Essays in Economics and Econometrics

Richard Finlay

A thesis in fulfilment of the requirements for the degree of

Doctor of Philosophy

UNSW Business School
Faculty of Economics

September 2015

PLEASE TYPE

THE UNIVERSITY OF NEW SOUTH WALES
Thesis/Dissertation Sheet

Surname or Family name: Finlay

First name: Richard | **Other name/s: Andrew**

Abbreviation for degree as given in the University calendar: **PhD**

School: UNSW Business School | **Faculty: Economics**

Title: Essays in economics and econometrics

Abstract 350 words maximum: (PLEASE TYPE)

The body of this thesis consists of three published papers written while enrolled in a PhD; in addition, two published papers completed while enrolled in a Master of Economics are included in an appendix. Of the five papers, four are largely empirical in nature, with the subject matter reflecting my (and my employer the Reserve Bank of Australia's) interest in areas relevant to policy-making institutions in Australia. The findings of this body of work can be summarised as follows: housing wealth effects in Australia arise from an easing of collateral constrains and a common association between rising house prices and another factor such as rising income expectations, rather than through 'traditional wealth effects'; the rise in Australian household saving over the 2000s appears to have been driven by a reduction in permanent income expectations following the financial crisis and a desire to pay down debt and rebuild assets; negative credit-supply shocks explain one-third to one-half of the fall in credit growth seen over the financial crisis, and around one-sixth of the fall in GDP growth, and so played an important but not dominant role in economic outcomes over the period; and much of the quarter-to-quarter volatility in Australian GDP data appears due to measurement error. The fifth paper is very different — it is a theoretical econometrics paper with little obvious application to policy-making. The paper is nonetheless a contribution to the literature in that it makes available via construction a new class of highly flexible (in terms of both permissible marginal distribution and permissible correlation structure) non-Gaussian random field, for possible use in economic and/or financial modelling.

Originality Statement

'I hereby declare that this submission is my own work and to the best of my knowledge it contains no materials previously published or written by another person, or substantial proportions of material which have been accepted for the award of any other degree or diploma at UNSW or any other educational institution, except where due acknowledgement is made in the thesis. Any contribution made to the research by others, with whom I have worked at UNSW or elsewhere, is explicitly acknowledged in the thesis. I also declare that the intellectual content of this thesis is the product of my own work, except to the extent that assistance from others in the project's design and conception or in style, presentation and linguistic expression is acknowledged.'

Signed ......R. A. Finley..............

Date......10/8/2015..................

Authenticity Statement

'I certify that the Library deposit digital copy is a direct equivalent of the final officially approved version of my thesis. No emendation of content has occurred and if there are any minor variations in formatting, they are the result of the conversion to digital format.'

Signed ......R. A. Finlay......

Date......10/8/2015......

Copyright Statement

'I hereby grant the University of New South Wales or its agents the right to archive and to make available my thesis or dissertation in whole or part in the University libraries in all forms of media, now or here after known, subject to the provisions of the Copyright Act 1968. I retain all proprietary rights, such as patent rights. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertation.

I also authorise University Microfilms to use the 350 word abstract of my thesis in Dissertation Abstract International (this is applicable to doctoral theses only).

I have either used no substantial portions of copyright material in my thesis or I have obtained permission to use copyright material; where permission has not been granted I have applied/will apply for a partial restriction of the digital copy of my thesis or dissertation.'

Signed R. A. Finley

Date 10/8/2015

Publication Declaration Statement

Publication I: Windsor, C., Jaaskela, J. and Finlay, R. (2015). Housing wealth effects: evidence from an Australian panel. *Economica*, **82**, 552-577.

Declaration: I certify that this publication was a direct result of my research towards this PhD, and that reproduction in this thesis does not breach copyright regulations.

R. A. Finlay 10/8/15

Publication II: Finlay, R. and Price, F. (2015). Household saving in Australia. *The B.E. Journal of Macroeconomics*, **15**, 677-704.

Declaration: I certify that this publication was a direct result of my research towards this PhD, and that reproduction in this thesis does not breach copyright regulations.

R. A. Finlay 10/8/15

Publication III: Rees, D., Lancaster, D. and Finlay, R. (2015). A state-space approach to Australian GDP measurement. *Australia Economic Review*, **48**, 133-149.

Declaration: I certify that this publication was a direct result of my research towards this PhD, and that reproduction in this thesis does not breach copyright regulations.

R. A. Finlay 10/8/15

Publication A.I: Finlay, R. and Jaaskela, J. (2014). Credit supply shocks and the global financial crisis in three small open economies. *Journal of Macroeconomics*, **40**, 270-276.

Declaration: I certify that this publication was a direct result of my research towards this PhD, and that reproduction in this thesis does not breach copyright regulations.

R. A. Finlay 10/8/15

Publication A.II: Finlay, R. and Seneta, E. (2014). Random fields with Pólya correlation structure. *Journal of Applied Probability*, **51**, 1037-1050.

Declaration: I certify that this publication was a direct result of my research towards this PhD, and that reproduction in this thesis does not breach copyright regulations.

R. A. Finlay 10/8/15

Table of Contents — Page

Introduction

The body of this thesis consists of three published papers completed over the course of my study for a PhD at UNSW; in addition, two published papers completed while studying for a Master or Economics are included in the appendix. Of the five papers, four are largely empirical in nature, with topics reflecting my (and my employer the Reserve Bank of Australia's) interest in areas relevant to policy-making institutions in Australia. In particular, two concern household saving and consumption in Australia, one suggests a new way to estimate quarterly GDP growth that has some advantages over the existing simple average of expenditure, income and product estimates used as the headline GDP figure by the Australian Bureau of Statistics (ABS), and the last (completed while studying for a Master of Economics and included in the appendix) examines the role of credit supply shocks during the Global Financial Crisis (GFC). The fifth paper, with was completed while studying for a Master of Economics and is included in the appendix, is different — it is a theoretical econometrics paper whose aim is to make available a new and flexible class of non-Gaussian random fields to applied researches in economics and finance. Below I give more detail on each paper; as all are co-authored, I also clarify my contribution to the publications.

Housing wealth effects: evidence from an Australian panel (Economica).

This paper explores the positive relationship between house prices and household spending by following a panel of Australian households using the HILDA Survey. We examine three potential explanations for the observed wealth effects: (1) a traditional housing wealth effect, whereby those with 'excess' housing increase spending as house prices rise due to higher lifetime resources (this explanation suggests that older households, who are more likely to own excess housing, will have the largest wealth effects); (2) a collateral channel, whereby rising house prices increase the value of housing collateral and so loosen credit constraints (this explanation suggests that younger households, who are more likely to be credit constrained, will have the largest wealth effects); and (3) that house prices and spending are both influenced by a common third factor such as rising permanent income expectations (this explanation suggests that younger households, who face more years of life, will have the largest wealth effects, and that renters as well as owners will benefit).

We find no evidence for 'traditional housing wealth effects'. Rather, young homeowners exhibit the largest wealth effects. Young renters also exhibit a positive consumption response to house prices, although less so than young homeowners. This suggests that increasing house prices raise spending via easing credit constraints and a common association between house prices and a third factor.

We contribute to the literature in a number of ways. First, unlike most past studies we follow an actual panel of households rather than a 'pseudo-panel' based on birth-cohorts. This allows for more precise estimates, and allows us to control for unobservable, time-invariant differences between households. Another novel feature of our study is the introduction of time-fixed effects to control for common macroeconomic influences. Finally, the availability of panel data allows us to compare hypotheses (2) and (3) in a more satisfactory manner than previous papers relying on pseudo-panels. The distribution of housing wealth effects by homeownership status is potentially biased when using pseudo-panels — as acknowledged by Attanasio et al. (2009) among others — because the split between homeowners and renters is possibly not constant over time, and because it is difficult to control for selection into homeownership with these data. To assess the possibility of problems with pseudo-panels we perform our estimation as an actual panel and as a pseudo-panel. The results for the age distribution of housing wealth effects of homeowners from the pseudo-panel are qualitatively similar to those obtained from the equivalent actual panel, although the necessary use of aggregate house prices (rather than self-assessed house prices) in pseudo-panels appears to inflate estimated housing wealth effects. This suggests that pseudo-panels may provide a partial, although imperfect, substitute for actual panels when estimating the age distribution of housing wealth effects. Likewise, the results for the age–tenure distribution of housing wealth effects are qualitatively similar to those obtained from the actual panel, but the differences in housing wealth effects between groups are much more muted. This result should be borne in mind when interpreting housing wealth effects by age and homeownership status estimated using pseudo-panels.

This paper was written in a very collaborative manner, so it is difficult to be precise about each author's contribution. Broadly, Jarkko Jaaskela and myself were the 'senior authors': together we were in the main responsible for planning the paper, deciding on methodology, interpreting and writing-up the results. These tasks were split fairly evenly between Jarkko and myself. Callan Windsor for the most part played the role of 'junior author': he implemented the code in STATA and was responsible for generating the results, although Callan also contributed to the methodology and write-up.

Household saving in Australia (The B.E. Journal of Macroeconomics).

This paper explores the drivers of the rise in household saving in Australia seen over the 2000s using household-level micro data. Although the HILDA Survey is generally regarded as the best panel dataset in Australia, it's spending and income data are not sufficiently comprehensive to be able to calculate accurate saving figures. Instead we used the 2003/04 and 2009/10 Household Expenditure Surveys (HES) from the ABS. These surveys do not constitute a panel (they are repeated cross-sections), but the offsetting benefit is very comprehensive information on expenditure and income, and therefore saving (amongst other things, the ABS use these surveys in constructing the CPI basket). To explore the drivers of the rise in saving we perform cross-sectional regressions on the 2003/04 and 2009/10 surveys, including many control variables, and examine changes in saving patterns over the survey years.

Our results suggest that the rise in household saving between 2003/04 and 2009/10 was driven by changes in behaviour rather than changes in population characteristics: in particular, more educated households, as well as households with high debt and/or wealth increased their propensity to save. Our interpretation of these results is that a reduction in future income growth expectations for more highly educated households after the financial crisis, and an associated effort to rebuild wealth and repay debt, drove the aggregate rise in household saving.

The paper's contribution to the literature consists solely in its results as discussed above, which are the first to examine the recent rise in saving in Australia using micro data. We make no methodological contribution (for methodology we essentially follow Chamon and Prasad, 2010).

This paper followed a fairly traditional 'senior author' (myself) / 'junior author' (Fiona Price) model: I was responsible with formulating the idea for the paper, deciding on methodology and writing up the results; Fiona was responsible for the data manipulation and estimation of the model in STATA.

A state-space approach to Australian GDP measurement (Australia Economic Review).

This paper uses state-space methods to construct new estimates of Australian GDP growth from the published national accounts estimates of expenditure, income and production. Across a range of specifications, our measures are substantially less volatile than headline GDP growth, while also having superior real-time properties and roughly equal value in forecasting models. We conclude that much of the quarter-to-quarter volatility in Australian GDP growth reflects measurement error rather than true shifts in the level of economic activity, with a smoother measure potentially useful for policy-makers looking to abstract from quarter-to-quarter noise.

Our work represents an application to Australian data of the techniques derived by Aruoba et al (2013), who construct a state-space measure of US GDP growth. The Australian dimension of our study is of interest for two reasons, aside from our natural curiosity as Australian researchers. First, whereas the US statistical authorities only construct income and expenditure measures of GDP at a quarterly frequency, the ABS also publishes a production measure. We show that the methods of Aruoba et al (2013) extend to this environment. Second, the Australian economy differs in several respects from that of the United States in ways that may make GDP measurement more challenging. In particular, Australia is a smaller, more trade-exposed economy with a large resource sector. Our results support the idea that these variations in economic structure translate into a different pattern of GDP measurement errors in Australia.

My contribution to this paper was unusually well demarcated: I wrote the first two sections of the paper and supplied the identification proof to show that the Aruoba et al (2013) result could be applied in the Australian context. Although I was not directly involved in data aggregation or code implementation, I also contributed to big-picture decisions on methodology.

Credit supply shocks and the global financial crisis in three small open economies (Journal of Macroeconomics).

This paper explores the contribution of credit supply shocks (as well as other identified shocks) to the evolution of various macroeconomic variables during the global financial crisis in Australia, Canada and the UK. In particular, for each of the three countries mentioned we estimate a sign-restricted VAR on that country and the United States (representing the 'world') that imposes the small open economy assumption (US variables are allowed to impact Australia, for example, but Australian variables cannot impact the US).

We find that negative domestic and foreign credit supply shocks together explain, on average, one-third to one-half of the fall in business credit and rise in credit spreads seen in the three countries during the financial crisis; other identified non-credit-supply shocks explain the rest. Credit supply shocks also explain around one-sixth of the fall in output in the three countries, and one-quarter of the fall initially seen in UK inflation. This suggests that credit supply shocks played an important role in the financial crisis, but not a dominant one.

A number of authors have analysed the importance of credit shocks (see for example Helbling et al., 2011 and Gilchrist et al., 2009 for two widely cited papers). Our analysis is closely related to this growing literature but distinct in its approach. In particular, we use a sign-restricted VAR with domestic and foreign blocks to identify credit supply (and other) shocks in three small open economies, and their impact on credit growth, credit spreads, GDP and inflation, among other key variables. Three other closely related papers, Helbling et al. (2011), Meeks (2012) and Fornari and Stracca (2013), also use sign restrictions to identify credit shocks, but they differ from us in a few important dimensions. First, we identify a rich set of macroeconomic shocks based on a simple argument regarding the effect that various demand, supply and other shocks will have on observed quantities and prices (Helbling et al. identify only credit supply and productivity shocks; Meeks identifies only a credit shock; Fornari and Stracca identify aggregate demand, financial and monetary policy shocks). Our identification strategy allows us to gauge the importance of not just domestic and foreign credit supply shocks, but also domestic and foreign credit demand shocks, as well as other standard shocks that are left largely unidentified in the other papers. Moreover, we use the quantity of credit and credit spreads to identify different types of credit shocks. Helbling et al. also use the quantity of credit, along with credit spreads and default rates, to identify credit shocks, whereas Meeks uses only the (US) credit spread and default rate; Fornari and Stracca follow an alternative strategy by identifying a financial shock that has an impact on the quantity of credit and the relative share price of the financial sector. Finally, we focus on small open economies that are affected by exogenous foreign (US) shocks. Meeks focuses only on the US, while Helbling et al. and Fornari and Stracca consider a larger class of countries that includes both small open economies and large, relatively closed economies without distinguishing between them.

As with the Economica paper, this paper was written in a very collaborative manner, with Jarkko Jaaskela and myself contributing roughly equally along each stage of the research process and to the final published product.

Random fields with Pólya correlation structure (Journal of Applied Probability).

This paper is very different from the preceding four. Rather than being policy-relevant, empirical and focused on Australia, it is an entirely theoretical econometrics paper. In particular, in recognition that the traditional Gaussian assumption is becoming increasingly untenable in finance in particular, and to a lesser extent economics, the paper constructs a new class of highly flexible, non-Gaussian, multi-dimensional random field for use in modelling and estimation. The marginal distribution of the new random field can be taken as any infinitely divisible distribution (this class includes, amongst many others, the Gaussian, Poisson, Gamma, negative binominal and Student's t distributions), while the correlation function can also be specified very flexibly. In addition, the paper fully characterises the random fields in terms of their joint characteristic functions, which allows for efficient estimation via the empirical characteristic function (ECF) method.

Other authors have constructed non-Gaussian random fields. Marfè (2012, 2014), for example, constructs a multivariate Lévy process that can accommodate a flexible range of linear and nonlinear dependencies across the spatial dimension and for which the marginal distribution may approximate any Lévy type. Our construction has a number of advantages over alternatives in the literature, however. The marginal distribution of our random fields may be taken as any infinitely divisible distribution with finite variance, whereas, for example, the marginal distributions of many alternative constructions are restricted to be of normal variance-mixing type and so exclude any non-symmetric distribution or any distribution that does not have support on $(-\infty,\infty)$, such as a distribution on the positive half-line. Our random fields can also be endowed with a rich and dynamic correlation structure across both the spatial and temporal dimensions. Although endowed with a rich dependence structure along the spatial dimension, the Lévy process constructed by Marfè has independent increments, so the increments lack a dependence structure along the time domain. Finally, our method of construction, based on carefully chosen sums of independent and identically distributed (i.i.d.) random variables, lends itself particularly easily to numerical simulation, while the random fields are fully characterized in terms of their joint characteristic function, allowing for efficient estimation.

I was the main author of this paper, with Professor Eugene Seneta providing advice and guidance but leaving most of the formulation of results and write-up to me.

Publication I

First published in Windsor, C., Jaaskela, J. and Finlay, R. (2015). Housing wealth effects: evidence from an Australian panel. *Economica*, **82**, 552-577.

Housing Wealth Effects: Evidence from an Australian Panel

By Callan Windsor†, Jarkko P. Jääskelä† and Richard Finlay‡

† *Reserve Bank of Australia* ‡ *Reserve Bank of Australia and University of New South Wales*

We explore the positive relationship between house prices and household spending by following a panel of Australian households. No evidence for 'traditional housing wealth effects' is found, with young homeowners exhibiting the largest wealth effects. Young renters also exhibit a positive consumption response to house prices, although less so than young homeowners. This suggests that increasing house prices raise spending via easing credit constraints and a common association between house prices and a third factor. Results from a cohort-level panel are similar to those using household-level data, suggesting 'pseudo-panels' may be used as a partial substitute for actual panels.

JEL Classification Numbers: E21, R21, R31

Keywords: house prices, consumption, micro data

INTRODUCTION

Although house prices and consumption tend to move together, understanding the relationship between the two has proven a vexing task for policymakers and commentators.[1] Taking the log difference of the series in Figure 1 and regressing non-housing consumption on house prices implies a marginal propensity to consume (MPC) of around 2½ cents per dollar change in house prices; more sophisticated estimates for Australia suggest a MPC of around 3 cents (see, for example, Dvornak and Kohler 2007).[2]

Figure 1: Real House Prices and Consumption
Per household, June 2003 = 100

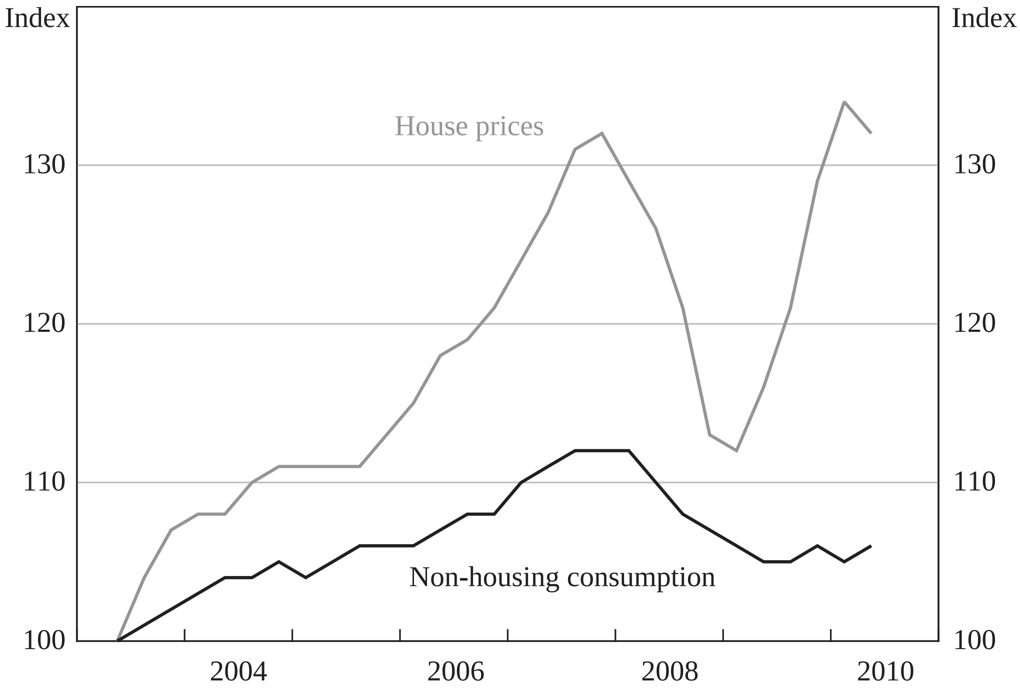

Notes: June 2009/10 dollars; deflated using trimmed mean CPI; consumption excludes rent and other dwelling services

Sources: ABS; RBA; RP Data-Rismark; authors' calculations

Nonetheless, simple bivariate regressions tell us very little about why there is a positive relationship between house prices and consumption. In particular, there are strong arguments against interpreting the aggregate MPC as a ‘traditional housing wealth effect’, whereby spending rises with house prices due to an increase in homeowners’ lifetime resources.

Housing assets, like consumer durables, are different from financial assets. A dwelling is both an asset and a consumption item that provides a stream of services over its lifetime. Accordingly, increasing house prices have a distributional impact on wealth, creating winners and losers whose spending responses may differ. For those who own more housing than they foresee needing in the future (for example, an older household looking to trade down), increasing house prices increase lifetime resources; for those who own less housing than they foresee needing in the future (for example, a young family who will need a larger house in the future), increasing house prices decrease expected lifetime resources. In aggregate, therefore, the causal relationship between house prices and spending is ambiguous, and depends on the MPCs of different groups.

Muellbauer (2007a) shows in a classical model with no credit constraints that it is possible to have a small positive housing wealth effect on non-housing consumption, but this exists only under specific assumptions, and does not exist for measures of aggregate consumption that include imputed housing services. In the context of an overlapping generations model, Buiter (2010) concludes that there is no traditional housing wealth effect on aggregate consumption. Both authors, however, argue that changes in house prices can affect spending when credit constraints are taken into account; for instance financial products enable housing to be used as collateral against which people can borrow to finance consumption. The effect on current spending could be quite large if homeowners were credit constrained before any increase in house prices. In the long run, however, there would be no wealth effect: an increase in house prices would stimulate debt-financed spending in the short run while depressing it in later periods as homeowners repay debt.

