includes hours worked by agricultural workers and workers that are otherwise self-employed, but the former study excludes earnings of workers in those sectors. Under the assumption that a typical farmer or entrepreneur earns the same hourly labor income as that prevailing on average in other sectors, our measure of labor income captures the compensation of employees as well as entrepreneurial income. We re-scaled our quarterly proxy for labor income to ensure that it is consistent with the annual labor income series from the SNZ national accounts.

We construct a series of quarterly tax payments using our measure of labor income as well as data on tax rates by income bracket. For the latter, we obtained an annual series on implied effective tax rates from the annual national accounts produced by SNZ, and interpolated it to obtain a quarterly series.

As the final variable necessary to construct after-tax labor income, we computed a quarterly measure of transfer receipts. For unemployment and pension benefits, the Ministry of Social Development provides benefit rates, from which we computed an estimate of the number of beneficiaries. For other transfer receipts, there are no available estimates on benefit rates. We account for these other payments by interpolating the corresponding annual data from Work and Income New Zealand. Since these

payments likely vary gradually over time (e.g. it includes the sickness and invalidity benefit), we do not believe to omit a significant amount of quarterly fluctuations by using an interpolated series.

We obtained the value of the housing stock from official Reserve Bank of New Zealand estimates. The primary input into these estimates is the Quotable Value New Zealand (QVNZ) house price index. QVNZ obtains capital values from local authorities which conduct periodic revaluations for the purpose of levying rates. QVNZ uses the ratio of actual sales prices to the capital values to produce a price index for each local authority area, which it then aggregates.

We use annual estimates of March quarter total residential dwelling values from Quotable Value New Zealand. We construct a quarterly series for housing values by using Statistics New Zealand data on dwelling consents (i.e. measuring the total number of dwellings) and the house price index.

We obtained quarterly data on financial assets after 1995 from Reserve Bank of New Zealand estimates. The annual version of this series is publicly available. The Reserve Bank does not compute a quarterly series of financial assets before 1995. We compute a quarterly proxy by interpolating available annual data. To construct a quarterly measure of directly held equities, we matched the growth rates each quarter with a weighted average of the New Zealand and Australian capital price indices, both denominated in New Zealand Dollars. The time-varying weight for the New Zealand capital price index is equal to the proportion of direct equity that is domestic, while the weight for the Australian index equals the proportion of direct equity that is held abroad. The other types of financial assets evolve gradually over time, such that we construct a quarterly measure for these assets before 1995 by interpolating their respective annual series with a cubic spline.

Our series for household debt reflects Reserve Bank of New Zealand estimates.

A caveat to our analysis is that in New Zealand, neither investments in farms and other unincorporated businesses nor investments in privately held corporations are accounted for in the available statistics on household net worth. We did not impute wealth held in unincorporated businesses and privately held firms.

References

- Ando, A and F Modigliani (1963), “The Life Cycle Hypothesis of Saving: Aggregate Implications and Tests”, *American Economic Review*, 53, 55-84.
- Aoki, K, J Proudman and G Vlieghe (2004), “House Prices, Consumption, and Monetary Policy: a Financial Accelerator Approach”, *Journal of Financial Intermediation*, 13, 414-435.
- Attanasio, O P (1999), “Consumption”, in *Handbook of Macroeconomics*, Volume 1B, eds J B Taylor and M Woodford, Elsevier.
- Attanasio, O P, J Banks, C Meghir and G Weber (1999), “Humps and Bumps in Lifetime Consumption”, *Journal of Business and Economic Statistics*, 17, 22-35.
- Calza, A, T Monacelli and L Stracca (2007), “Mortgage Markets, Collateral Constraints, and Monetary Policy: Do Institutional Factors Matter?”, CEPR Discussion Paper 6231.
- Campbell, J Y (1996), “Understanding Risk and Return”, *Journal of Political Economy*, 104, 298-345.
- Campbell, J Y and N G Mankiw (1989), “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence”, in *NBER Macroeconomics Annual 1989*, University of Chicago Press.
- Campbell, J R and Z Hercowitz (2005), “The Role of Collateralized Household Debt in Macroeconomic Stabilization”, NBER Working Paper W11330.
- Carroll, C D (1997), “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis”, *Quarterly Journal of Economics*, 112, 1-55.
- Case, K E, J M Quigley and R J Shiller (2005), “Comparing Wealth Effects: the Stock Market versus the Housing Market”, *B.E. Journal of Macroeconomics (Advances)*, 5, article 1.
- Coleman, A (2007), “Credit Constraints and Housing Markets in New Zealand”, Reserve Bank of New Zealand Discussion Paper 2007/11.