Factors other than credit conditions may also affect the relationship between house prices and spending. If older households plan to leave money or even a house to their children or grandchildren, and/or younger households anticipate such bequests, then irrespective of house prices, younger and older households may not perceive any change in their lifetime resources available for spending (see, for instance, Mishkin 2007). Households may also perceive housing wealth as a precautionary saving vehicle against unanticipated future events such as redundancy (Carroll *et al* 2003).[3] Finally, the relationship between housing market turnover and consumption – high turnover leading to increased spending on furnishings, audio-visual equipment and the like – could drive co-movements between house prices and spending, since turnover tends to increase when house prices rise.

Against this background, it is perhaps not surprising that there is no clear consensus on the cause of the correlation between house prices and household spending. Broadly speaking there are three hypotheses that have predictions for how households with certain characteristics should respond to changes in house prices. Under hypothesis (1), increases in perceived house prices raise spending via a traditional housing wealth effect. This channel points to a stronger effect on the spending of older homeowners (who are most likely to own 'excess' housing). Under hypothesis (2), increases in house prices increase in the value of collateral and so loosen credit constraints. This raises spending via the availability of financial products that enable house equity redraws, the refinancing of debt, and/or a reduction in the necessary level of buffer-stock, or precautionary, saving. Younger homeowners are more likely to be credit constrained (Disney *et al* 2010) as well as buffer-stock savers (Gourinchas and Parker 2001). Accordingly, this credit constraints hypothesis suggests a stronger link between house prices and spending for younger homeowners (in what follows we will use the term 'credit constraints hypothesis' to encompass all effects discussed above, including the use of house equity redraws and the buffer-stock saving motive). And under hypothesis (3), house prices and spending are influenced by a common third factor such as something that affects expectations regarding future income. A common influence like higher income expectations, or a reduction in income uncertainty, should have a stronger effect on the spending of younger households, regardless of house tenure status; that is, this hypothesis implies that the spending of young homeowners and young renters should both rise, as both have relatively more years of work ahead of them and so benefit the most from a rise in the wages they may expect to earn in the future. Although income expectations are probably the key common factor, there are other factors – such as shifts in the monetary policy regime and associated changes in inflation expectations – that would also affect the spending of renters. However, these have remained stable over our sample period, which covers 2003 to 2010.

It is difficult to discriminate between these competing hypotheses based on the aggregate relationship between house prices and non-housing consumption, although Aron *et al* (2012) is a notable exception (they show how to distinguish between the collateral, pure wealth effects and common factor hypotheses on aggregate data by employing relevant controls, for instance, for credit conditions). A number of studies have used micro data to understand the co-movement between house prices and consumer spending, but with mixed results. Using a survey of UK households, Attanasio *et al* (2009) argue that income expectations, as per hypothesis (3), have played an important role, because the association between house prices and spending is stronger for younger households irrespective of house tenure type. Using the same UK survey, Campbell and Cocco (2007) draw the opposite conclusion. They find housing wealth effects are largest for older homeowners and lowest for renters. They interpret this heterogeneity in housing wealth effects as being consistent with a traditional wealth effect.

The major methodological difference between Attanasio *et al* (2009) and Campbell and Cocco (2007) is the empirical specification of the model: the former use an equation for the level of consumption while the latter use an equation for consumption growth. Cristini and Sevilla (2014) suggest that differences in the empirical specification may be the cause of the conflicting results. The results presented in this paper are largely based on the methodology employed in Attanasio *et al* (2009), but as a robustness test we also consider the model specified in Campbell and Cocco (2007), and draw very similar conclusions in regard to the likely cause of housing wealth effects.[4] This is not surprising – we estimate household-level fixed-effects regressions in a panel dataset, which if estimated consistently should yield similar results to a first-differenced regression (Wooldridge 2002).

Our findings are most similar to those presented in Muellbauer (2009) and Duca *et al* (2011) who argue that a housing collateral effect is the key to understanding the role of house prices in explaining consumption fluctuations. While Muellbauer (2009) agrees with the results presented by Attanasio *et al* (2009), there is disagreement over interpretation. In addition to the common association between house prices, income innovations and spending, Muellbauer finds that credit constraint effects are positive for young homeowners and negative for the old.

In this paper we use the Household, Income and Labour Dynamics in Australia (HILDA) Survey to examine housing wealth effects, using household-level data for the eight years to 2010. Our analysis contributes to the literature in a number of ways. To begin, we fully exploit the panel nature of our dataset that follows individual households through time (see also Browning *et al* 2013, who use a panel of Danish households to study the relationship between household spending and innovations to house prices). That is, we estimate the response of a household's spending to changes in the perceived price of their house while controlling for unobservable, time-invariant differences between households (such as their level of optimism or thriftiness). Another departure from the identification strategy used in previous papers is the introduction of time-fixed effects to control for common macroeconomic influences. Finally, the availability of panel data allows us to compare hypothesis (2) and hypothesis (3) in a more satisfactory manner than previous papers relying on synthetic panels constructed using a time series of cross-sectional data have been able to do. The distribution of housing wealth effects by homeownership status is potentially biased when using such ‘pseudo panels’ – as acknowledged by Attanasio *et al* (2009) among others – because the split between homeowners and renters is possibly not constant over time and because it is difficult to control for selection in to homeownership with these data.

At the household level, we estimate housing wealth effects that are positive and large for young homeowners and fall to zero for old homeowners; for renters, consumption responses to changing local house prices are positive but small for the young, essentially zero for the middle-aged, and become negative for older households. We suggest that young homeowners’ relatively strong spending response to an increase in house prices supports the credit constraints hypothesis, while the age profile for renters suggests that a third common factor, most likely income expectations, is also playing a role, with higher expected future incomes offsetting higher expected future housing costs for the young, but not for the old (low down-payment constraints may also be important for young renters if they choose to enter the housing market at some point in the future).

We also examine whether these results can be replicated in a more parsimonious, but less informative, model that relies on a pseudo-panel of birth cohorts instead of household-level data. This is done to assess the effect that aggregating may have had on earlier studies using UK data (see Attanasio *et al* 2009). For instance, Muellbauer (2007b) has argued that some of the studies cited above fail to control for cross-sectional variation across households; our dataset allows us to assess this criticism directly. The results for the age distribution of housing wealth effects of homeowners from the cohort pseudo-panel are qualitatively similar to those obtained from the equivalent actual panel, although the necessary use of aggregate house prices rather than self-assessed house prices in pseudo-panels appears to inflate estimated housing wealth effects. This suggests that pseudo-panels may provide a partial, although imperfect, substitute for actual panels when estimating the age distribution of housing wealth effects (at least with the sample used in this paper). Likewise, the results for the age-tenure distribution of housing wealth effects are qualitatively similar to those obtained from the actual panel, but the differences in housing wealth effects between groups are much more muted. This result should be borne in mind when interpreting housing wealth effects by age and homeownership status estimated using pseudo-panels.

The remainder of this paper is set out as follows. Section I introduces the dataset used in this study and presents some stylised features of the variables of interest. Section II presents our methods and Section III details results. Section IV concludes.

I. DATA

The HILDA data

The HILDA Survey is a nationally representative annual household panel of Australian households. It began in 2001 with around 7 700 responding households. It asks questions regarding families, household financial conditions, employment and wellbeing. Special modules provide another layer of detailed information on household wealth every four years.

We define a household as one that does not change tenure type, with households that shift from renting to owning, or vice versa, treated as separate households pre- and post-tenure change. Additionally, we allow moving house to affect spending in a flexible way by treating any homeowner that moves as a new household, and including a dummy for the year of the move (renters who purchase a house are also assigned a moving dummy). This is important because housing transactions may be associated with higher spending and house prices if homeowners purchase goods and services when they trade-up in the housing market. Moving house also provides an easy opportunity to add or reduce housing equity. In general, moving has a small positive effect on spending, although the inclusion or exclusion of the moving dummy does not significantly affect estimated housing wealth effects.

We use two panels of responding households over the period 2003 to 2010. Panel one comprises all homeowners who responded to the survey at least twice over 2003 to 2010. The criteria for selection into panel one are detailed in Table 1: from an initial sample of 57 027 we are left with 34 191 observations.[5] Panel two adds renters to panel one and drops those households who do not live in Sydney, Melbourne or Brisbane, the three largest Australian cities (our postcode-level data on house sale prices, which we will need to use with this panel, only covers these cities).

Table 1: The Panels – 2003–2010		
	Number of observations	
	Added/Dropped	Remaining
Criteria for selection into panel one – homeowners		
Responded in any given year		57 027
Homeowner	-19 706	37 321
Responded at least twice over 2003–2010	-3 130	34 191
Sample size		**34 191**
Criterion for selection into panel two – households		
Add renters	15 473	49 664
Live in Sydney, Melbourne or Brisbane	-28 502	21 162
Sample size		**21 162**

Sources: HILDA Release 10.0; authors' calculations

Demographic variables considered are the age of the household head (the person most likely to make financial decisions for the household), number of children and adults in the household, education, occupation, region of residence and labour force status (see Table A1 for the distribution of these variables). Household financial variables considered include household disposable income, household expenditure and house prices (see Table B1 for the distribution of these variables).

The notion that young homeowners are more likely to be credit constrained than older homeowners is crucial to the interpretation of our results. This argument is supported by the responses of homeowners to a self-assessed measure of credit constraints available within HILDA – the ability to raise $3 000 in an emergency – that are significantly positively correlated with age.[6]

The cross-tabulations in Figure 2, which show the mean ratios of home loans to house prices and the mean ratios of unsecured credit card debts to house prices for panel one (homeowners), also lend support to the idea that younger homeowners are relatively more credit constrained.

Younger homeowners have both high secured and unsecured debt, relative to older homeowners. Given that unsecured debt is likely to be more costly, if younger homeowners were not credit constrained, one would expect them to substitute costly unsecured debt for less costly secured debt, and therefore for the right-hand panel of Figure 2 to show no clear age pattern.

Figure 2: Credit Constraints by Age

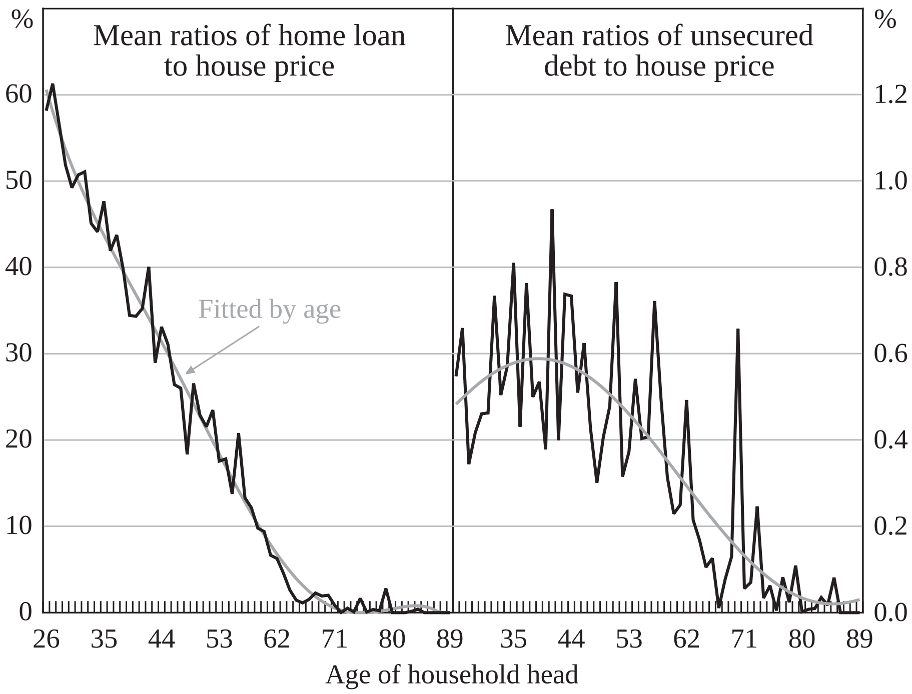

Notes: Calculated using all homeowners in panel one, defined in Table 1; fitted lines obtained by regressing ratios on a polynomial in age; estimated using the 2006 wealth module (results for the 2010 wealth module are very similar)

Sources: HILDA Release 10.0; authors' calculations

The finding that young homeowners have high levels of secured and unsecured debt relative to the value of their homes is important. Disney *et al* (2010) find that only a sub-sample of homeowners with both high secured and unsecured debt increase their indebtedness (to potentially fund spending) following a rise in house prices. Figure 2 show that these homeowners are more likely to be young. The intuition behind the Disney *et al* (2010) result is straightforward: an increase in house price allows a household to refinance – by substituting relatively expensive unsecured debt for secured debt – and potentially borrow more to spend.

The main household financial variables used in this study are self-reported non-housing expenditure and self-reported house prices. These data are discussed in the next two sections.

The HILDA Spending Estimates

The sample covers the period 2003 to 2010. Over this period, the ratio of HILDA non-housing spending to an aggregate measure of household consumption from the Australian Bureau of Statistics (ABS) has been steady at around one-half, and the relationship between movements in the HILDA spending numbers and the aggregate consumption figures has also been broadly stable with a correlation coefficient between the growth rates in these series of around 0.75 (Table 2).[7]

Table 2: Real per Household Spending and Consumption

	2003	2004	2005	2006	2007	2008	2009	2010
HILDA	41 237	41 569	42 576	43 692	43 719	44 273	41 995	42 912
ABS	78 083	79 949	81 807	82 741	85 233	87 403	84 548	84 834
Ratio	0.53	0.52	0.52	0.53	0.51	0.51	0.50	0.51

Notes: 2009/10 dollars; deflated using trimmed mean CPI; HILDA data are from panel one
Sources: ABS; HILDA Release 10.0; authors' calculations

Over the period 2006 to 2010, the HILDA spending estimates were calculated as the sum of 25 self-reported spending categories defined according to the usual amount spent on weekly, monthly and annual items. However, from 2003 to 2005 self-reported figures are only available for three components: meals eaten out, groceries and childcare costs. The relationship between real spending on these items, the age of the household head and real total expenditure in the years 2006 to 2010 was used to impute real total spending for households from

2003 to 2005 (with the imputation performed separately for owners and renters). The estimated imputation regressions for panel one are presented in Table 3, where *total spending*$_{it}$ is real total spending by household i in time t, meo_{it} is real spending on meals eaten out, gro_{it} is real spending on groceries, cc_{it} is real spending on childcare costs and age_{it} is the age of the household head.

Table 3: Spending Imputation

$$total\ spending_{it} = \alpha_0 + \alpha_1 meo_{it} + \alpha_2 gro_{it} + \alpha_3 cc_{it} + \alpha_4 age_{it} + \alpha_5 age^2_{it} + E_{it}$$

	Linear model	Log-linear model[(a)]
Meals eaten out	3.31***	0.58***
Groceries	2.15***	0.51***
Childcare costs	0.30***	0.06***
Age	895.88***	0.03***
Age squared	–9.93***	–0.0004***
Constant	–3629.61*	9.28***
Obs (2006–2010)	21 475	21 475
R^2	0.29	0.49
RMSE[(b)]	0.90	0.48

Notes: 2009/10 dollars; deflated using trimmed mean CPI; regression output above is for panel one; ***, ** and * indicate significance at the 1, 5 and 10 per cent level, respectively; robust standard errors

(a) Coefficients on meals eaten out, groceries and childcare costs show the expected percentage change in total spending from a $100 increase in spending on these items, holding all other predictors constant

(b) Data has been transformed to make the *RMSEs* comparable across models

Sources: HILDA Release 10.0; authors' calculations

The first column shows the estimated coefficients from a linear specification, and the second column reports the estimated coefficients from a log-linear specification. Based on tests of functional form, the log-linear model was chosen to impute spending in years 2003 to 2005. The results from this regression are consistent with other papers implementing a similar imputation method (see, for example, Skinner 1989; Lehnert 2004; and Contreras and Nichols 2010).

This point notwithstanding, the results presented in Section III are also robust to restricting the sample to the period 2006 to 2010, when no imputation is necessary to obtain a measure of total expenditure in the HILDA Survey.

HILDA Self-reported House Prices

The house price variable used throughout this analysis is a household's self-reported house price every year from 2003 to 2010.[8] To check the consistency of these self-reported house prices we compare the mean of all self-reported house prices in each period to an independent nationwide measure of house prices (Table 4). The series appear to move together closely, albeit with a level difference; the correlation coefficient between the growth rates in each series is around 0.65. There is, however, one notable exception: the self-reported house price series misses the decline in nationwide prices that occurred between 2008 and 2009.

Table 4: Self-reported House Prices and Independent House Prices

$'000

	2003	2004	2005	2006	2007	2008	2009	2010
HILDA[(a)]	356	398	416	458	489	509	524	559
Aggregate[(b)]	316	352	371	404	446	462	441	517
Ratio	1.13	1.13	1.12	1.13	1.10	1.10	1.19	1.08

Notes: (a) Unweighted mean from panel one

(b) Calculated as the total value of household dwelling assets from RBA Statistical Table B20 (Selected Assets and Liabilities of the Private Non-financial Sectors), divided by the number of dwellings owned by households

Sources: ABS; HILDA Release 10.0; RBA; authors' calculations

For Australia's three largest cities – Sydney, Melbourne and Brisbane – Figure 3 shows a scatter-plot of median self-assessed house prices by postcode over the period 2003 to 2010 against independent house prices from a near-census of sales data. The two line up very closely, with an R^2 of around 0.6, suggesting that homeowners are a good judge of the true price of their houses.

Figure 3: Self-assessed House Prices vs. Sale Prices

2003 to 2010; postcode level

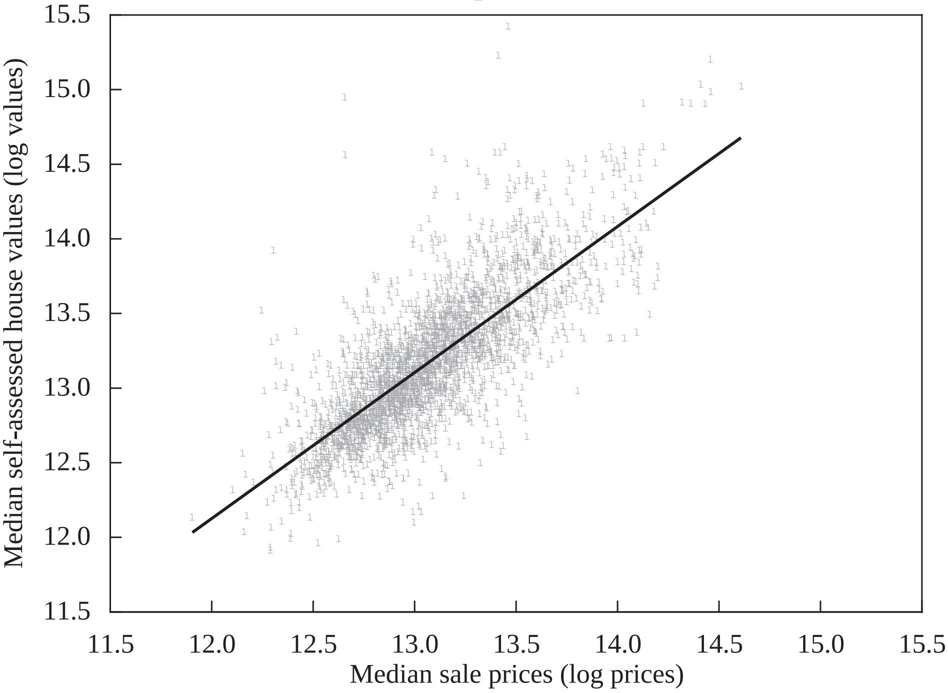

Sources: Australian Property Monitors (APM); HILDA Release 10.0; authors' calculations

Finally, Figure 4 plots the mean of self-reported house prices within 12 major statistical regions for panel one (see Table 1). These show large variations over time, across cities, and between cities and regions.

Figure 4: Self-reported House Prices
By major statistical region

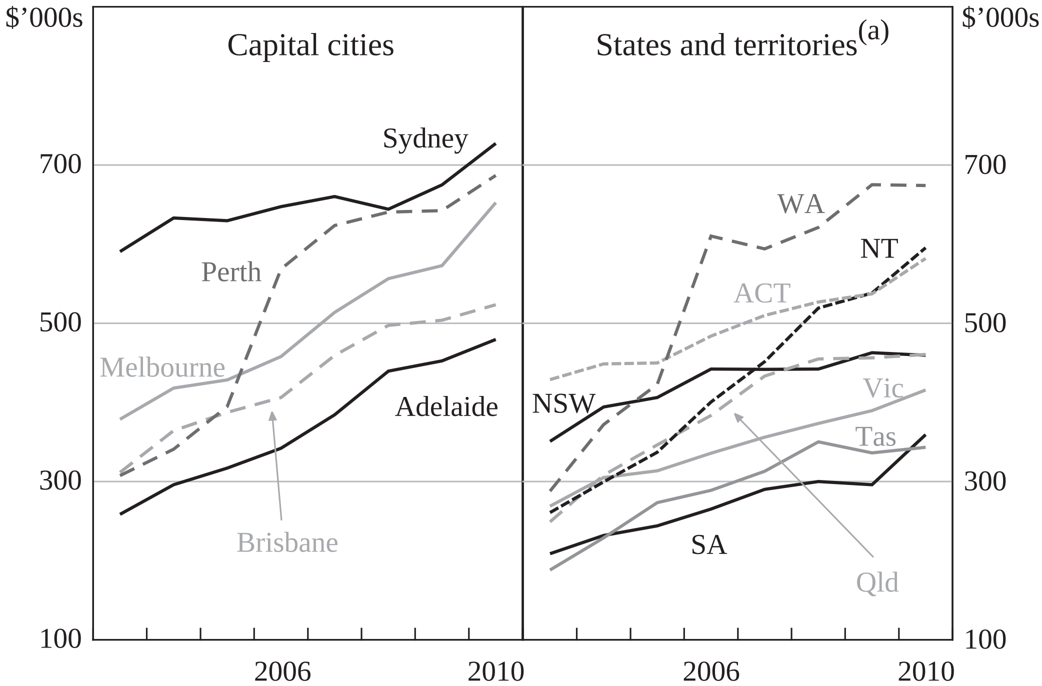

Note: (a) States exclude respective capital cities

Sources: HILDA Release 10.0; authors' calculations

II. METHODOLOGY

To study the nature of housing wealth effects, we use the HILDA panel and the framework proposed by Attanasio and Weber (1994) (see also Attanasio *et al* 2009). Specifically, we estimate housing wealth effects by examining how the value of goods and services consumed by households responds to changes in house prices, controlling for a number of other factors such as education levels and income. Further, we split households up into young, middle and old households in order to examine differences in housing wealth effects between age groups. The estimation is performed using a panel regression where each household's house price and spending level is tracked through time.

The main advantage of our study over previous studies is that we use an actual panel rather than a pseudo-panel of birth cohorts constructed from a series of cross-sections. This enables us to move by degrees – from household-level data to cohort-level data – by first tracking the same households through time and then tracking the same 'cohorts' (defined as a group of households with fixed membership) through time. By doing so, any differences in results due to different levels of data aggregation can be identified.

The appeal of Attanasio and Weber's framework is the lack of structure it imposes upon empirically estimated relationships, although it can be seen as an approximation to the life-cycle model. The life-cycle model predicts that real spending is equal to an annuity value of lifetime resources and its interaction with the life-cycle of the household:

$$\left(total\ spending_{it}\right) = \omega_{it}\kappa(lifecycle_{it})\ exp(E_{it})\ ,$$

where $total\ spending_{it}$ is real annual spending of household i at time t, and ω_{it} is some fraction of total wealth that includes, for instance, financial wealth and housing wealth. The function $\kappa(lifecycle_{it})$ captures the age and composition of household members. What is left unexplained, $exp(E_{it})$, is unexplained variation in lifetime earnings including temporary shocks/measurement error in current earnings. Taking logs of the above equation yields:

$$(1) \qquad ln(total\ spending_{it}) = ln(\omega_{it}) + ln(\kappa(life\ cycle_{it})) + E_{it}.$$

Equation (1) can be estimated using proxies for log lifetime wealth $ln(\omega_{it})$ and for the life-cycle function $\kappa(lifecycle_{it})$ as per Equation (2):

$$(2) \qquad ln(total\ spending_{it}) = \alpha_i + B'W_{it} + A'Z_{it} + E_{it}.$$

Log lifetime wealth is proxied with the constant α_i and a vector of variables, W_{it}, which includes: dummy variables for the highest level of education achieved by the household head;[9] the occupational classification of the household head; and the log of real household disposable income, $HHDY_{it}$ (we also include the log of real housing wealth, detailed below, but for presentational purposes we consider it separately from the other wealth variables contained in W_{it}). The coefficients in vector B will represent a log-level shift in spending for changes in categorical variables and, for the continuous variables, the elasticity of spending. The life-cycle function is proxied with a vector of variables, Z_{it}, including: the number of adults and the number of children in the household; a dummy for households with three or more adult members; labour force status of the household head; and region of residence. Finally, when estimating Equation (2) time dummies are included to control for common macro factors and moving dummies are included to control for the effect of moving houses on spending.

While Attanasio *et al* (2009) exclude current income from their lifetime wealth controls, we include it here for two reasons. First, excluding current income could bias estimated housing wealth effects to the extent that current income is correlated with house prices.[10] Second, for many households current income may be the best estimate of permanent income and hence should be an important determinant of spending. However, in testing our third hypothesis – that of common factors – it is arguably the case that one should not control for current income because of the information it contains about innovations to permanent income. Against this background, it is worth noting that our results are robust to the exclusion of current income from the vector of lifetime controls.[11]

The impact of changes in real house prices on spending is the key focus of this paper. We use self-assessed house prices from homeowners' responses to the following question:

Do you know what the approximate value of your house is? I mean, how much would it bring if you sold it today? Include land, house improvements, and fixtures (such as curtains and light fittings) usually sold with a home. Exclude house contents.

The house-price variable, $\ln(HP_{it})$, is then interacted with a vector of dummies, Age_i, indicating the age group of the household head in the first year they were surveyed as either young (up to 39 years), middle (40 to 54 years) or old (over 55 years).

It is the existence, or lack thereof, of differences in housing wealth effects across different age groups that will allow us to distinguish between the various hypotheses put forward for the cause of these wealth effects. Larger housing wealth effects for older homeowners would be consistent with a traditional wealth effect, while larger effects for younger homeowners could reflect credit constraints and/or common factors.

To examine the relative merit of these latter two explanations, the panel including renters is considered (panel two in Table 1) and the term $C'(ln(HP_{it}) \times Age_i \times Tenure_i)$ is added to the baseline model, where $Tenure_i$ is a dummy variable indicating the tenure type of the household. If housing wealth effects are entirely due to common factors then consumption should increase for young renters as well as young homeowners following an increase in house prices within their neighbourhood. If these effects are due entirely to credit constraints then the consumption of young homeowners should again increase, but the consumption of young renters should not.

For ease of interpretation, we present results in Section III in a form that is comparable to the aggregate MPCs discussed in the introduction and commonly referred to in the literature. Estimated elasticities for each age group are converted to MPCs by multiplying the elasticities by the sample average ratio of non-housing consumption to housing wealth over the period 2003 to 2010 for each age group.[12]

III. RESULTS

Household-level Analysis

At the household level, Equation (2) becomes:

$$(3) \qquad ln\left(total\ spending_{it}\right) = \alpha_i + T_t + M_{it} + B^{'}W_{it} + D^{'}(ln(HP_{it}) \times Age_i) + A^{'}Z_{it} + E_{it},$$

where α_i are household-fixed effects that control for unobserved time-invariant differences between households, T_t are time-fixed effects and M_{it} are moving dummies. In this specification, we omit dummies for occupation and education from our wealth term since these are typically time-invariant and so are captured by the household-fixed effect.

For panel one (our panel of homeowners), we estimate the model using self-assessed house prices, ln(HP_{it}), where the coefficient on house prices is allowed to vary across age groups (full regression output is given as Model 1 in Table C1).[13] For our panel of homeowners, Figure 5 compares predicted real spending from this model (dashed lines) to actual spending (solid lines), using real spending averages over the eight years to 2010 for all homeowners within each age group. This allows us to assess the functional form of the model by examining whether the life-cycle pattern of spending follows a hump-shape; such patterns are well-known and widely reported in the literature – see, for instance, Attanasio and Weber (2010). From a visual examination it seems that this specification provides a good fit to the data in the spending-age space for each age group.

Figure 5: Real Household Spending
By age within age group

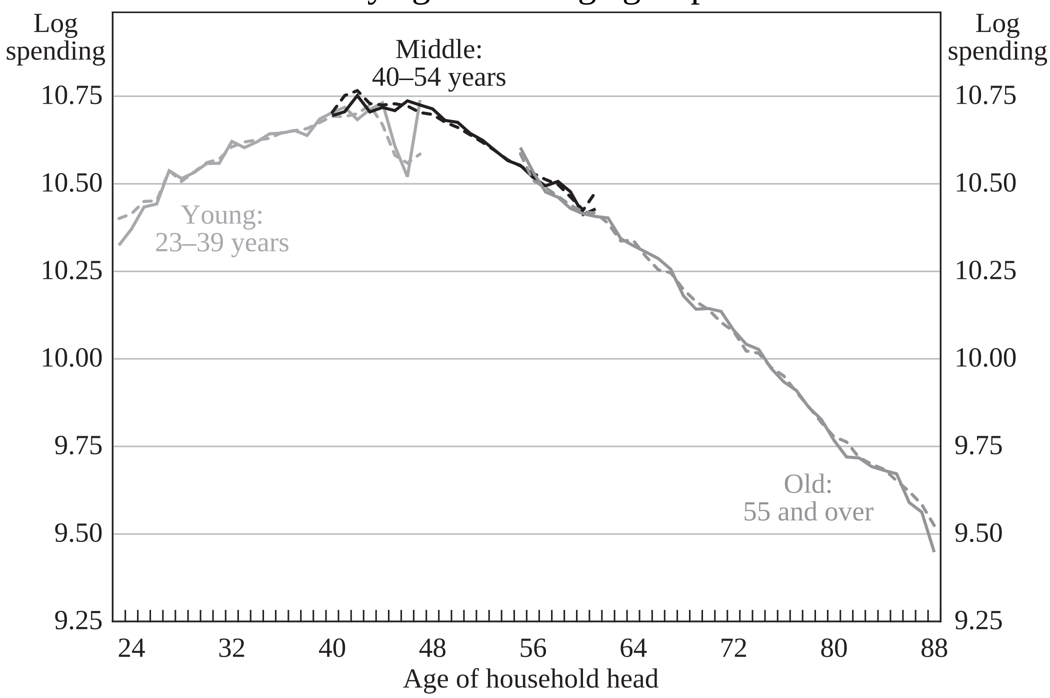

Notes: 2009/10 dollars; deflated using trimmed mean CPI; data are for panel one, defined in Table 1; dashed lines are fitted values

Sources: HILDA Release 10.0; authors' calculations

The first column of Table 5 shows that spending responses to a change in house prices are estimated to be largest (and most statistically significant) for young homeowners at around 3 cents per dollar. For old homeowners we find no significant spending responses. The difference between spending responses for young and middle-aged homeowners is statistically significant, as is the difference between young and old homeowners. The second column of Table 5 shows that the age distribution is not sensitive to restricting spending to non-durable items.[14] Likewise, the third column shows that the results are robust to restricting the sample to the period 2006 to 2010, when no imputation for total spending is necessary (full regression output is given as Model 1, Model 3 and Model 4 in Table C1). Although not shown, all the results presented henceforth are also robust to this restriction (results available upon request).

Table 5: Household-level Housing Wealth Effects: Homeowners

Cents per dollar change in wealth

	Total spending	Non-durable spending	Non-imputed spending
ln(*HP_{it}*)			
Young	3.13***	2.79***	2.86**
Middle	1.08***	0.93***	0.57
Old	-0.06	-0.05	0.22
H_0: Young = Middle[(a)]	R***	R***	R*
H_0: Young = Old[(b)]	R***	R***	R**
H_0: Middle = Old[(c)]	R**	R**	F

Notes: R refers to a rejection of the null hypothesis H_0, F refers to a failure to reject H_0; ***, ** and * indicate significance at the 1, 5 and 10 per cent level, respectively

(a) H_0 is that housing wealth effects for young and middle-aged homeowners are not statistically different from one another

(b) H_0 is that housing wealth effects for young and old homeowners are not statistically different from one another

(c) H_0 is that housing wealth effects for middle-aged and old homeowners are not statistically different from one another

Sources: HILDA Release 10.0; authors' calculations

The finding of large and significant housing wealth effects for younger households is consistent with the credit constraints hypothesis and the common third factor hypothesis. To distinguish between these hypotheses, the panel with renters is considered (panel two from Table 1). Of course renters do not provide a self-assessed house price. To address this issue, all households are attributed with the log of average house prices within their postcode, derived from a near-census of sales data.[15] However, these data are only available for Australia's three largest cities – Sydney, Melbourne and Brisbane. Therefore, the results presented in Table 6 are restricted to households residing in these cities.

Table 6: Household-level Housing Wealth Effects: Homeowners and Renters

Cents per dollar change in wealth

	Total spending	Non-durable spending	Total spending: excluding $\ln(HHDY_{it})$	Non-durable spending: excluding $\ln(HHDY_{it})$
$\ln(HP_{it})$				
Young owner	6.34***	5.16***	6.30***	5.13***
Middle owner	2.30*	1.32	2.38*	1.39
Old owner	-0.05	-0.31	-0.09	-0.35
Young renter	2.45***	1.68*	2.45***	1.68*
Middle renter	0.20	0.04	0.25	0.08
Old renter	-3.91**	-1.67	-3.88**	-1.64
H_0: Young = Middle[(a)]	R*	R*	R*	R*
H_0: Young = Old[(b)]	R***	R***	R***	R***
H_0: Middle = Old[(c)]	F	F	F	F
H_0: Young renter = Young owner[(d)]	R**	R*	R**	R*
H_0: Middle renter = Middle owner[(d)]	F	F	F	F
H_0: Old renter = Old renter[(d)]	R*	F	R*	F

Notes: R refers to a rejection of the null hypothesis H_0, F refers to a failure to reject H_0; ***, ** and * indicate significance at the 1, 5 and 10 per cent level, respectively

(a) H_0 is that housing wealth effects for young and middle-aged homeowners are not statistically different from one another

(b) H_0 is that housing wealth effects for young and old homeowners are not statistically different from one another

(c) H_0 is that housing wealth effects for middle-aged and old homeowners are not statistically different from one another

(d) H_0 is that housing wealth effects for renters in a given age group are not statistically different from owners in the same age group

Sources: HILDA Release 10.0; authors' calculations

The results again show that housing wealth effects for homeowners fall with age; housing wealth effects for young homeowners are positive, large and statistically significantly different from those for middle-aged and old homeowners. For renters we find a similar age pattern: the consumption response to an increase in house prices is positive and significant for young renters, indistinguishable from zero for middle-aged renters, and negative for old renters (see Model 5 and Model 6 in Table C1 for full regression output; the larger housing wealth effects for homeowners in this model are largely due to the use of aggregate house prices, with aggregate prices tending to inflate the estimated effects).

The results for renters suggest some role for a third common factor such as income expectations in addition to credit constraints, but not a dominant one. The spending response of young renters to an increase in house prices is positive, but statistically significantly less than that of young homeowners, suggesting that a third common factor is not the main driver of spending choices. The age pattern of the response for renters, however – older renters reduce their spending in response to a house price increase, while middle-aged renters do not respond and young renters increase their spending – does suggest that a common factor is having an impact, with the negative impact of rising house prices on renters (for example via higher future expected rental costs) offset by higher income expectations for younger renters, but not for older renters. This observation is supported by the distribution of (proxies for) income expectations by age for renters: the share of old renters in our sample headed by a person without a secondary education is around 50 per cent, versus 25 per cent for middle-aged renters and 20 per cent for young renters; likewise, the share of old renters unemployed or marginally attached to the labour force is 41 per cent, versus 17 per cent for middle-aged renters and 16 per cent for young renters.[16]

Cohort-level Analysis

In this section the results are replicated using the synthetic cohort techniques applied by Attanasio *et al* (2009). While this is a less informative dataset, comparing results at different levels of data aggregation allows the effect of aggregating data on model results to be examined. However, it is important to acknowledge that the HILDA data set retains its panel structure even if panel data methods are not applied to it. Accordingly, our comparison abstracts from the additional noise associated with putting together a time series of cross-sectional data, which could also affect results.

This approach controls for unobservable time-constant differences between *cohorts* rather than *households*, thereby reducing the number of parameters in Equation (2). Fourteen birth cohorts are defined, from before 1926 to 1995, and are entered into Equation (2) as dummy variables (Table 7).

Table 7: Households per Cohort

Data are for panel two

Cohort dummy	Birth year	Average cohort size by year
Cohort 0	1986 to 1995	172
Cohort 1	1981 to 1985	383
Cohort 2	1976 to 1980	434
Cohort 3	1971 to 1975	541
Cohort 4	1966 to 1970	614
Cohort 5	1961 to 1965	694
Cohort 6	1956 to 1960	698
Cohort 7	1951 to 1955	582
Cohort 8	1946 to 1950	525
Cohort 9	1941 to 1945	418
Cohort 10	1936 to 1940	336
Cohort 11	1931 to 1935	289
Cohort 12	1926 to 1930	277
Cohort 13	Pre-1926	248

Sources: HILDA Release 10.0; authors' calculations

With these cohorts, the model becomes:

$$\text{(4)} \qquad ln(total\ spending^{c}_{it}) = \alpha_c + T_t + M_{it} + B^{'}W_{it} + D^{'}(ln(HP_{it}) \times Age_i) + A^{'}Z_{it} + u^{c}_{it} + E_{it}\ ,$$

where α_c are the cohort dummies (for $c = 0,\ldots,13$) and u_{it}^{c} is household i's deviation from the cohort average. Again our proxy for wealth includes the highest education level of the household head, the occupation of the household head, real disposable income and real housing assets.

When estimating Equation (4), $u_{it}^{c} + E_{it}$ is treated as a composite error term which is uncorrelated with the explanatory variables. If self-assessed house prices are used, however, this assumption becomes tenuous. It is likely that the cohort dummies, α_c, capturing unobserved cohort heterogeneity, and households' deviations from these values, u_{it}^{c}, will be correlated with self-assessed house prices. In this case, estimates of housing wealth effects will be biased. In light of this we estimate Equation (4) using independent house prices (again taken as the log of average sale prices within a household's postcode; and again we drop households who do not live in Sydney, Melbourne or Brisbane). This breaks the link between a household's house price and any unobserved household heterogeneity.[17] (This is another advantage of using panel data in that it allows one to use self-assessed house prices, which are more relevant when estimating housing wealth effects).

The fit of the model using independent house prices is examined in Figure 6, which compares predicted real spending from Equation (4), averaged across birth cohorts, to actual spending. It seems that this specification provides a good fit to the data within the spending-age space for each cohort (regression output is given in Table D1).

Figure 6: Real Household Spending

By age within birth cohort

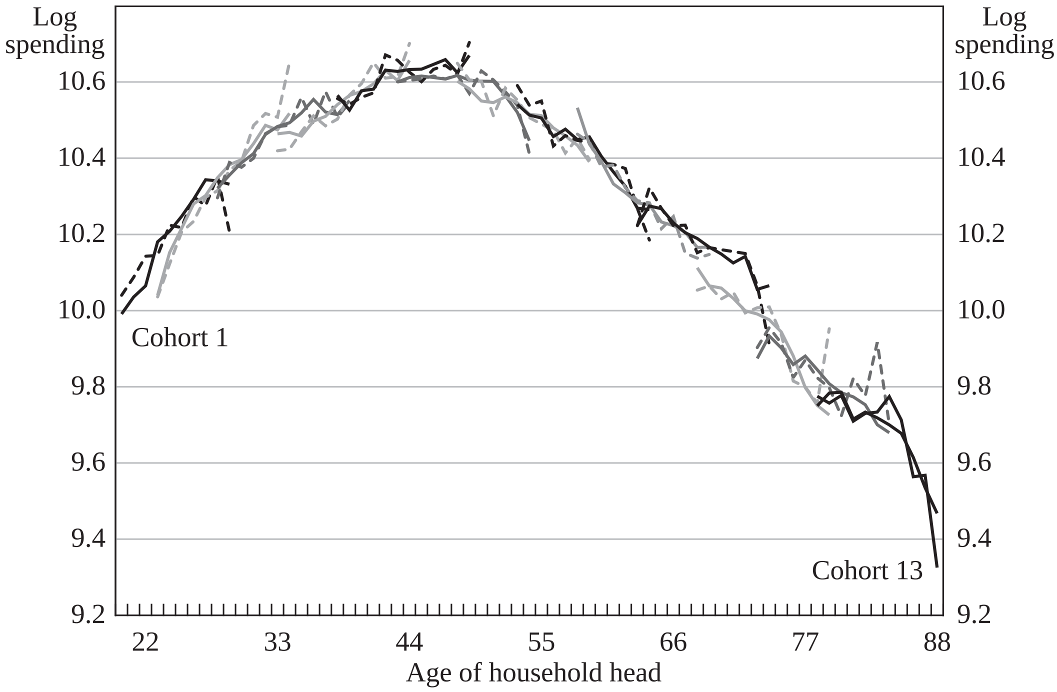

Notes: 2009/10 dollars; deflated using trimmed mean CPI; data are for panel one households living in Sydney, Melbourne or Brisbane, defined in Table 1; series represent birth cohorts as defined in Table 7; dashed lines are fitted values

Sources: APM; HILDA Release 10.0; authors' calculations

Table 8 shows the coefficients on house prices across the age distribution using panel one (homeowners), and for comparison, the equivalent coefficients from a household-level regression using independent house prices. Similar to the results of Table 5, the housing wealth effects estimated in the cohort model decline with age; differences in housing wealth effects by age are less pronounced in the cohort model, however, and older homeowners now show significantly positive responses. In particular, housing wealth effects across the age distribution as estimated in the cohort model are not statistically significantly different from each other. This suggests that 'pseudo-panels' can act as partial, although imperfect, substitutes for actual panels. Moreover, the use of aggregate house prices appears to inflate estimated spending responses, evidenced in a comparison between the second column of Table 8 and the first column of Table 5.