- Corsetti, G and P T Konstantinou (2009), “What Drives US Foreign Borrowing? Evidence on External Adjustment to Transitory and Permanent Shocks”, CEPR Discussion Paper 7134.
- Davis, M A and J Heathcote (2005), “Housing and the Business Cycle”, *International Economic Review*, 46, 751-784.
- Davis, M A and M 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.
- Deaton, A (1991), “Saving and Liquidity Constraints”, *Econometrica*, 59, 1221-1248.
- Deaton, A (1992), “*Understanding Consumption*”, Oxford University Press.
- De Veirman, E and A Dunstan (2008), “How Do Housing Wealth, Financial Wealth and Consumption Interact? Evidence from New Zealand”, Reserve Bank of New Zealand Discussion Paper 2008/05.
- Flavin, M A (1981), “The Adjustment of Consumption to Changing Expectations about Future Income”, *Journal of Political Economy*, 89, 974-1009.
- Friedman, M (1957), “*A Theory of the Consumption Function*”, Princeton University Press.
- 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.
- Gourinchas, P and J A Parker (2002), “Consumption over the Life Cycle”, *Econometrica*, 70, 47-89.
- Hall, R E (1978), “Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence”, *Journal of Political Economy*, 86, 971-987.
- Hansen, H and S Johansen (1999), “Some Tests for Parameter Constancy in Cointegrated VAR Models”, *Econometrics Journal*, 2, 306-333.
- Hayashi, F (2000), “*Econometrics*”, Princeton University Press.

- Iacoviello, M (2005), “House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle”, *American Economic Review*, 95, 739-764.
- Iacoviello, M and S Neri (2010), “Housing Market Spillovers: Evidence from an Estimated DSGE Model”, *American Economic Journal: Macroeconomics*, 2, 125-164.
- Japelli, T and M Pagano (1994), “Saving, Growth, and Liquidity Constraints”, *Quarterly Journal of Economics*, 109, 83-109.
- Kimball, M S (1990), “Precautionary Saving in the Small and in the Large”, *Econometrica*, 58, 53-73.
- Kiyotaki, N, A Michaelides and K Nikolov (2006), “From Shirtsleeves to Shirtsleeves in a Long Lifetime”, mimeo, Princeton University.
- Krusell, P and A A Smith, Jr. (1999), “On the Welfare Effects of Eliminating Business Cycles”, *Review of Economic Dynamics*, 2, 245-272.
- Lettau, M and S Ludvigson (2001), “Consumption, Aggregate Wealth, and Expected Stock Returns”, *Journal of Finance*, 56, 815–849.
- Lettau, M and S Ludvigson (2004), “Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption”, *American Economic Review*, 94, 276-299.
- Ludwig, A and T Slok (2004), “The Relationship between Stock Prices, House Prices and Consumption in OECD Countries”, *B.E. Journal of Macroeconomics (Topics)*, 4, Article 4.
- Modigliani, F and R Brumberg (1954), “Utility Analysis and the Consumption Function: an Interpretation of Cross-Section Data”, in *Post Keynesian Economics*, ed K K Kurihara, Rutgers University Press.
- Monacelli, T (2009), “New Keynesian Models, Durable Goods, and Collateral Constraints”, *Journal of Monetary Economics*, 56, 242-254.
- Muellbauer, J (1994), “The Assessment: Consumer Expenditure”, *Oxford Review of Economic Policy*, 10, 1-41.
- Ortalo-Magne, F and S Rady (2006), “Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints”, *Review of Economic Studies*, 73, 459-485.

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

Slacalek, J (2009), “What Drives Personal Consumption? The Role of Housing and Financial Wealth”, *B.E. Journal of Macroeconomics (Topics)*, 9, Article 37.

Zeldes, S P (1989), “Consumption and Liquidity Constraints: an Empirical Investigation”, *Journal of Political Economy*, 97, 305-346.