Table 8: Cohort-level Housing Wealth Effects by Age

Cents per dollar change in wealth

	Cohort regression	Household-level regression
	Aggregate house prices	Aggregate house prices
ln(*HP*$_{it}$)		
Young	4.17***	6.78***
Middle	3.87***	2.85**
Old	2.93***	0.36
H_0: Young = Middle[(a)]	F	R*
H_0: Young = Old[(b)]	F	R***
H_0: Middle = Old[(c)]	F	F

Notes: R refers to a rejection of the null hypothesis H_0, F refers to a failure to reject H_0; ***, ** and * indicate significance at the 1, 5 and 10 per cent level, respectively

(a) H_0 is that housing wealth effects for young and middle-aged homeowners are not statistically different from one another

(b) H_0 is that housing wealth effects for young and old homeowners are not statistically different from one another

(c) H_0 is that housing wealth effects for middle-aged and old homeowners are not statistically different from one another

Sources: APM; HILDA Release 10.0; authors' calculations

Finally, the cohort model is used to examine the age-tenure distribution of housing wealth effects. The results presented in the first column of Table 9 are qualitatively consistent with those estimated from the equivalent actual panel (these results are repeated in the second column) in that the consumption response to an increase in house prices falls with age for both owners and renters, although this pattern is considerably dampened in the cohort framework. These weaker results may in part be caused by biases associated with splitting the sample by age and homeownership status in the cohort model: within the cohort framework it is difficult to control for selection in to homeownership, whereas in the actual panel one can flexibly control for homeownership choice.

Table 9: Cohort-level Housing Wealth Effects: Homeowners and Renters

Cents per dollar change in wealth

	Cohort regression	**Household-level regression**
	Aggregate house prices	Aggregate house prices
ln(*HP_{it}*)		
Young owner	4.70***	6.34***
Middle owner	4.15***	2.30*
Old owner	2.90***	-0.05
Young renter	3.29***	2.45***
Middle renter	2.74***	0.20
Old renter	1.69**	-3.91**
H_0: Young = Middle[(a)]	F	R*
H_0: Young = Old[(b)]	F	R***
H_0: Middle = Old[(c)]	F	F
H_0: Young renter = Young owner[(d)]	R**	R**
H_0: Middle renter = Middle owner[(d)]	R**	F
H_0: Old renter = Old renter[(d)]	R***	R*

Notes: R refers to a rejection of the null hypothesis H_0, F refers to a failure to reject H_0; ***, ** and * indicate significance at the 1, 5 and 10 per cent level, respectively

(a) H_0 is that housing wealth effects for young and middle-aged homeowners are not statistically different from one another

(b) H_0 is that housing wealth effects for young and old homeowners are not statistically different from one another

(c) H_0 is that housing wealth effects for middle-aged and old homeowners are not statistically different from one another

(d) H_0 is that housing wealth effects for renters in a given age group are not statistically different from owners in the same age group

Sources: HILDA Release 10.0; authors' calculations

IV. CONCLUSION

We use a household-level dataset, the HILDA Survey, to explore the relationship between house prices and household spending in Australia. Three main arguments have been put forward in the literature to explain the apparent co-movement between house prices and spending: (1) a 'traditional housing wealth effect', whereby spending rises with house prices due to an increase in households' lifetime resources; (2) the removal of credit constraints, whereby spending rises with house prices due to households' ability to borrow more, given more valuable collateral, and the related buffer-stock savings argument, whereby higher house prices act as a form of precautionary savings for low-saving households, allowing them to increase spending; and (3) that spending and house prices move together due to a common third factor, such as changing perceptions of lifetime income.

Our analysis most strongly supports a combination of the second and third explanations – that credit constraints and/or buffer-stock saving, in combination with a common association between house prices and a third factor, most likely income expectations, provide the vehicle through which house prices affect spending. At both the cohort and household level we find that the spending by younger (and so more credit constrained) households is more responsive to changes in house prices than that of older households. This argues against the traditional housing wealth effect, which should be stronger for older households who typically own more housing than they will need over their remaining lifetimes. We also find that young homeowners respond more than young renters to rising house prices. This suggests that a common third factor is not the dominant source of the observed housing wealth effects. The age pattern of spending responses for renters, however – rising house prices reduce the spending of older renters, leave the spending of middle-aged renters unchanged, and increase the spending of young renters – suggests a role for income expectations or another similar factor.

Finally, by analysing the same dataset at two different levels of aggregation we are able to assess the effect that aggregating data has on model results. We find that household-level and cohort regressions imply qualitatively similar spending reactions in response to changes in house prices across age groups, although the age patterns are less marked in the cohort framework. This suggests that 'pseudo-panels' are a partial, although imperfect, substitute for actual panels. The necessary use of aggregate house prices in pseudo-panels also appears to inflate estimated housing wealth effects.

REFERENCES

ARON J., DUCA JV., MUELLBAUER J., MURATA K. and MURPHY A (2012). Credit, Housing Collateral, and Consumption: Evidence from Japan, the U.K., and the U.S. *The Review of Income and Wealth*, **58(3)**, 397–423.

ATTANASIO OP., BLOW L., HAMILTON R. and LEICESTER A (2009). Booms and Busts: Consumption, House Prices and Expectations. *Economica*, **76(301)**, 20–50.

ATTANASIO OP. and WEBER G (1994). The UK Consumption Boom of the Late 1980s: Aggregate Implications of Microeconomic Evidence. *The Economic Journal*, **104(427)**, 1269–1302.

ATTANASIO OP. and WEBER G (2010). Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy. *Journal of Economic Literature*, **48(3)**, 693–751.

BENITO A. (2006), Does Job Insecurity Affect Household Consumption. *Oxford Economic Papers*, **58**, 157–181.

BROWNING M., GORTZ M. and LETH-PETERSEN S (2013). Housing Wealth and Consumption: A Micro Panel Study. *The Economic Journal*, **123**(**May**), 401–28.

BUITER WH (2010). Housing Wealth isn't Wealth. *Economics: The Open-Access, Open-Assessment E-Journal*, **4(2010-22)**.

CAMPBELL JY. and COCCO JF (2007). How do House Prices Affect Consumption? Evidence from Micro Data. *Journal of Monetary Economics*, **54(3)**, 591–621.

CARROLL CD., DYNAN KE. and KRANE SD (2003). Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets. *The Review of Economics and Statistics*, **85(3)**, 586–604.

CONTRERAS J. and NICHOLS J (2010). Consumption Responses to Permanent and Transitory Shocks to House Appreciation. Board of Governors of the Federal Reserve System Finance and Economics Discussion Series, **No 2010-32**.

CRISTINI A. and SEVILLA A (2014). Do House Prices Affect Consumption? A Re-assessment of the Wealth Hypothesis. *Economica*, **81(323)**.

DISNEY R., BRIDGES S. and GATHERGOOD J (2010). House Price Shocks and Household Indebtedness in the United Kingdom. *Economica*, **77(307)**, 472–96.

DUCA J, MUELLBAUER J and MURPHY A (2011). House Prices and Credit Constraints: Making Sense of the US Experience. *The Economic Journal*, **121(552)**, 533–51.

DVORNAK N and KOHLER M (2007). Housing Wealth, Stock Market Wealth and Consumption: A Panel Analysis for Australia. *Economic Record,* **83**, 117–130.

GOURINCHAS P-O. and PARKER JA (2001). The Empirical Evidence of Precautionary Saving. *The American Economic Review*, **91(2)**, 406–12.

LEHNERT A (2004). Housing, Consumption, and Credit Constraints. Board of Governors of the Federal Reserve System Finance and Economics Discussion Series, **No 2004-63**.

MISHKIN FS (2007). Housing and the Monetary Transmission Mechanism. NBER Working Paper, **No 13518**.

MUELLBAUER J (2007a). Housing, Credit and Consumer Expenditure, in Housing, Housing Finance, and Monetary Policy. A Symposium sponsored by the Federal Reserve Bank of Kansas City, Kansas City, 267–334.

MUELLBAUER J (2007b). Housing and Personal Wealth in a Global Context. WIDER (World Institute for Development Economics Research) Research Paper, **No 2007/27**.

MUELLBAUER J (2009). Household Decisions, Credit Markets and the Macroeconomy: Implications for the Design of Central Bank Models. Paper prepared for the Eighth BIS Annual Conference. Financial System and Macroeconomic Resilience: Revisited, Basel, 25–26 June.

SKINNER J (1989). Housing Wealth and Aggregate Saving. *Regional Science & Urban Economics*, **19(2)**, 305–24.

WOOLDRIDGE JM (2002), Econometric Analysis of Cross Section and Panel Data, MIT Press, Cambridge.

WINDSOR C, JÄÄSKELÄ J and FINLAY R (2013), Home Prices and Household Spending. *RBA Discussion Paper*, 2013-04.

APPENDIX A: DEMOGRAPHICS OF THE HOUSEHOLD HEAD

Table A1: Demographics of the Household Head
All households (panel one plus renters)
(*continued next page*)

	Obs	Mean	
Age groups			
All ages	49 664	49	
Young	18 151	32	
Middle	15 821	49	
Old	15 692	70	
	Obs	**Per cent**	**Cumulative per cent**
Education			
Postgraduate – masters or doctorate	2 217	4	4
Grad diploma/grad certificate	2 863	6	10
Bachelor or honours	6 576	13	23
Advanced diploma	4 628	9	33
Certificate III or IV	11 288	23	56
Certificate I or II	643	1	57
Certificate not defined	272	1	57
Year 12	5 991	12	69
Year 11 and below	15 166	31	100
Undetermined	20	0	100
Occupation			
Non-response	16 290	33	33
Agriculture, forestry and fishing	1 231	2	35
Mining	731	1	37
Manufacturing	3 655	7	44
Electricity, gas, water and waste services	403	1	45
Construction	2 942	6	51
Wholesale trade	1 145	2	53
Retail trade	2 362	5	58
Accommodation and food services	1 353	3	61
Transport, postal and warehousing	1 790	4	64
Information, media and telecommunications	881	2	66
Financial and insurance services	1 281	3	69
Rental, hiring and real estate services	461	1	70

Table A1: Demographics of the Household Head
All households (panel one plus renters)
(*continued next page*)

	Obs	Per cent	Cumulative per cent
Occupation			
Professional, scientific and tech services	2 593	5	75
Administrative and support services	900	2	77
Public administration and safety	2 796	6	82
Education and training	3 256	7	89
Health care and social assistance	3 879	8	97
Arts and recreation services	530	1	98
Other services	1 185	2	100
Labour force status			
Employed full time (FT)	26 721	54	54
Employed part time (PT)	6 813	14	68
Unemployed looking for FT work	803	2	69
Unemployed looking for PT work	271	1	70
Not in labour force, marginally attached	1 871	4	73
Not in labour force, not marginally attached	13 170	27	100
Employed, but usual hours worked unknown	15	0	100
Number of adults			
1	16 561	33	33
2	24 495	49	83
3	5 511	11	94
4	2 415	5	99
5	565	1	100
6	92	0	100
7 or more	25	0	100
Number of children aged 0–14			
0	35 198	71	71
1	5 909	12	83
2	5 870	12	95
3	2 028	4	99
4	476	1	100
5	118	0	100
6 or more	65	0	100

Table A1: Demographics of the Household Head
All households (panel one plus renters)
(*continued*)

	Obs	Per cent	Cumulative per cent
Region of residence			
Sydney	8 144	16	16
NSW excluding Sydney	6 807	14	30
Melbourne	8 514	17	47
Vic excluding Melbourne	3 661	7	55
Brisbane	4 504	9	64
Qld excluding Brisbane	5 708	11	75
Perth	3 214	6	82
WA excluding Perth	1 481	3	85
Adelaide	3 500	7	92
SA excluding Adelaide	1 292	3	94
ACT	1 587	3	97
NT	286	1	98
Tasmania	966	2	100

Sources: HILDA Release 10.0; authors' calculations

APPENDIX B: HOUSEHOLD FINANCES

Table B1: Household Finances

All households (panel one plus renters)

	Mean ($'000, 2009/10 dollars)		
	Household disposable income	Household expenditure	Self-assessed house prices
Percentile			
Less than 20	18	15	281
20—39.9	41	25	396
40—59.9	64	34	502
60—79.9	90	46	660
80—100	155	82	1 192

Notes: Households with reported disposable income or expenditure of 0 are dropped; household disposable income is for regular and recurrent income and is after tax; household expenditure includes imputed values for 2003—2005, detailed in Section I; persons in the household responsible for household bills are asked to complete the household-level expenditure questions; for households with more than one respondent to particular expenditure questions, the average response it taken; mean self-assessed house prices calculated for homeowners' primary residence

Sources: HILDA Release 10.0; authors' calculations

APPENDIX C: REGRESSION OUTPUT – HOUSING WEALTH EFFECTS

Table C1: Household-level Wealth Effects – Panels One and Two
(*continued next page*)

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
No of adults	0.117*** (12.64)	0.021*** (5.32)	0.095*** (7.25)	0.119*** (13.84)	0.125*** (10.58)	0.133*** (11.19)
No of children (aged 0–14)	0.078*** (9.75)	0.011*** (5.12)	0.080*** (6.46)	0.075*** (10.20)	0.079*** (7.97)	0.084*** (8.43)
Dummy: more than 2 adults	-0.003 (-0.20)	-0.048*** (-5.33)	-0.022 (-1.03)	-0.004 (-0.27)	-0.009 (-0.40)	-0.010 (-0.45)
M_{it}	0.015 (1.54)		0.053*** (3.54)	-0.010 (-1.14)	0.016 (1.01)	0.016 (1.02)
ln($HHDY_{it}$)	0.013*** (4.06)	0.009** (2.46)	0.016*** (3.83)	0.009*** (3.00)	0.025*** (5.32)	
ln(HP_{it}) × young	0.158*** (5.75)	0.098*** (2.95)	0.145*** (2.71)	0.141*** (5.55)		
ln(HP_{it}) × middle	0.061*** (2.76)	0.004 (0.17)	0.032 (0.92)	0.053*** (2.63)		
ln(HP_{it}) × old	-0.004 (-0.23)	-0.003 (-0.18)	0.016 (0.58)	-0.004 (-0.23)		
ln(HP_{it}) × young × renter					0.124*** (2.71)	0.124*** (2.70)
ln(HP_{it}) × young × owner					0.321*** (3.78)	0.319*** (3.74)
ln(HP_{it}) × middle × renter					0.012 (0.14)	0.014 (0.16)
ln(HP_{it}) × middle × owner					0.130* (1.77)	0.135* (1.83)
ln(HP_{it}) × old × renter					-0.293** (-2.04)	-0.291** (-2.02)
ln(HP_{it}) × old × owner					-0.004 (-0.05)	-0.007 (-0.09)

Table C1: Household-level Wealth Effects – Panels One and Two (*continued*)

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Constant	9.242***	-0.027***	9.379***	9.245***	8.488***	8.735***
	(53.33)	(-3.30)	(33.17)	(58.12)	(17.63)	(18.15)
Obs	34 177	24 594	20 675	34 177	20 065	20 065
Within R^2	0.023		0.015	0.025	0.033	0.031
R^2	0.723	0.006	0.777	0.736	0.745	0.744
RMSE	0.353	0.449	0.373	0.326	0.355	0.355

Notes: ***, ** and * indicate significance at the 1, 5 and 10 per cent level, respectively; *t* statistics in parentheses; dummies for year and labour force status omitted from table; robust standard errors clustered at the household level

Model 1 – panel one with self-assessed house prices

Model 2 – model 1 in first differences

Model 3 – model 1 (2006 to 2010)

Model 4 – model 1 with non-durable spending as the dependent variable

Model 5 – panel two with postcode-level independent house prices

Model 6 – model 5 excluding ln($HHDY_{it}$)

Sources: APM; HILDA Release 10.0; authors' calculations

APPENDIX D: REGRESSION OUTPUT – COHORT-LEVEL HOUSING WEALTH EFFECTS

Table D1: Cohort-level Wealth Effects

(*continued next page*)

	Model 1	Model 2
Dummy: more than 2 adults	-0.088* (-1.96)	-0.084* (-2.08)
No of adults	0.091*** (11.04)	0.177*** (11.51)
No of children (aged 0–14)	0.073*** (11.55)	0.089*** (11.11)
Dummy: cohort 1	-0.001 (-0.15)	0.006 (1.02)
Dummy: cohort 2	0.079*** (5.93)	0.035*** (3.69)
Dummy: cohort 3	0.062*** (4.66)	0.034** (2.75)
Dummy: cohort 4	0.115*** (8.71)	0.060*** (3.76)
Dummy: cohort 5	0.114*** (8.30)	0.071*** (4.35)
Dummy: cohort 6	0.097*** (7.34)	0.053*** (3.26)
Dummy: cohort 7	0.075*** (6.46)	0.032* (2.05)
Dummy: cohort 8	0.046* (2.16)	0.013 (0.49)
Dummy: cohort 9	0.053 (1.74)	0.003 (0.08)
Dummy: cohort 10	-0.024 (-0.66)	-0.041 (-1.06)
Dummy: cohort 11	-0.149*** (-4.03)	-0.144*** (-3.74)

Table D1: Cohort-level Wealth Effects

(*continued next page*)

	Model 1	Model 2
Dummy: cohort 12	-0.288*** (-7.33)	-0.271*** (-6.92)
Dummy: cohort 13	-0.453*** (-11.13)	-0.444*** (-11.22)
Education dummy: postgraduate	0.193*** (5.24)	0.195*** (6.32)
Education dummy: graduate	0.132*** (6.76)	0.150*** (9.59)
Education dummy: bachelor	0.143*** (6.98)	0.150*** (8.18)
Education dummy: diploma	0.115*** (5.02)	0.117*** (6.32)
Education dummy: certificate	0.053** (2.46)	0.054*** (3.06)
Education dummy: Year 12	0.042* (1.80)	0.044** (2.38)
M_{it}	0.037** (2.92)	0.032** (2.30)
$\ln(HHDY_{it})$	0.095*** (9.74)	0.077*** (12.72)
$\ln(HP_{it}) \times$ young	0.211*** (20.17)	
$\ln(HP_{it}) \times$ middle	0.220*** (7.01)	
$\ln(HP_{it}) \times$ old	0.220*** (4.78)	
$\ln(HP_{it}) \times$ young $\times$ renter		0.176*** (10.54)
$\ln(HP_{it}) \times$ young $\times$ owner		0.245*** (13.75)
$\ln(HP_{it}) \times$ middle $\times$ renter		0.155*** (5.59)

Table D1: Cohort-level Wealth Effects
(*continued*)

	Model 1	Model 2
$\ln(HP_{it})$ × middle × owner		0.233***
		(7.36)
$\ln(HP_{it})$ × old × renter		0.119**
		(2.55)
$\ln(HP_{it})$ × old × owner		0.208***
		(4.50)
Constant	6.122***	7.117***
	(8.72)	(9.77)
Obs	13 668	20 065
R^2	0.443	0.471
RMSE	0.442	0.450

Notes: ***, ** and * indicate significance at the 1, 5 and 10 per cent level, respectively; *t* statistics in parentheses; dummies for year, region, occupation, young, middle and labour force status omitted from table; robust standard errors clustered at the cohort level
Model 1 – panel one
Model 2 – panel two

Sources: APM; HILDA Release 10.0; authors' calculations

COPYRIGHT AND DISCLAIMER NOTICES

that the State of Victoria accepts no responsibility or liability whatsoever for any errors, faults, defects or omissions in the information supplied.

State of Queensland

© State of Queensland (Department of Natural Resources and Mines) 2012. In consideration of the State permitting use of this data you acknowledge and agree that the State gives no warranty in relation to the data (including accuracy, reliability, completeness, currency or suitability) and accepts no liability (including without limitation, liability in negligence) for any loss, damage or costs (including consequential damage) relating to any use of the data. Data must not be used for direct marketing or be used in breach of the privacy laws.

HILDA

The Household, Income and Labour Dynamics in Australia (HILDA) Survey was initiated and is funded by the Australian Government Department of Social Services (DSS), and is managed by the Melbourne Institute of Applied Economic and Social Research (Melbourne Institute). The findings and views based on these data should not be attributed to either DSS or the Melbourne Institute.

ACKNOWLEDGMENTS

The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Reserve Bank of Australia. We would like to thank Ellis Connolly, Frank Cowell, Alex Heath, Christopher Kent, Marion Kohler, Peter Tulip and three anonymous referees for useful comments on earlier drafts. We are also grateful to seminar participants at the Reserve Bank of Australia, the University of Tasmania, the Australian Conference of Economists 2013 and the HILDA Survey Research Conference 2013.

NOTES

1 Throughout this paper we use the term 'house' to refer to both detached houses and units/apartments.

2 Here, the estimated elasticity is converted to a MPC using the sample average ratio of non-housing consumption to house prices of 16 per cent.

3 Carroll *et al* (2003) show that the precautionary saving response does not exist for measures of wealth that exclude housing. This result appears counter-intuitive as housing is sometimes considered an illiquid asset. However, the emergence of products such as house loans that allow for redraw, as well as reverse mortgages, has made it easier for households to tap their housing wealth, thereby increasing its liquidity.

4 See, for example, Model 2 in Table C1.

5 In a working paper version of this paper (Windsor *et al* 2013), we used a more restrictive balanced-panel sample. The results are broadly similar to those presented here.

6 Arguably the best test of the credit constraints hypothesis would be to examine differences in housing wealth effects for homeowners who could raise \$3 000 in an emergency versus homeowners who could not raise \$3 000 in an emergency. Unfortunately, non-response rates for this particular question are quite high. It is also difficult to control for a household's selection into subjective self-assessed categories. For these reasons, even the positive correlation between credit constraints and age reported in the text should be treated with some caution.

7 In these comparisons, the ABS figures are not adjusted to make them more comparable to the HILDA Survey measure. However, differences in the concept and scope of these data should be borne in mind. The ABS

data constitute the broadest, accruals-based measure, while the survey data measure only regular and recurring spending. Aside from these differences, notable omissions from the HILDA spending data include: entertainment expenses, non-fee education expenses, gifts and donations, personal and household services, health and beauty products, ornaments, art and jewellery, and financial service charges.

8 Note that although we use the term 'house price' throughout this paper, respondents to the HIDLA survey in fact provide a self-assessed house value, which will incorporate any capital improvements made to the dwelling. Ideally we would like to use house prices, rather than values, thereby abstracting from variation caused by capital improvements. We can go some way towards constructing such a variable by deducting from house values the expenditure component 'home repairs/renovations and maintenance'. This expenditure component is only available from 2006 to 2010, however, and we cannot isolate expenditure solely on capital improvements (one has to deduct the aggregate component 'home repairs/renovations and maintenance'). Nonetheless, the results using this definition of house prices are very similar to those reported.

9 Education is generally considered to be an effective proxy for permanent income. Attanasio and Weber (2010), for instance, document that more highly educated households tend to have higher (and steeper) income profiles than those headed by less educated individuals.

10 In reality, the extent of this omitted variable bias appears limited. Replacing the dependent variable, ln(*total spending*$_{it}$) with ln(*HHDY*$_{it}$), and dropping ln(*HHDY*$_{it}$) from the list of explanatory variables, yields insignificant coefficients on the house price terms across the age distribution.

11 Attanasio *et al* (2009) note that their results are similarly robust to the inclusion of current income.

12 These ratios are 0.2, 0.18 and 0.13 for young, middle and old age groups respectively.

13 Model 2 in Appendix C1 shows the results, using self-assessed house prices, from estimating Equation (3) in first differences, similar to the approach of Campbell and Cocco (2007). As expected, these results are similar to the fixed-effect results, allaying concerns one might have about violating the strict exogenity assumption (i.e. correlation between our controls and the residual in Equation (3)).

14 The following items are classified as durable: new and used motor vehicles, motorbikes or other vehicles; computers and related devices; televisions, house entertainment systems and other audio-visual equipment; whitegoods such as ovens and fridges; and furniture.

15 We also estimated the model using hedonically-adjusted prices at the postcodes level in order to control for possible biases due to changes in the quality and composition of houses sold in a given postcode and time period. The results were very similar.

16 At the suggestion of a referee we also tested our interpretation of the age-tenure distribution of housing wealth effects in a two-stage regression exercise. In the first stage, we replaced *ln*(*total spending*$_{it}$) in Equation (3) with real household disposable income in period *t+1*. If house prices are associated with income expectations, then the expectation was for a positive coefficient on house prices in this regression. In the second stage, *ln*(*total spending*$_{it}$) was regressed on fitted values of real household disposable income in period *t+1* obtained from the first-stage regression, interacted with age-tenure dummies and using household-fixed effects. If income expectations affect the spending of young households the most, the expectation was for a positive coefficient on fitted values of income in period *t+1* for young households (with an appropriate transformation of the standard errors in the second-stage regression). In the first stage, we found some evidence that current house prices are positively associated future income. In the second-stage, fitted values of future income were positively associated with the current spending of young owners, young renters and middle-aged owners. These results are broadly consistent with our interpretation of house prices for young renters capturing a common association between house prices and a third factor, most likely income expectations. A related third common factor that could jointly affect house prices and spending is the local unemployment rate, which may be interpreted as a proxy for income insecurity. Including the local unemployment rate in our model, interacted with age-tenure dummies, suggests that higher local unemployment is indeed associated with lower consumption for young homeowners and young renters. The inclusion of the local unemployment (level or change) rate does not, however, substantially change our results, with the age-tenure pattern of house price effects largely unchanged from those presented in Table 6. See also Benito (2006) for similar evidence with micro data on British households.

17 The same rationale is used by Attanasio *et al* (2009) to justify using the level of regional house prices in their analysis rather than homeowners' estimates of the price of their homes, which are available in the UK Family Expenditure Survey data that they use.

Publication II

First published in Finlay, R. and Price, F. (2015). Household saving in Australia. *The B.E. Journal of Macroeconomics*, **15**, 677-704.

Household Saving in Australia*

Richard Finlay

Reserve Bank of Australia, 65 Martin Place, NSW 2000, Australia *and*
School of Economics, University of New South Wales, NSW 2052, Australia
finlayr@rba.gov.au

Fiona Price

Reserve Bank of Australia, 65 Martin Place, NSW 2000, Australia
pricef@rba.gov.au

Abstract

This paper investigates household saving behaviour in Australia, as well as the drivers behind the recent rise in the aggregate household saving ratio. Our results explaining differences in saving behaviour across households are consistent with theory and previous findings. As might be expected, households' saving ratios tend to increase with income, but decrease with wealth and gearing. More at-risk households such as single-parent and migrant households tend to save more than other households, all else equal. While saving differs substantially across age groups we find that, at least in part, this reflects differing circumstances.

Our results suggest that the rise in household saving between 2003/04 and 2009/10 was driven by changes in behaviour rather than changes in population characteristics: in particular, more educated households, as well as households with high debt and/or wealth increased their propensity to save. Our interpretation of these results is that a reduction in future income growth expectations for more highly educated households after the financial crisis, and an associated effort to rebuild wealth and repay debt, drove to the aggregate rise in household saving.

JEL Classification Numbers: D14, E21

Keywords: household saving, micro data

* The authors would like to thank Tiago Cavalcanti, Alex Heath, Jonathan Kearns, Christopher Kent, Gianni La Cava, Peter Tulip and an anonymous referee for helpful comments. 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 Australia.

1. Introduction

Between the early 1970s and the early 2000s the aggregate household saving ratio in Australia steadily declined, from around 20 per cent to around zero (Figure 1).[1] This trend was likely driven by a number of factors, including an increased availability of credit, falling real interest rates, more stable economic outcomes, rising asset prices, and rising household income and income expectations. While the importance of various factors waxed and waned over the three decades, it is likely that all contributed to some extent to a higher rate of growth in consumption compared with income, and so the fall in the saving ratio seen over this period.

Figure 1: Household Income, Consumption and Saving

Per capita annual rolling average

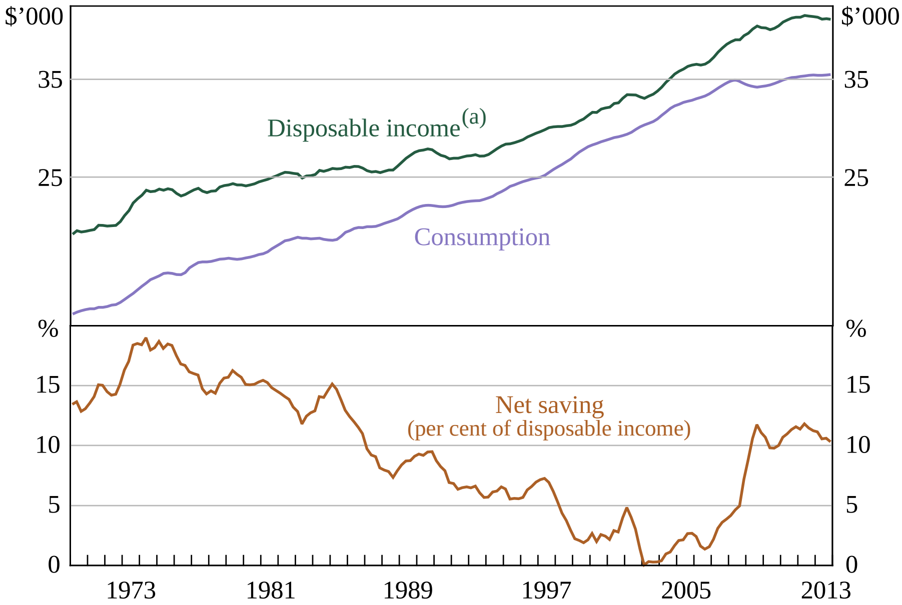

Notes: In 2011/12 dollars; deflated using the household final consumption expenditure implicit price deflator; disposable income and consumption are plotted using a log scale

(a) After tax and net interest payments

Sources: ABS; RBA

However, in the latter half of the 2000s, the household saving ratio reversed this decline, and is now at a level similar to that of the mid 1980s. This rise in saving coincided with a deterioration in Australia's (and the world's) economic environment, with GDP growth, credit growth and the exchange rate falling sharply over late 2007 and 2008, while official interest rates were cut sharply after a period of steady increase (Figure 2). As well as suggesting more caution on the part of borrowers given a riskier economic environment, the fall in credit growth could also be indicative of more restrictive credit availability,

---

1 The *gross* saving ratio is defined as disposable income (income less tax and interest payments) minus consumption, divided by disposable income (this is the saving ratio we calculate from micro data and base our regressions on; see Section 2.1 for more details). The *net* saving ratio shown in the bottom panel of Figure 1 additionally deducts depreciation.

which may have led to higher saving.[2] More generally, the rise in saving represents an important change in the economic environment given that household consumption accounts for a little over half of GDP.

Figure 2: Australian Economic Indicators

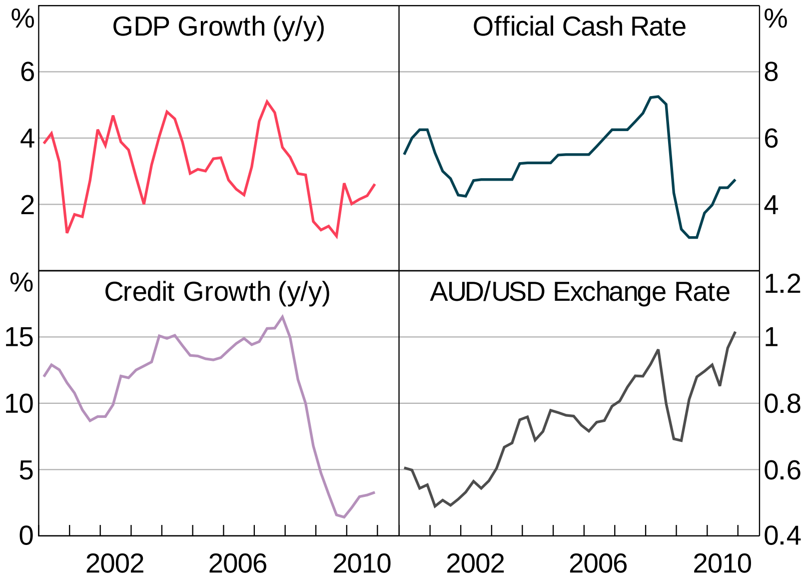

Sources: ABS; RBA

The extent to which the higher saving ratio is sustained will depend on what caused the change in saving. For example, if saving rose due to an unexpected boost to income that households believed to be temporary, standard theory would suggest that saving will fall again as the boost to income dissipates. Conversely, if household saving initially fell due to expectations of high future income and asset price growth, as well as an associated run-up in housing debt, then a downward reassessment of those expectations may lead to a more enduring rise in saving.[3]

In this paper we investigate both the static determinants of household saving and the drivers of the rise in saving over the 2000s. By examining aggregate data on household income and consumption, we can observe that saving rose around the mid to late 2000s, reflecting an increase and then levelling off in average per capita income, and a temporary fall and then levelling-off in consumption (Figure 1). Even if one assumes that changes in income were due to factors outside of the household sector's control, it is difficult to draw conclusions from the aggregate data about what drove the changes in consumption behaviour. Instead we turn to household-level data and examine the link between various household

---

2 Central bank commentary at the time pointed to both supply and demand factors as driving the decline in credit growth, with lending standards tightened for marginal borrowers but credit remaining readily available for higher-quality borrowers (RBA 2008). Finlay and Jääskelä (2014) also find that both credit supply and non-credit-supply shocks played a role in the fall in credit growth over the late 2000s, with credit supply shocks explaining around one-third of the total fall in credit growth.

3 See, for example, Dynan and Kohn (2007) for a discussion of the link between house prices, borrowing and saving in the United States, or Iacoviello (2004) for a general equilibrium model of house prices, debt and consumption.

characteristics and saving behaviour. To do this, we use the Household Expenditure Survey (HES) from the Australian Bureau of Statistics (ABS) detailing household income and expenditure in 2003/04 and 2009/10.

We are agnostic about the drivers of saving in the cross-section, and consider a number of factors that may drive the saving decision including life-cycle factors, credit constraints, precautionary motives, income and wealth. Regarding the rise in saving seen over the 2000s, our hypothesis is that lower income growth expectations following the financial crisis, as well as an associated downward revision to asset price growth (primarily house price growth) expectations, are the primary drivers, although we do not preclude other factors in our modelling.

It is important to note, however, that we cannot directly test the effect of, for example, income growth expectations on saving behaviour, since there is no variable that exactly measures a household's expectations. Rather, we examine how saving relates to household characteristics that are correlated with our driver of interest, for example education level. To the extent that saving varies with education, we draw the inference that it is underlying income growth expectations that are driving the behaviour. We acknowledge that while the effect of household characteristics on saving can be estimated, the interpretation that we place on these estimates is subject to debate.

While ours is the first study looking at the recent rise in household saving in Australia using household-level data, other papers have analysed household saving behaviour. For Australia, Harris, Loundes and Webster (2002) use household-level data from Melbourne Institute surveys to consider the household characteristics that lead a household to identify with a type of saving behaviour that ranges from 'running into debt' to 'saving a lot'. The authors find that households with higher income and wealth, households that own their own home and households with a more positive economic outlook tend to identify themselves as active savers. Their findings suggest several saving hypotheses help to explain variation in household saving behaviour.

More recently, Berger-Thomson, Chung and McKibbin (2009) use the Household, Income and Labour Dynamics in Australia (HILDA) Survey to examine how uncertainty affects households' consumption decisions. The authors find that households that are worried about their future employment status have lower marginal propensities to consume out of current income compared with households that are not concerned about their future employment status, and so save more.

Chamon and Prasad (2010) examine household saving behaviour in China using household-level data between 1995 and 2005. Similar to our study, one of their aims is to uncover the reasons behind the rise in the Chinese household saving ratio over this period. The authors find that precautionary saving motives are an important determinant of this rise, with younger and older households increasing saving due to rising uncertainty and increasing housing, education and healthcare costs in China.

Attanasio and Weber (1994) examine two popular hypotheses for the sharp fall that occurred between 1986 and 1988 in the United Kingdom's household saving ratio: that it was due to a substantial rise in house prices; and that it was due to a rise in perceived permanent income. While wealth effects may have boosted consumption growth in the 1980s, the authors conclude that the sharp fall in saving is best explained by younger households upwardly revising their expectations of permanent income.

The remainder of this paper is organised as follows. In Section 2 we describe the household-level datasets that we use, and examine how they compare with aggregate data available in the Australian national accounts. Section 3 presents cross-sectional results on the drivers of savings behaviour from a model of the median household's saving that is similar to those employed in Chamon and Prasad (2010) and Islam *et al* (2013). Section 4 presents results on the rise in saving seen over the 2000s from the median regression model, as well as a decomposition of the change in the mean saving ratio into parameter and characteristic effects. Modelling median saving allows us to assess determinants of the saving behaviour of a 'typical' household, while modelling mean saving allows us to quantify the size of various influences on the aggregate saving ratio. Section 5 concludes.

2. Data

The 2003/04 and 2009/10 Household Expenditure Surveys are cross-sectional surveys of a nationally representative sample of households in Australia during the survey period.[4] For each household, the surveys collect information on income and consumption, as well as a range of socio-demographic characteristics. These socio-demographic characteristics allow us to assess the saving behaviour of particular groups of households, which is not possible with aggregate data.

The ABS also conducted expenditure surveys in 1975/76, 1984, 1988/89, 1993/94 and 1998/99. We do not use these earlier surveys in our analysis since: (i) methodological changes render surveys conducted before 1998/99 less comparable to those from 1998/99 on; and (ii) the surveys conducted before 2003/04 omit important variables, such as household wealth, which can play a large role in influencing saving behaviour.

---

4 The 2003/04 HES surveyed around 7 000 households, while the 2009/10 HES surveyed around 10 000 households. The sample we use excludes those who give zero or negative values for income, the unemployed, and households where the household head is aged over 75 years. We also trim the top and bottom 2 per cent of the sample based on the saving ratio distribution.

2.1 Definition of Income, Consumption and Saving

The most important quantitative data that we use are household income, consumption and saving.

Disposable income includes: labour income; farm income; income of unincorporated enterprises; net rental income; imputed rent for owner-occupiers[5]; interest on savings; dividends; transfer income from the government, private institutions and other households; superannuation contributions by employers on behalf of employees[6]; superannuation drawdowns by self-funded retirees; inheritance; gifts and other income from family members. Income is after tax and interest payments.

Note that the national accounts definition of income includes a number of items that are unavailable in the HES, the largest of which are imputed interest and current transfers to non-profit institutions serving households. We also cannot separately identify (and therefore exclude) capital draw-downs from investment earnings for self-funded retirees in the HES, so that income for self-funded retirees is overstated.

Consumption includes total expenditure on goods and services as well as imputed rent for owner-occupiers. Principal and interest repayments on debt, home capital improvement expenditure and life insurance and superannuation related expenses are not included in consumption.

Saving is calculated as the difference between disposable income and consumption. The main difference between our definition of saving and that from the national accounts stems from the different definition of income, as noted above. Note that our definition of saving captures only active saving and does not include any capital gains or losses.

2.2 Comparison of Aggregate and Micro Data

In order to use the household surveys to analyse the drivers behind the increase in the aggregate saving ratio, the surveys must be comparable with each other and with data from the national accounts. There were no major methodological changes between the 2003/04 and 2009/10 Surveys, so the two surveys should be comparable, and while the surveys do not capture all household consumption or income when compared with national accounts data, they capture a similar proportion of each. This implies that the aggregate saving ratio from the HES datasets should be consistent with the aggregate saving ratio from the national accounts, as indeed it is, with both measures showing a similar increase in saving between 2003/04 and 2009/10.

---

5 Imputed rent for owner-occupiers is determined using the methodology outlined in ABS (2008) for 2003/04; imputed rent using this methodology is already included in the 2009/10 HES.

6 Superannuation is the name given to defined contribution pension accounts in Australia, into which employers must contribute 9.25 per cent of employees' wages and salaries; employees can also contribute additional funds if they wish to.

2.3 Descriptive Analysis

Looking at the distribution of the saving ratio across households, the median saving ratio in 2003/04 was 5 per cent, while in 2009/10 it was 9 per cent (Figure 3). The shift up in the saving ratio evident across most of the distribution is consistent with the rise in the aggregate saving ratio over this period. The distribution of the saving ratio displays a long tail of negative saving ratios. This is unsurprising as consumption is always positive, but income, which is the denominator of the saving ratio, can sometimes be close to zero, which leads to large negative saving ratios for some households.

Figure 4 shows how income, consumption and saving vary by household age in the 2003/04 and 2009/10 Surveys. Household consumption tends to track income closely, with both varying significantly over the life cycle, suggesting that households do not fully smooth their consumption, although Attanasio (1999) points out that the hump-shaped consumption profile is less pronounced after controlling for family size and composition. Between the 2003/04 and 2009/10 Surveys, saving increased especially for younger and older households, with income rising more than consumption for these groups.

Changes in household saving behaviour do not appear to be specific to certain levels of household wealth, with the saving ratio increasing across all wealth quintiles between 2003/04 and 2009/10 (Figure 5). Most (age-matched) income quintiles also saw a rise in saving between 2003/04 and 2009/10, with only the lowest income group recording a fall in saving (Figure 6).[7]

This simple descriptive analysis suggests that relatively young and old households, but not middle-aged households, considerably increased their saving between 2003/04 and 2009/10, while a change in saving behaviour was evident across most wealth and income groups. While this could in part reflect the ageing of the population, we find that this ageing effect is not large enough to explain the large increase in the saving ratio over the 2000s. (Appendix A describes a simple decomposition model based on age and birth cohorts, and shows that the increase in saving cannot be attributed to these factors).[8] Given this, we need to consider other explanations for the rise in the saving ratio.

7 Age-matching controls for age-related effects when comparing income quintiles. For example, since post-retirement households are typically in the lower income quintiles, the saving behaviour of older households will have a significant influence on the saving behaviour of the lower (non age-matched) income quintiles. Age-matching is done by splitting the households in each age group into separate income quintiles. Income quintiles from each age group are then recombined, so that, for example, the lowest age-matched income quintile consists of all those households that make up the lowest income quintile within each age group.

8 This is not surprising given that most of the rise in the saving ratio occurred over a relatively short period, whereas the ageing of the population is a slow-moving process. Chamon and Prasad (2010) find a similar result in their study, while Browning and Lusardi (1996) argue that ageing is too slow to provide a sufficient explanation for the large decline in the US aggregate household saving ratio.

Figure 3: Distribution of Saving Ratio

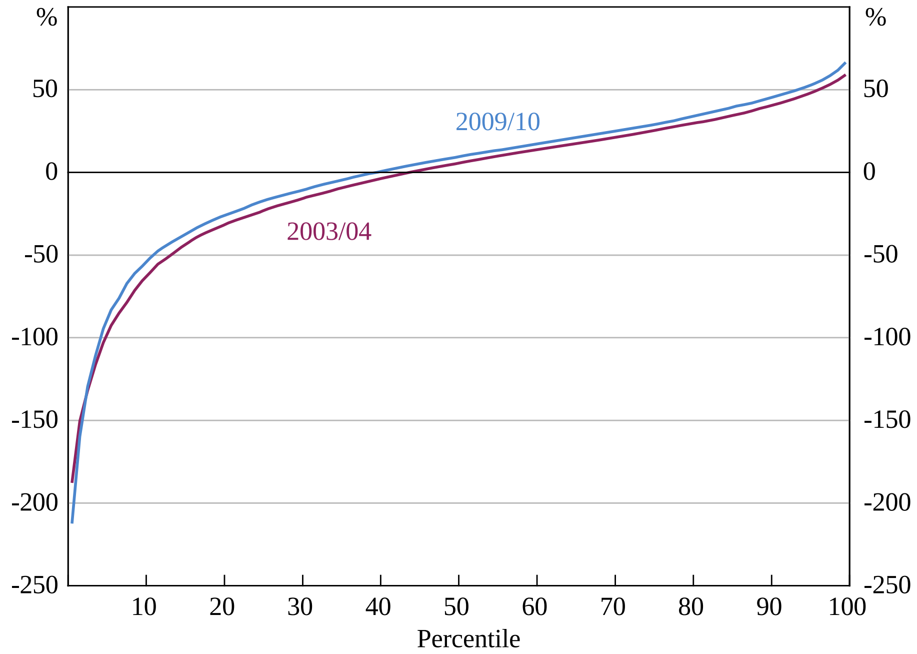

Sources: ABS; authors' calculations

Figure 4: Household Income, Consumption and Saving

By age of household head

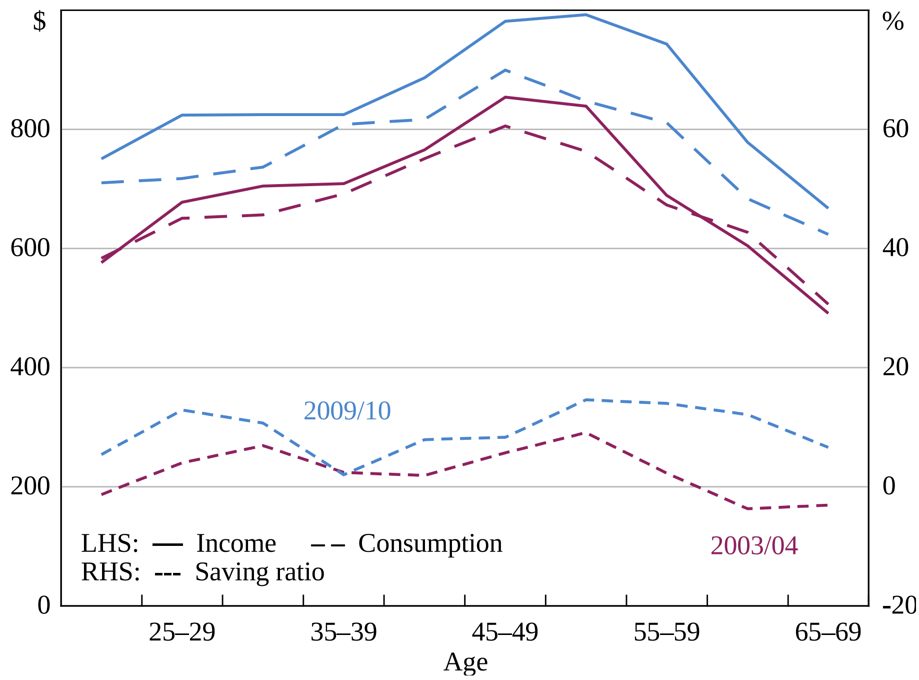

Notes: Saving is gross of depreciation; income and consumption are weekly and in 2009/10 dollars; weighted averages across age groups

Sources: ABS; authors' calculations

Figure 5: Household Saving Ratio

By wealth quintile

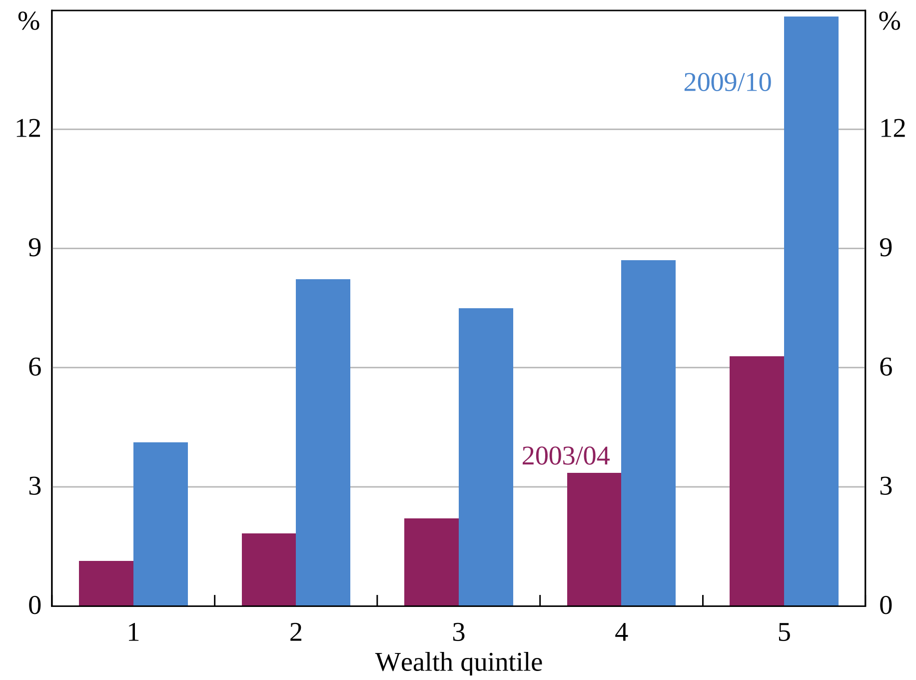

Notes: Gross of depreciation; weighted average

Sources: ABS; authors' calculations

Figure 6: Household Saving Ratio

By income quintile

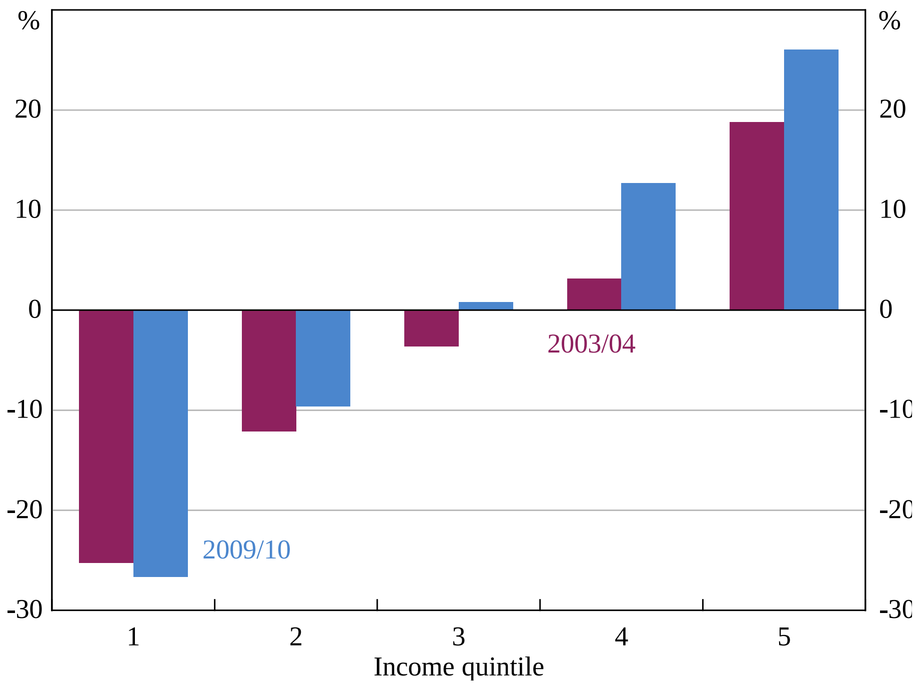

Notes: Gross of depreciation; weighted average; age-matched

Sources: ABS; authors' calculations

3. Cross-sectional Analysis – Determinants of Saving

In this section we investigate the drivers of saving in the cross-section, rather than the increase in saving seen over the 2000s. To do this we estimate a model of the median saving ratio that takes into account a range of household characteristics. The median saving ratio gives a better indication of how much a 'typical' household saves than the mean saving ratio, which can be heavily influenced by a small number of extreme values. The mean saving ratio is nonetheless important since it determines economy-wide household saving, and we return to it in Section 4.

3.1 Determinants of Saving

Income is a particularly important determinant of household consumption, although there is some debate as to how it effects saving. Economic orthodoxy would suggest that a household's permanent or long-run level of income should *not* affect the saving ratio, since households with relatively high levels of permanent income would also have relatively high levels of consumption. Aggregate time series data on national saving supports this proposition: as countries grow richer, household incomes trend higher but saving ratios do not. Conversely, the evidence from cross-sectional, household-level studies is less clear; for example, Dynan *et al* (2004) find that individual households' saving ratios *are* affected by their level of permanent income.

Our main results are estimated under the assumption that households' permanent income levels do not affect saving ratios, although our results are robust to relaxing this assumption (see Model (2) in Table B3). In particular, we assume that a household's saving ratio is a function of the deviation of their current level of income from their permanent level of income:

$$saving\ ratio_i = \beta_1(y_i - y_i^*) + \beta_2 X_i + \varepsilon_i.$$

Here y_i is the natural logarithm of household *i*'s current income, y_i^* is the logarithm of their permanent income, and X_i represents other household characteristics pertinent to the saving decision such as age, labour force status and household composition. This model implies that a household will increase their saving ratio if their current level of income rises and/or their permanent level of income falls, for example due to a one-off bequest or a downward reassessment of expected future income growth.

In practice we cannot observe the permanent income of a household, and so must estimate it. We do this by regressing current income on proxies for permanent income, including households' education level, occupation and age, and taking the fitted values as measuring permanent income. We estimate separate models for labour income and non-labour income (see Table B1 for model results). We then use the percentage deviation of current income from the modelled estimate of permanent income as our income variable, $\left(y_i - y_i^*\right)$.

As our permanent income regressions are cross-sectional, they will capture the income level that, for example, a highly educated household 'should' be earning given earnings of similar households at the present time. They will *not* capture any time-series dimension, such as economy-wide expectations of future income growth for highly educated households, however, which should also form a part of permanent income. Given this, we treat the separate effect of education on a household's saving ratio, after controlling for deviations of current from modelled permanent income, as a measure of future income growth expectations (Attanasio and Weber (2010), for instance, document that more educated households have steeper income profiles than those headed by less-educated individuals).

Some authors have argued that including a measure of income in models such as ours may introduce measurement error and endogeneity issues, resulting in biased estimates. For example, Sabelhaus and Groen (2000), Brzozowski and Crossley (2011) and Meyer and Sullivan (2011) argue that large dissaving at the bottom of the income distribution in household surveys is more likely to be due to households under-reporting their income than genuine dissaving, although Browning and Lusardi (1996) argue that reporting bias in household income is unlikely to be a serious issue for most households. There is growing recognition, however, that income is too important as a driver of household saving to be excluded – see, for example, Dynan *et al* (2004) and Muellbauer (2007) – and so we choose to include it in the form discussed above in our main model. As a robustness check we also estimate a model excluding any measure of income – Model (2) in Table B3 – and find similar results.

In addition to income, the other drivers of saving that we explore are outlined below.

- **Credit constraints**. Credit constrained households would be expected to save more than otherwise similar households. This follows since credit constrained households that wished to borrow to fund consumption would be precluded from doing so, in effect forcing them to save more than they wish. Credit constrained households are identified from households' answers to questions regarding financial stress; households are assumed to be credit constrained if they answer in the affirmative to at least two out of seven financial stress questions. The reference household is not credit constrained.

- **Precautionary motives**. Similar to Chamon and Prasad (2010), we construct a variable that seeks to measure a household’s risk of unemployment, a risk that is likely to influence a household’s saving behaviour. (Chamon and Prasad, in their study of Chinese households, estimate the risk of incurring a large health expense). One might expect that employed households that face a relatively high risk of becoming unemployed in the future will save more than other households (see, for example, the models outlined in Zeldes (1989), Deaton (1991), Carroll (1992) and Carroll and Samwick (1997)). Each household’s risk of unemployment is estimated using a logit model of the probability of a household containing one or more unemployed people. If a household’s fitted probability of future unemployment is greater than 10 per cent, the risk of unemployment variable is set equal to 1 (see Table B2 for model details). Precautionary motives may also be captured in other variables that describe households with less secure incomes or those who are more vulnerable to income shocks, such as migrant and single-parent households.

- **Wealth effects**. Higher wealth has been found to have a significantly positive effect on household consumption in Australia, and therefore a negative effect on saving, all else equal (Dvornak and Kohler 2003; Yates and Whelan 2009; Windsor, Jääskelä and Finlay 2015). We include the ratio of household wealth relative to income and the gearing ratio (debt relative to assets) to capture wealth effects in our model, as well as a dummy variable indicating whether a household is living off retirement savings or obtains more than 20 per cent of their income from investments.

- **Life-cycle motives**. Although an ageing population cannot explain the rise in the aggregate saving ratio, age is an important determinant of household saving in the cross-section. Age dummy variables are used to capture the saving behaviour of different age groups: young (less than 30 years old), pre-retirement (50–64 years) and old (65 years and over). The reference household is middle-aged (30–49 years).

Other controls include household size; the number of children in the household (relative to household size); state or territory of usual residence; region of residence (rural/urban); skill level of occupation; marital status; gender of the household head; dummy variables for owning one’s home outright or with a mortgage; and dummy variables that identify if a household obtains more than 20 per cent of their income from wages and salaries, business income, government payments, or other income.

3.2 Cross-Sectional Results

Table 1 shows selected results from the median regressions for 2003/04 and 2009/10, where the dependent variable is the saving ratio and the independent variables are as described above. The differences in coefficients across the two time periods are also presented, and will be discussed in Section 4 (full regression outputs are presented in Table B3).

Table 1: Median Regression Model

	Coefficients		
Variable	2003/04	2009/10	Difference over time
Income elasticity	0.04***	0.04***	0.0
Highly educated	–4.7***	1.9	6.6***
Single-parent household	8.0**	5.2**	–2.8
Migrant	2.4	2.8*	0.4
Self-funded retiree	–19.0***	–8.3**	10.8*
Wealth-to-income ratio	–0.1	–0.2***	–0.1
Gearing ratio	–13.8***	–2.4	11.4**
Young	–0.8	–0.3	0.6
Pre-retired	0.3	6.8***	6.5**
Old	10.1**	7.9*	–2.2

Notes: ***, ** and * represent significance at the 1, 5 and 10 per cent level, respectively; HES household weights used; 2003/04 and 2009/10 samples are pooled and 300 bootstrapped repetitions of the regression are used to obtain the standard errors; coefficients on other variables are reported in Table B3

Sources: ABS; authors' calculations

Income

As expected we find that the coefficients on the deviation of current income from permanent income are significant and positive, meaning that households whose current level of income is above their permanent level of income save more, all else equal. The value of the coefficient on income suggests that in the cross-section, a 1 percentage point increase in current income relative to permanent income is associated with a 0.04 percentage point increase in the saving ratio, all else equal; this is within the range of estimates presented in Dynan *et al* (2004) using US data, although a little lower than those presented in Chamon and Prasad (2010) using Chinese data.

Precautionary motives

At-risk households such as single-parent households tend to save more than other households, all else equal. Households where the household head was not born in an English-speaking country also tend to save more. This is consistent with the results of Islam *et al* (2013), who examine the saving behaviour of migrants in Australia and find that they have a higher propensity to save compared with Australian-born households with similar characteristics. While this effect could reflect the differing priorities of migrants compared with existing residents, it could also be evidence of precautionary saving if being born in a non-English-speaking country is associated with less certainty regarding employment (and indeed being a migrant is associated with a higher risk of unemployment, all else equal – see Table B2 in Appendix B).

Wealth

We find that higher wealth-to-income ratios and higher gearing ratios are associated with lower saving ratios (and therefore more consumption), holding all else equal. Similarly, we find that those either living off or deriving a substantial part of their income from investments save less than otherwise similar households would.[9]

Life-cycle

Perhaps unsurprisingly we find that holding all else equal, pre-retirement households save more than middle-aged households (the reference group), who in turn tend to save the same or more than the young. Older households, all else equal, tend to save more than middle-aged or younger households would, were they to face similar living circumstances, suggesting that the low level of saving seen in the data by older households is predominantly due to their circumstances rather than their age *per se*.

Other factors discussed in Section 3.1 such as education level, credit constrains and the effect of being at greater risk of unemployment were either not statistically significant or changed sign over the two sample periods.

4. Time Series Analysis – The Rise in Saving

This section examines the rise in the saving ratio between 2003/04 and 2009/10. To do this, we look at changes in households' propensity to save using the median regression model from Section 3; we also decompose the total change in the mean saving ratio (the concept of saving reported in the national accounts) into changes in households' propensity to save and changes in household characteristics.

9 Note that self-funded retirees are likely to dissave more than suggested by our results. As discussed in Section 2.1, in our dataset we cannot separately identify capital draw-downs from investment earnings for self-funded retirees. As such, some of the income attributed to self-funded retirees is actually dissaving from their accumulated assets.

4.1 Changes in the Median Saving Ratio

The last column of Table 1 in Section 3 shows the change in model coefficients between the 2003/04 Survey and the 2009/10 Survey. We interpret changes in these coefficients, where they are statistically significant, as indicating changing preferences regarding saving for those households with the corresponding characteristics. As noted in Section 1, however, since we cannot directly measure household preferences, other interpretations of the data are possible.

Income

There is no change in the coefficient on deviations of current relative to permanent income between the two surveys. There is a significant change in the coefficient on education, however. Relative to high school educated households, more educated households significantly increased their propensity to save between 2003/04 and 2009/10. If we interpret education as a measure of future income growth expectations, this suggests that more educated households downgraded their income growth expectations between 2003/04 and 2009/10; this implies current income being high relative to permanent income, which would lead households to spend less and save more.

Wealth

The negative effect on saving of a high gearing ratio fell significantly between 2003/04 and 2009/10. This suggests that households adopted a more prudent attitude to debt between 2003/04 and 2009/10, and accords with other data sources that suggest households have increased their voluntary mortgage repayments over the past few years, aided by lower interest rates. We also note that households in the 'pre-retirement' age group (50-64 years), self-funded retirees and those earning at least 20 per cent of their income from investments – that is, those households most exposed to movements in asset prices – increased their propensity to save between 2003/04 and 2009/10, suggesting a reaction to the large fall in asset prices that occurred during the financial crisis.

For other factors, the change in propensity to save between 2003/04 and 2009/10 was not statistically significant.

4.2 Changes in the Mean Saving Ratio

Using the same model, but applied to the mean, the model-implied mean saving ratio in year *i* can be expressed as

$$\overline{saving\ ratio}_i = \bar{X}_i{}'\hat{\beta}_i$$

where $\bar{X}_i$ is a vector of the averages of variables used in the saving model in year *i*, including the constant term, and β_i is a vector of the coefficient terms associated with the variables in $\bar{X}_i$ in year *i*. Given this, we can express the change in the mean saving ratio as

$$\begin{aligned}\Delta saving\ ratio_{21} &= \bar{X}_2{}'\hat{\beta}_2 - \bar{X}_1{}'\hat{\beta}_1 \\ &= \left(\bar{X}_1{}'\hat{\beta}_2 - \bar{X}_1{}'\hat{\beta}_1\right) + \left(\bar{X}_2{}'\hat{\beta}_2 - \bar{X}_1{}'\hat{\beta}_2\right) \\ &= \bar{X}_1{}'\left(\hat{\beta}_2 - \hat{\beta}_1\right) + \left(\bar{X}_2{}' - \bar{X}_1{}'\right)\hat{\beta}_2 \\ &= parameter\ effect + characteristic\ effect,\end{aligned}$$

where year 1 represents 2003/04 and year 2 represents 2009/10. That is, the change in the model-implied mean saving ratio can be decomposed into changes in model parameters and changes in population characteristics. This follows the method introduced by Blinder (1973) and Oaxaca (1973).[10]

This decomposition enables us to separately estimate the roles that population characteristics and model parameters have played in the rise of household saving. The results suggest that changes in population characteristics played very little role in the increase in the saving ratio between 2003/04 and 2009/10, with changes in model parameters – and in particular those related to education and wealth – dominating (Table 2).

Table 2: Contribution to Change in Mean Saving Ratio between 2003/04 and 2009/10

	Total Change	of which		
		Education	Wealth	Other
Due to Model Parameters	5.8***	3.6***	5.0**	–2.8
Due to Characteristics	0.1	0.0	–0.1	0.2
Total	5.9***	3.6***	4.9**	–2.6

Notes: ***, ** and * represent significance at the 1, 5 and 10 per cent level, respectively; HES household weights used; 2003/04 and 2009/10 samples are pooled and 300 bootstrapped repetitions of the regression are used to obtain the standard errors; the wealth category includes the wealth-to-income ratio, the gearing ratio, home ownership dummies, the self-funded retiree dummy and the pre-retirement age dummy

Sources: ABS; authors' calculations

Consistent with the results from the median analysis in Section 4.1, more educated households increased their propensity to save in an economically and statistically significant way between 2003/04 and 2009/10. Given our interpretation of education as a measure of future income growth expectations, the rise in saving for more educated households suggests a downward reassessment by these households of their future income prospects. Wealthy households and those with high debt levels

10 As noted, the model used here is very similar to the model for the median saving ratio in Section 3, except that it is estimated by least squares. As such, the model is of conditional mean saving ratios rather than conditional median saving ratios. See Table B4 for output from the least squares regression.

(included in the wealth grouping) also tended to increase their propensity to save between 2003/04 and 2009/10, suggesting an effort to rebuild wealth after the effects of the financial crisis and changed attitudes to debt.

Overall, the results from the median and mean time series analysis are consistent with the rise in saving seen over the 2000s being driven by a downward reassessment of future income growth and asset price growth expectations following the financial crisis, with households adopting a more prudent attitude towards debt over this period.

5. Conclusion

This paper investigates household saving behaviour in Australia, as well as the drivers behind the recent rise in the aggregate household saving ratio. Our results explaining household saving behaviour in the cross-section are consistent with theory and previous findings. As might be expected, households' saving ratios tend to increase with income, while saving is found to decrease with wealth and gearing. Single parent and migrant households tend to save more than other households, all else equal. While saving differs substantially across age groups we find that, at least in part, this reflects differing circumstances.

Our results suggest that the rise in household saving seen over the 2000s was driven by changing behaviour rather than changing population characteristics. In particular, more highly educated households as well as households with high debt and/or wealth increased their propensity to save between 2003/04 and 2009/10. Our interpretation of these results is that a more prudent attitude towards debt and an effort to rebuild wealth after the financial crisis contributed to the rise in the household saving ratio; the sharp increase in saving for higher educated households also suggests a reduction in future income growth expectations.

Appendix A: A Simple Model of Age, Cohort and Time Effects

This model follows the approach of Deaton and Paxson (1994) and Chamon and Prasad (2010), and provides a simple way to disentangle age and birth cohort effects to find their 'pure' effect on saving.[11]

With no shocks to income and a constant real interest rate, the life-cycle hypothesis suggests household consumption can be expressed as:

$$C_{ab;h} = f\left(a\right) \times W_b \times e^{\varepsilon_{ab;h}}.$$

Here $C_{ab;h}$ denotes consumption for household *h* where the household head is aged *a* and belongs to birth cohort *b*, $f(a)$ describes how consumption varies with age, W_b denotes the average lifetime resources of households from birth cohort *b*, and $e^{\varepsilon_{ab;h}}$ is a (multiplicative) household-specific idiosyncratic shock.

Taking logs and averaging consumption over households in the same age (*a*) and birth cohort (*b*) gives

$$\overline{\ln\left(C_{ab}\right)} = \overline{\ln f\left(a\right)} + \overline{\ln W_b},$$

where the age effect – $f(a)$ – is assumed to depend on age but not birth cohort, while lifetime resources – W_b – are assumed to depend on birth cohort but not age. We then use dummy variables to decompose the age, birth cohort and time (i.e. unexplained) components of consumption

$$\overline{\ln\left(C_{ab}\right)} = D^a\alpha_c + D^b\beta_c + D^t\gamma_c,$$

Where D^a, D^b and D^t correspond to age, birth cohort and time dummy variables, and α_c, β_c and γ_c correspond to the coefficients capturing age, birth cohort and time effects on consumption.

Since a household's birth cohort is simply a function of the survey year and their age, we need to place some restrictions on the coefficients in this model to enable identification. Following Chamon and Prasad (2010), the birth cohort effects are constrained to sum to zero and be orthogonal to a linear trend:[12]

11 In this exercise we use the 1988/89, 1993/94, 1998/99, 2003/04 and 2009/10 HES, because a longer time period is needed to determine birth cohort effects precisely. While there were some major methodological changes to pre-1998/99 surveys which make it difficult to compare surveys across time, we assume that the cohort and age effects on consumption and income remain comparable.

$$\Sigma_{i=1}^{n}\beta_c(i)=0 \text{ and } \Sigma_{i=1}^{n}\left(\beta_c(i)\times i\right)=0.$$

Household income (Y) can be modelled in a similar way as

$$\overline{\ln(Y_{ab})}=D^a\alpha_y+D^b\beta_y+D^t\gamma_y,$$

where α_y, β_y and γ_y correspond to the coefficients capturing age, birth cohort and time effects on income. Similar constraints apply: $\Sigma_i\beta_y(i)=0$ and $\Sigma_i\left(\beta_y(i)\times i\right)=0$.

Combining the results of the income and consumption models gives the effect that age, birth cohort and time have on household saving, where household saving ratios are calculated as the difference between the fitted values of the dependent income and consumption variables. Figures A1 to A3 show the estimated effect of age, birth cohort and time respectively, assuming the other effects are held constant. Our reference household for this analysis is a household head aged 30 to 34 surveyed in 2009/10. Note that the level of saving shown in the figures depends on the reference household chosen, but the profile of saving does not, so one should focus on how saving changes for different age, birth cohort or time groups, rather than the level of saving *per se*.

Focusing on the age effect, Figure A1 shows how the average household's saving ratio varies with age, holding the survey year and birth cohort constant. The distribution of the age effect partially exhibits the concave relationship predicted by the standard life-cycle model; saving is low early and late in life, and high during a household's working years. One anomaly stands out from the standard life-cycle prediction, however: the dip around middle-age (30 to 50 years), when households reduce their saving before building it back up when they enter the pre-retirement age group.[13]

A possible explanation for this that accords with a slightly amended life-cycle model is simply that costs increase around middle age. Younger households have relatively few living costs and so are able to save for a down-payment on a house, while middle-aged households have children and must pay mortgage interest. The behaviour is also consistent with a myopic model of household behaviour. For example, Thaler and Shefrin (1981) argue that hyperbolic discounting can explain why

12 As argued in Chamon and Prasad, constraining the time effects would force the decomposition to attribute rising consumption and income to age and/or birth cohort effects, rather than an economy-wide rise in productive capacity. Likewise, restraining the age effects would prevent us from examining the life-cycle hypothesis, which makes predictions about how consumption and income should vary with age.

13 As noted in Section 2.1, the saving of self-funded retirees, and so the older age groups, is likely to be overstated.

younger households tend not to save enough for retirement, while Carroll and Samwick (1997) argue that younger households place more weight on saving for large purchases and emergencies to smooth near-term consumption rather than saving for longer-term (retirement) consumption.

Figure A1: Effect of Age on Saving
Holding survey year and birth cohort constant

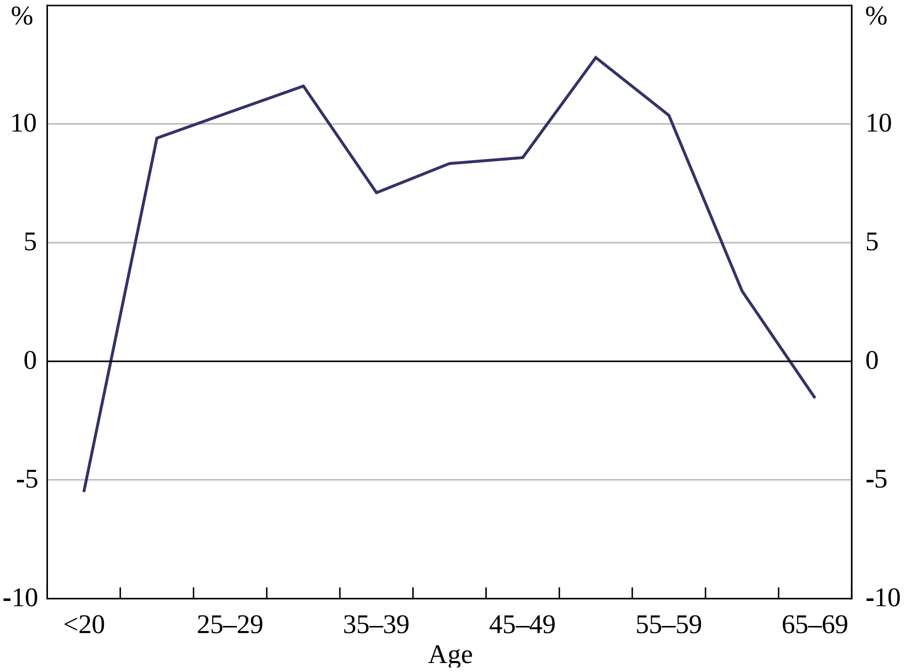

Notes: Gross of depreciation; average by group; survey year = 2009/10, birth cohort = 1975–1980

Sources: ABS; authors' calculations

Figure A2 shows how the average household saving ratio varies with birth cohort, holding the survey year and age of the household head constant; the effects are less clear than those for age, although they suggest that the baby boomer cohort (born between 1946 and 1964) saves more than other birth cohorts throughout their lives.

Time effects in this model represent all determinants of saving not relating to age or birth cohort. Between the 1998/99 and 2003/04 Surveys, the time effect on saving is found to be negligible; on the other hand, the time effect between the 2003/04 and 2009/10 Surveys is large and positive (Figure A3).

Figure A2: Effect of Birth Cohort on Saving

Holding survey year and age constant

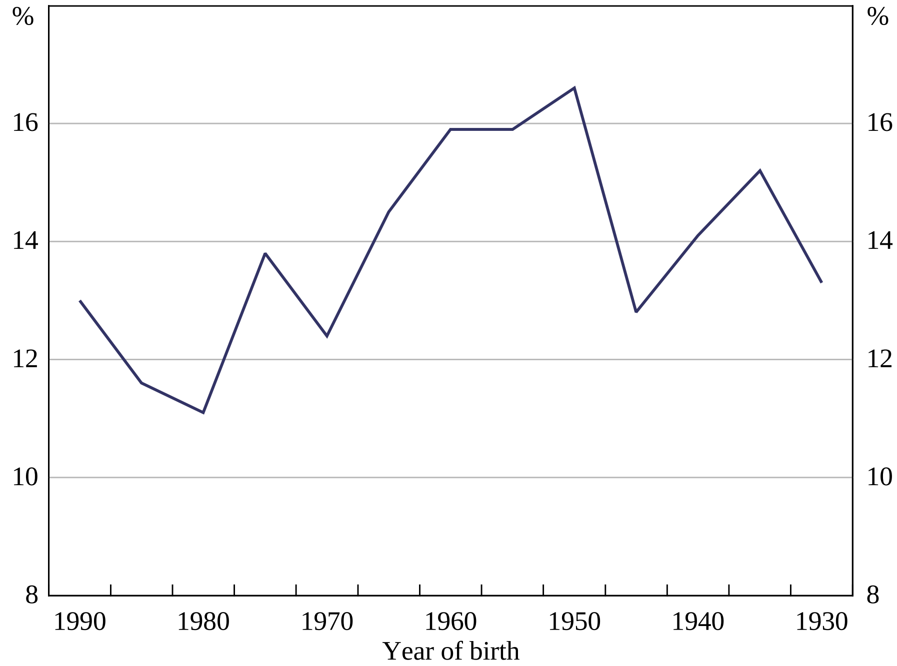

Notes: Gross of depreciation; average by group; survey year = 2009/10, age = 30–34

Sources: ABS; authors' calculations

Figure A3: Effect of Survey Year on Saving

Holding age and birth cohort constant

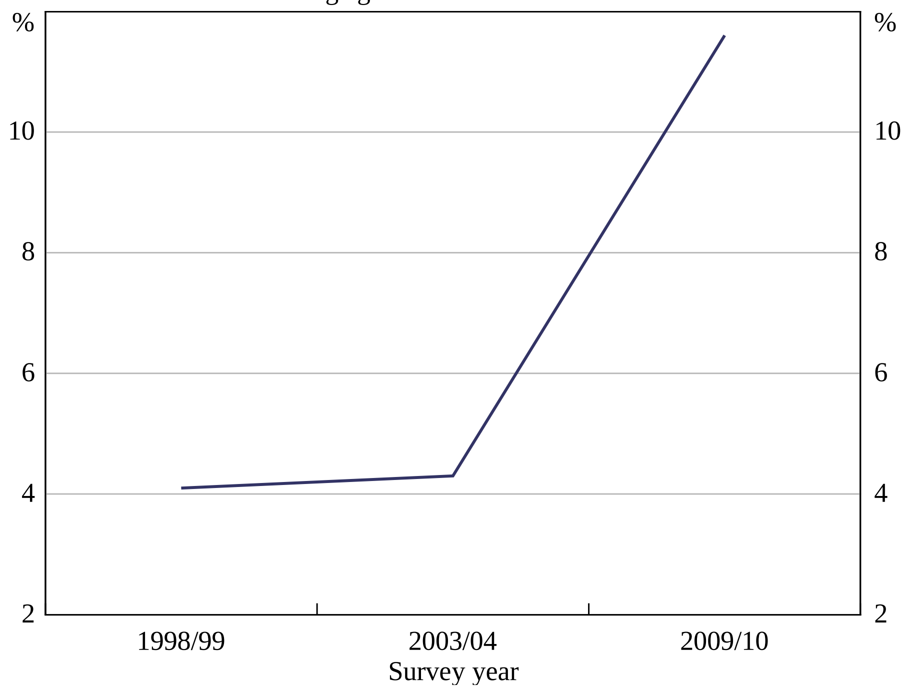

Notes: Gross of depreciation; average by group; birth cohort = 1975–1980, age = 30–34

Sources: ABS; authors' calculation

Appendix B: Auxiliary Regressions

B.1 Permanent Income Model

Since we cannot observe a household's permanent level of income, we estimate it by separately regressing current labour income (for those in the labour force) and non-labour income on proxies for permanent income, and taking the fitted values as measuring permanent income (Table B1).

Table B1: Permanent Income Models

Coefficients

Variable	**2003/04**		**2009/10**	
	Labour income	Non-labour income	Labour income	Non-labour income
Highly educated	0.1***	–0.1***	0.2***	0.0
Migrant	–0.1*	0.0	–0.2***	0.0
Female	–0.2***	0.2***	–0.1*	0.2***
State				
– Vic	–0.1	0.2***	0.0	–0.1***
– Qld	–0.1*	0.2***	–0.1	–0.1**
– SA	–0.2*	0.4***	–0.1	0.1
– WA	–0.3***	0.2***	0.0	–0.2***
– TAS	0.0	0.3***	–0.3	0.1
– ACT and NT	0.0	0.3**	0.4***	0.1
Non-urban	–0.4***	0.1**	–0.2***	–0.1***
Middle-skilled occupation	–0.3***		–0.3***	
Low-skilled occupation	–0.3***		–0.2***	
Young	0.2***	–0.4***	0.2***	–0.3***
Pre-retired	–0.2***	0.5***	–0.1**	0.5***
Old	–2.2***	0.7***	–1.0***	0.8***
Number in work	1.0***		1.1***	
Household size		0.4***		0.5***
Self-funded retiree		0.6***		0.4***
Small business owner		1.2***		1.2***
Non-financial wealth		1.6***		0.1***
Financial wealth		0.3***		0.8***
Government payments		1.3***		1.0***
Constant	5.9***	2.9***	5.8***	3.9***
R^2	0.15	0.42	0.13	0.30

Note: ***, ** and * indicate significance at the 1, 5 and 10 per cent level, respectively

Sources: ABS; authors' calculations

B.2 Risk of Unemployment Model

For households with no unemployed members and a household head aged less than 65 years, the risk of unemployment variable is set equal to one if the fitted value of a logit regression of unemployment status on a range of household characteristics is greater than 10 per cent. In particular, for $unemployed_{it}$ representing a dummy variable that equals one if household *i* has at least one unemployed person in survey *t*, and zero otherwise, we model $unemployed_{it}$ using a number of independent variables as detailed in Table B2.

Table B2: Risk of Unemployment Models

	Coefficients	
Variable	2003/04	2009/10
Highly educated	–0.6***	–0.2
Migrant	0.3*	0.4***
Female	0.5***	0.4***
State		
– Vic	–0.1	0.0
– Qld	0.0	0.1
– SA	0.1	0.1
– WA	0.1	0.0
– TAS	0.0	0.1
– ACT and NT	–0.2	–0.7**
Non-urban	0.0	0.1
Young	0.2	0.1
Pre-retired	0.0	0.4**
Old	–1.0***	–1.6***
Household size	0.3***	0.4***
Constant	–3.3***	–4.0***
R^2	0.06	0.07

Note: ***, ** and * indicate significance at the 1, 5 and 10 per cent level, respectively

Sources: ABS; authors' calculations

B.3 Median and Mean Regression Models

Tables B3 and B4 present full median and mean regression outputs for two models: the model used in the main text (Model (1)), and an alternate model where we drop the income variable (Model (2)). For both tables, ***, ** and * indicate significance at the 1, 5 and 10 per cent level, respectively, where 300 repetitions of bootstrapped HES household weights are used to obtain standard errors.

Table B3: Median Regression Models

Coefficients

Variable	2003/04 (1)	2003/04 (2)	2009/10 (1)	2009/10 (2)
Income elasticity	0.04***	na	0.04***	na
Highly educated	–4.7***	–4.7***	1.9	1.6
Income (>20%)				
- Business	8.4***	3.5	4.3	–2.9
- Salary	10.6***	16.3***	2.8	12.7***
- Government	–4.5*	–3.3	–0.8	–0.5
- Other	–5.1**	–6.2**	–10.1***	–6.8***
Risk of unemployment	0.7	0.9	–1.0	–2.2
Low-skilled occupation	–5.1**	–5.5***	–2.5	–1.6
Middle-skilled occupation	–6.5***	–7.2***	–2.9**	–2.3
Not in the labour force	–9.1**	–2.3	–9.9**	0.7
Self-funded retiree	–19.0***	–15.9***	–8.3**	–3.4
Pensioner	–10.2**	–4.4	–11.7***	–3.4
Household size	2.7***	2.4***	2.3**	2.6**
Share of children	–23.0***	–22.5***	–24.4***	–22.9***
Female	–3.8***	–4.8***	–3.4**	–3.0**
Single-parent household	8.0**	10.7***	5.2**	5.4**
Married	–3.9**	–4.6**	–0.4	–1.9
Migrant	2.4	2.3	2.8*	2.5
Wealth-to-income ratio	–0.1	0.0	–0.2***	–0.1*
Gearing ratio	–13.8***	–15.8***	–2.4	–2.0
Mortgage	–2.9	3.5*	–3.9*	2.8
Own home outright	1.2	6.1***	3.9	9.0***
Credit constrained	2.9	2.2	0.1	0.3
Worse off than a year ago	–6.3***	–5.6***	–4.1***	–3.8***
No of credit cards	–1.8***	–1.6***	–0.9	–0.7
Personal debt	–14.6***	–15.8***	–15.6***	–17.9***
State				
- Vic	–0.7	–0.7	0.2	0.6
- Qld	1.0	1.4	1.4	1.6
- SA	1.4	1.0	5.6***	6.3***
- WA	0.2	0.1	4.1**	4.8***
- TAS	–0.9	–1.1	–0.6	–1.9
- ACT and NT	–3.6	–4.4*	4.7**	4.9**
Non-urban	0.6	1.1	–0.4	–0.3
Young	–0.8	–0.1	–0.3	2.0
Pre-retired	0.3	–0.7	6.8***	5.1***
Old	10.1**	7.5*	7.9*	7.1*
Constant	17.0***	9.0**	14.7***	–0.3
R^2	0.06	0.05	0.06	0.05

Table B4: Mean Regression Models

Coefficients

Variable	2003/04 (1)	2003/04 (2)	2009/10 (1)	2009/10 (2)
Income elasticity	0.05***	na	0.04***	na
Highly educated	–2.5**	–2.6**	3.0***	3.0***
Income (>20%)				
- Business	11.4***	6.5***	9.9***	5.1***
- Salary	4.2*	14.3***	2.6	11.4***
- Government	–9.5***	–8.5***	–3.0*	–1.3
- Other	–12.7***	–12.2***	–13.8***	–12.2***
Risk of unemployment	1.4	1.9	–0.6	–0.3
Low-skilled occupation	–4.6***	–5.0***	2.5	1.8
Middle-skilled occupation	–5.1***	–5.6***	–1.1	–1.5
Not in the labour force	–12.9***	–2.6	–8.5***	–0.4
Self-funded retiree	–8.4***	–8.3***	3.4*	2.6
Pensioner	–7.0*	–1.4	–9.9***	–4.7*
Household size	2.0***	1.3*	2.8***	2.5***
Share of children	–21.4***	–20.8***	–28.5***	–27.7***
Female	–3.9***	–4.4***	–2.5**	–3.1***
Single-parent household	8.8***	11.4***	4.8**	5.1**
Married	–3.0*	–2.8*	0.3	–0.7
Migrant	2.5*	2.2	4.6***	4.7***
Wealth-to-income ratio	–0.1***	–0.1***	–0.1***	0.0
Gearing ratio	–15.9***	–15.9***	–10.1***	–9.7***
Mortgage	–3.8**	2.6*	–1.5	4.0***
Own home outright	1.2	6.3***	3.6**	7.1***
Credit constrained	1.3	1.5	–2.2	–1.7
Worse off than a year ago	–5.1***	–5.2***	–3.0***	–3.3***
No of credit cards	–1.1**	–0.9*	–0.3	–0.2
Personal debt	–17.5***	–19***	–15.4***	–17.3***
State				
- Vic	–0.1	–0.4	–1.3	–1.4
- Qld	–0.7	–0.4	–0.4	–0.2
- SA	–0.4	–0.1	4.7***	4.8***
- WA	0.4	0.7	2.3	2.4
- TAS	–1.0	–1.7	–1.6	–1.0
- ACT and NT	–2.1	–1.9	4.3	3.9
Non-urban	0.7	1.2	–0.8	–0.6
Young	–1.1	–0.3	1.2	2.8**
Pre-retired	–0.5	–1.5	4.4***	3.9***
Old	4.3	5.7	5.8**	6.2**
Constant	19.8***	7.3**	8.5***	–3.0
R^2	0.12	0.10	0.12	0.09

References

ABS (Australian Bureau of Statistics) (2008), 'Experimental Estimates of Imputed Rent Australia', Information Paper, ABS Cat No 6525.0.

Attanasio OP (1999), 'Consumption', in JB Taylor and M Woodford (eds), *Handbook of Macroeconomics, Volume 1B*, Handbooks in Economics 15, Elsevier Science, Amsterdam, pp 741–812.

Attanasio OP and G Weber (1994), 'The UK Consumption Boom of the Late 1980s: Aggregate Implications of Microeconomic Evidence', *The Economic Journal*, 104(427), pp 1269–1302.

Attanasio OP and G Weber (2010), 'Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy', *Journal of Economic Literature*, 48(3), pp 693–751.

Berger-Thomson L, E Chung and R McKibbin (2009), 'Estimating Marginal Propensities to Consume in Australia Using Micro Data', RBA Research Discussion Paper No 2009-07.

Blinder AS (1973), 'Wage Discrimination: Reduced Form and Structural Estimates', *The Journal of Human Resources*, 8(4), pp 436–455.

Browning M and TF Crossley (2001), 'The Life-Cycle Model of Consumption and Saving', *The Journal of Economic Perspectives*, 15(3), pp 3–22.

Browning M and A Lusardi (1996), 'Household Saving: Micro Theories and Micro Facts', *Journal of Economic Literature*, 34(4), pp 1797–1855.

Brzozowski M and TF Crossley (2011), 'Viewpoint: Measuring the Well-Being of the Poor with Income or Consumption: A Canadian Perspective', *The Canadian Journal of Economics*, 44(1), pp 88–106.

Carroll CD (1992), 'The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence', *Brookings Papers on Economic Activity*, 1992(2), pp 61–156.

Carroll CD and AA Samwick (1997), 'The Nature of Precautionary Wealth', *Journal of Monetary Economics*, 40(1), pp 41–71.

Chamon MD and ES Prasad (2010), 'Why are Saving Rates of Urban Households in China Rising?', *American Economic Journal: Macroeconomics*, 2(1), pp 93–130.

Deaton AS (1991), 'Saving and Liquidity Restraints', *Econometrica*, 59(5), pp 1221–1248.

Deaton AS and CH Paxson (1994), 'Saving, Growth, and Aging in Taiwan', in DA Wise (ed), *Studies in the Economics of Aging*, National Bureau of Economic Research Project Report, The University of Chicago Press, Chicago, pp 331–357.

Dvornak N and M Kohler (2003), 'Housing Wealth, Stock Market Wealth and Consumption: A Panel Analysis for Australia', RBA Research Discussion Paper No 2003-07.

Dynan KE and DL Kohn (2007), 'The Rise in US Household Indebtedness: Causes and Consequences', in C Kent and J Lawson (eds), *The Structure and Resilience of the Financial System*, Proceedings of a Conference, Reserve Bank of Australia, Sydney, pp 84–113.

Dynan KE, J Skinner and SP Zeldes (2004), 'Do the Rich Save More?', *Journal of Political Economy*, 112(2), pp 397–444.

Finlay R, and JP Jääskelä (2014), 'Credit Supply Shocks and the Global Financial Crisis in Three Small Open Economies', *Journal of Macroeconomics*, 40, pp 270–276.

Harris MN, J Loundes and E Webster (2002), 'Determinants of Household Saving in Australia', *Economic Record*, 78(241), pp 207–223.

Iacoviello M (2004), 'Consumption, House Prices, and Collateral Constraints: A Structural Econometric Analysis', *Journal of Housing Economics*, 13(4), pp 304–320.

Islam A, J Parasnis and D Fausten (2013), 'Do Immigrants Save Less than Natives? Immigrant and Native Saving Behaviour in Australia', *Economic Record*, 89(284), pp 52–71.

Meyer BD and JX Sullivan (2011), 'Viewpoint: Further Results on Measuring the Well-Being of the Poor using Income and Consumption', *The Canadian Journal of Economics*, 44(1), pp 52–87.

Muellbauer J (2007), 'Housing and Personal Wealth in a Global Context', WIDER (World Institute for Development Economics Research) Research Paper No 2007/27.

Oaxaca R (1973), 'Male-Female Wage Differentials in Urban Labor Markets', *International Economic Review*, 14(3), pp 693–709.

Reserve Bank of Australia (2008), Financial Stability Review, September.

Sabelhaus J and JA Groen (2000), 'Can Permanent-Income Theory Explain Cross-Sectional Consumption Patterns?', *The Review of Economics and Statistics*, 82(3), pp 431–438.

Thaler RH and HM Shefrin (1981), ‘An Economic Theory of Self-Control’, *Journal of Political Economy*, 89(2), pp 392–406.

Windsor C, J Jääskelä and R Finlay (2015), ‘Housing Wealth Effects: Evidence from an Australian Panel’, *Economica*, forthcoming.

Yates J and S Whelan (2009), *Housing Wealth and Consumer Spending*, AHURI Final Report No 132, Australian Housing and Urban Research Institute, Melbourne.

Zeldes SP (1989), ‘Consumption and Liquidity Constraints: An Empirical Investigation’, *Journal of Political Economy*, 97(2), pp 305–346.

Publication III

First published in Rees, D., Lancaster, D. and Finlay, R. (2015). A state-space approach to Australian GDP measurement. *Australia Economic Review*, **48**, 133-149.

A State-Space Approach to Australian Gross Domestic Product Measurement

Daniel M. Rees, David Lancaster and Richard Finlay*

* Rees and Lancaster: Economic Group, Reserve Bank of Australia, New South Wales 2000 Australia; Finlay: Financial Markets Group, Reserve Bank of Australia, New South Wales 2000 Australia, and School of Economics, The University of New South Wales, New South Wales 2052 Australia. Corresponding author: Rees, email <reesd@rba.gov.au>. We would like to thank Adrian Pagan, James Holloway, Troy Gill, Jamie Hall, Christian Gillitzer, Gianni La Cava, Matthew Read, Peter Tulip, Alexandra Heath, Christopher Kent, Dan Smith, Paul Roberts, Michael Smedes and two referees, along with seminar participants at the Reserve Bank of Australia and the Australian Bureau of Statistics for discussions and helpful suggestions. The views expressed in this article are those of the authors and do not necessarily reflect the views of the Reserve Bank of Australia.

Abstract

We use state-space methods to construct new estimates of Australian gross domestic product growth from the published national accounts estimates of expenditure, income and production. Across a range of specifications, our measures are substantially less volatile than headline domestic product. We conclude that much of the quarter-to-quarter volatility in Australian domestic product growth reflects measurement error, rather than true shifts in the level of economic activity.

1. Introduction

The level and growth of real economic activity are of great interest to economic policy-makers as well as the general public. Increases in activity are typically associated with rising living standards and economic activity influences other economic outcomes, such as inflation and unemployment. But, measuring economic activity is difficult. In Australia, a key measure of activity, gross domestic product (GDP), is measured using three different approaches, based on expenditure (GDP(E)), income (GDP(I)) and production (GDP(P)).[1] Conceptually, the three measures should be equal, but in practice the measures differ because they are constructed from different data sources and have varying degrees of measurement error.[2]

In this article, we use state-space methods to combine the three ABS measures of GDP into an estimate of aggregate economic growth. In contrast to existing approaches, our method allows us to capture three salient features of GDP measurement. First, GDP(E), GDP(I) and GDP(P) should be equal. Second, all three of these quantities are measured with some degree of error. Third, because of overlap between the data sources that feed into the three published estimates of GDP, these measurement errors are likely to be correlated. Once we account for these features of the data, we generate an estimate of economic activity which is smoother than conventional measures of GDP. This suggests that many large quarterly fluctuations in the rate of economic growth reflect errors in measurement, rather than fundamental shifts in the pace of economic activity.

In Australia, the most common alternative to our approach is to take a simple average of the three measures, known as GDP(A).[3] The ABS considers this to be the most reliable estimate of final output, in part because independent errors in the underlying measures are often offsetting (Aspden 1990; ABS 2011). More broadly, the literature on model averaging suggests that if one possesses a set of estimates for some quantity being measured, then a combination of the estimates tends to perform better than any individual estimate.[4]

While using a simple average of the three GDP measures as an estimate for actual GDP is simple and transparent, it does not fully exploit all available information. For example, if one measure of GDP is particularly noisy, so that any given observation is likely to be quite different from actual GDP, then it may make sense to place less weight on that measure and more weight on the remaining two. The technique we explore in this article provides one way of achieving this: it uses the time-series properties of the three GDP measures to construct a composite GDP measure that more fully exploits the available information.

Our article builds on the existing literature on GDP measurement. Most directly, it represents an application to Australian data of the techniques derived by Aruoba et al. (2013), who construct a state-space measure of US GDP.[5] The Australian dimension of our study is of interest for two reasons, aside from our natural curiosity as Australian researchers. First, whereas the US statistical authorities only construct income and expenditure measures of GDP at a quarterly frequency, the ABS also publishes a production measure. We show that the methods of Aruoba et al. (2013) extend to this environment. Second, the Australian economy differs in several respects from that of the

United States, in ways that may make GDP measurement more challenging. In particular, Australia is a smaller, more trade-exposed economy with a large resource sector. Our results support the idea that these variations in economic structure translate into a different pattern of GDP measurement errors in Australia.

Our work is also related to research evaluating the relative merits of expenditure, income and production as measures of economic activity. The primary focus of the research to date has been on the US economy, for which the most widely reported measure of output is derived from the expenditure side of the accounts. Despite this, a common finding is that expenditure-side estimates of output in the United States suffer from more severe measurement issues than income-side estimates. In particular, estimates of US GDP(I) tend to be less variable than GDP(E), while also being more highly correlated with other indicators of economic conditions (Fixler and Grimm 2006; Nalewaik 2010, 2011). Furthermore, in the United States, GDP(E) tends to be revised towards GDP(I) over time.

Research using Australian national accounts data favours the use of the production-side rather than expenditure- or income-side estimates (Aspden 1990; ABS 2012; Bishop, Gill and Lancaster 2013). The relatively large share of resources in Australian GDP makes measures of output particularly responsive to trade data. Timing differences in imports and exports and variability in trade prices can introduce noise into estimates of expenditure and income (ABS 2012). In addition, GDP(I) and GDP(E) are reliant on the ABS' register of businesses, which is typically updated with a delay. Bishop, Gill and Lancaster (2013) found that GDP(P) tends to be revised less than the other two measures and is as reliable in real time as GDP(A). These factors provide a case for applying a larger weight on GDP(P) in model averaging.

While it is useful to know the relative merits of expenditure, income and production measures of economic activity, using just one measure is unlikely to be optimal. The techniques that we use in this article allow information from all three measures of GDP to be combined, with more weight being placed on the more reliable measures.

2. Estimating Gross Domestic Product Growth

We treat GDP growth as an unobserved variable that follows a first-order autoregressive (AR(1)) process:

$$\Delta y_t = \mu(1-\rho) + \rho y_{t-1} + \varepsilon_{G,t} \quad (1)$$

where Δy_t represents the growth rate of real GDP, μ is the mean growth rate of GDP, ρ indicates persistence and $\varepsilon_{g,t}$ is a normally distributed innovation. It is common to model GDP growth as an AR(1) process, as growth rates are typically assumed to be homoscedastic and moderately persistent.

We then assume that the three observed GDP measures—GDP(E), GDP(I) and GDP(P)—provide noisy readings of actual GDP. For example, in our model, the growth rate of GDP(E) is equal to the growth rate of actual GDP plus a measurement error term:

$$\Delta y_t^E = \Delta y_t + \varepsilon_{E,t}$$

Stacking the three observed measures in matrix form gives us our measurement equation:

$$\begin{bmatrix} \Delta y_t^E \\ \Delta y_t^I \\ \Delta y_t^P \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \Delta y_t + \begin{bmatrix} \varepsilon_{E,t} \\ \varepsilon_{I,t} \\ \varepsilon_{P,t} \end{bmatrix} \quad (2)$$

where Δy_t^E, Δy_t^I and Δy_t^P represent growth in GDP(E), GDP(I) and GDP(P) and $\varepsilon_{E,t}$, $\varepsilon_{I,t}$ and $\varepsilon_{P,t}$ represent their measurement errors.

Using this basic framework, we estimate three models that differ in their treatment of the observable variables, the shocks to GDP and the measurement errors.

2.1 Model 1: No Correlation

Our first model assumes that all stochastic terms are independent; that is, $(\varepsilon_{G,t}, \varepsilon_{E,t}, \varepsilon_{I,t}, \varepsilon_{P,t}) \sim N(0,\Sigma)$ where:

$$\Sigma = \begin{bmatrix} \sigma_G^2 & 0 & 0 & 0 \\ 0 & \sigma_E^2 & 0 & 0 \\ 0 & 0 & \sigma_I^2 & 0 \\ 0 & 0 & 0 & \sigma_P^2 \end{bmatrix}$$

2.2 Model 2: Correlation

Next, we allow for correlation between the various GDP measures; that is, for $\left(\varepsilon_{G,t}, \varepsilon_{E,t}, \varepsilon_{I,t}, \varepsilon_{P,t}\right) \sim N(0,\Sigma)$ where:

$$\Sigma = \begin{bmatrix} \sigma_G^2 & \sigma_{GE} & \sigma_{GI} & \sigma_{GP} \\ \sigma_{GE} & \sigma_E^2 & \sigma_{EI} & \sigma_{EP} \\ \sigma_{GI} & \sigma_{EI} & \sigma_I^2 & \sigma_{IP} \\ \sigma_{GP} & \sigma_{EP} & \sigma_{IP} & \sigma_P^2 \end{bmatrix}$$

This model allows the errors in the three observable GDP measures to be inter-related and for the size of the shock to actual GDP to affect the measurement error in the observed measures of GDP. For example, large innovations in actual GDP may be associated with less precise estimates of GDP(E), GDP(I) and/or GDP(P) than is the case for small innovations.

In order to identify the model, we must place at least one restriction on the Σ matrix.[6] In line with Aruoba et al. (2013), we impose this restriction by requiring that:[7]

$$\zeta = \frac{Var\left(\Delta y_t\right)}{Var\left(\Delta y_t^E\right)} = \frac{\frac{1}{1-\rho^2}\sigma_G^2}{\frac{1}{1-\rho^2}\sigma_G^2 + 2\sigma_{GE} + \sigma_E^2} = 0.6 \quad (3)$$

That is, we assume that the variance of actual GDP growth is equal to just over half of the variance of the observed GDP(E) growth series. Although intuitively appealing, the restriction is arbitrary and any number of alternative restrictions would also suffice.[8]

2.3 Model 3: Unemployment

Our third model includes an additional observable variable that depends on GDP growth but whose measurement error is unrelated to that of the other observable variables: the quarterly change in the unemployment rate. That is, we replace Equation (2) with:

$$\begin{bmatrix} \Delta y_t^E \\ \Delta y_t^I \\ \Delta y_t^P \\ \Delta U_t \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ \kappa \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \\ 1 \\ \lambda \end{bmatrix} \Delta y_t + \begin{bmatrix} \varepsilon_{E,t} \\ \varepsilon_{I,t} \\ \varepsilon_{P,t} \\ \varepsilon_{U,t} \end{bmatrix} \qquad (4)$$

where ΔU_t is the change in the unemployment rate. In this case, we assume that $(\varepsilon_{G,t}, \varepsilon_{E,t}, \varepsilon_{I,t}, \varepsilon_{P,t}, \varepsilon_{U,t}) \sim N(0, \Sigma)$ with:

$$\Sigma = \begin{bmatrix} \sigma_G^2 & \sigma_{GE} & \sigma_{GI} & \sigma_{GP} & \sigma_{GU} \\ \sigma_{GE} & \sigma_E^2 & \sigma_{EI} & \sigma_{EP} & 0 \\ \sigma_{GI} & \sigma_{EI} & \sigma_I^2 & \sigma_{IP} & 0 \\ \sigma_{GP} & \sigma_{EP} & \sigma_{IP} & \sigma_P^2 & 0 \\ \sigma_{GU} & 0 & 0 & 0 & \sigma_U^2 \end{bmatrix}$$

That is, we impose three restrictions on the Σ matrix: zero correlation between the innovation to the change in the unemployment rate $(\varepsilon_{U,t})$ and the measurement error for growth of GDP(E), GDP(I) and GDP(P) $(\varepsilon_{E,t}, \varepsilon_{I,t}, \varepsilon_{P,t})$. With these restrictions, the model is identified (in fact, the model is over-identified).

3. Estimation

We follow the approach of Aruoba et al. (2013) and estimate the models within a Bayesian framework. We work with Model 3 in this section; Models 1 and 2 are nested in Model 3 and can be recovered by setting appropriate parameters to zero.

First, we express our model in state-space form. Let $s_t = \left[\Delta y_t, \varepsilon_{E,t}, \varepsilon_{I,t}, \varepsilon_{P,t}, \varepsilon_{U,t}\right]'$, $m_t = \left[\Delta y_t^E, \Delta y_t^I, \Delta y_t^P, U_t\right]'$, $M = \left[\mu(1-\rho), 0,0,0,0\right]'$, $K = [0,0,0,\kappa]'$, $\psi = \left[\varepsilon_{G,t}, \varepsilon_{E,t}, \varepsilon_{I,t}, \varepsilon_{P,t}, \varepsilon_{U,t}\right]'$,

$$A = \begin{bmatrix} \rho & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \text{ and } C = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ \lambda & 0 & 0 & 0 & 1 \end{bmatrix}$$

Then, we can express Model 3 as:

$$\begin{aligned} s_t &= M + As_{t-1} + \psi_t \\ m_t &= K + Cs_t \end{aligned}$$

For ease of notation, we collect the parameters in the vector $\Theta = \left(\mu, \rho, \kappa, \lambda, \sigma_G^2, \sigma_{GE}, \sigma_E^2, \sigma_{GI}, \sigma_{EI}, \sigma_I^2, \sigma_{GP}, \sigma_{EP}, \sigma_{IP}, \sigma_P^2, \sigma_{GU}, \sigma_U^2\right)$.

We use the Metropolis–Hastings Markov Chain Monte Carlo (MCMC) algorithm to estimate model parameters.[9] We first maximise the posterior distribution of Θ, given the observed data:

$$p(\Theta \mid m_{1:T}) \propto p(m_{1:T} \mid \Theta) p(\Theta)$$

where $p(m_{1:T} \mid \Theta)$ is the density of the observable data, given the model parameters, and $p(\Theta)$ is the density of the priors over the parameter draw. This gives us an initial estimate of Θ, denoted as Θ^0. We use the inverse Hessian at the maximum to obtain an estimate of the covariance matrix of Θ, Ω^0. Θ^0 and Ω^0 are then used to initiate the MCMC algorithm: at each iteration *i*, we draw a proposed parameter vector $\Theta^* \sim N\left(\Theta^{i-1}, c\Omega^{i-1}\right)$.

Here, c is a scaling parameter set to achieve an acceptance rate of around 25 per cent. We accept Θ^* as Θ^i with probability

$$\min\left(1, \frac{p\left(m_{1:T} \mid \Theta^*\right)p\left(\Theta^*\right)}{p\left(m_{1:T} \mid \Theta^{i-1}\right)p\left(\Theta^{i-1}\right)}\right)$$

and set $\Theta^i = \Theta^{i-1}$ otherwise. We set $p\left(\Theta^*\right) = 0$ if Θ^* is not a valid draw; for example, if it implies a covariance matrix that is not positive definite.

In order to sample Θ^* from the $N\left(\Theta^{i-1}, c\Omega^{i-1}\right)$ distribution, we need to evaluate $p\left(m_{1:T} \mid \Theta\right)$, the density of the observable data given the model parameters. To do this, we use the standard Kalman filter and simulation smoother, as described in Durbin and Koopman (2012). We take 50,000 draws from the posterior distribution and discard the first 25,000.

3.1 Priors

Our prior for the mean growth rate of GDP, μ, follows a normal distribution, with mean of 0.80 and standard deviation of 10.[10] The mean of this prior corresponds to the average quarterly growth rate of GDP over our sample, while the standard deviation is extremely large relative to the volatility of the GDP series, indicating that this prior places only a very weak restriction on the range of potential values. For the persistence of shocks to GDP growth, ρ, we use a beta prior, with mean of 0.50 and standard deviation of 0.20. The prior restricts the value of this parameter to lie between 0 and 1, consistent with GDP growth being a stationary series.[11]

For the variances of the shocks to GDP and the measurement errors, we impose inverse-gamma priors, with mean of 2 and standard deviation of 4. These priors ensure that the variances of all shocks are greater than 0. Finally, for the covariance terms, the priors follow a normal distribution, with mean of 0 and standard deviation of 5.

In all cases, our priors are loose, ensuring that we place a large weight on information from the data, but rule out unreasonable parameter values.

3.2 Data

Our data span is 1980Q1–2013Q2. The starting date reflects the fact that, while Australian national accounts data are available on a quarterly basis from 1959Q3, the quality of the underlying data sources has changed over time so that the pattern of measurement errors in the early years of each GDP series may be unrepresentative of their current performance. The GDP and unemployment rate data that we use in our estimation are all seasonally adjusted by the ABS and the GDP data are expressed in real terms.[12]

4. Results

4.1 Model 1: No Correlation

In Model 1, we assume that shocks to GDP and the measurement errors are independent of each other. Table 1 shows the parameter estimates. The median estimate of μ is 0.78, which is close to the average growth rate of GDP(A) over the sample. The estimate of ρ is 0.40. This implies that the GDP growth process has relatively little persistence, although the parameter is larger than estimates from an AR(1) model of GDP(A) growth over our sample. Innovations to GDP growth are estimated to be similar in size to the measurement errors in the expenditure and production equations and smaller than the average measurement errors in the income equation.

By taking draws from the posterior distribution of the model's parameter values, we can recover an estimate of 'true' GDP growth over the sample. For each parameter draw, we apply a simulation smoother to the underlying GDP data to obtain an estimate of 'true' GDP growth conditional on the parameters. We call the median of this GDP series estimated from Model 1 'GDP(M1)'. Figure 1 plots this series against the published quarterly growth rates of GDP(A), GDP(E), GDP(I) and GDP(P). The GDP(M1) is highly correlated with GDP(A), but it is less volatile.[13] That is, our model suggests that some extreme readings of GDP(A) are likely to represent measurement error in one or more of the individual measures of GDP.[14]

4.2 Model 2: Correlation between Gross Domestic Product Innovations and Measurement Errors

In Model 2, we allow for correlation between innovations to GDP and the measurement errors. Table 2 presents the parameter estimates. The estimated parameters of the GDP process, μ and ρ, are similar to those in Model 1. However, the variance of innovations to GDP and the measurement errors are larger. This is most notable in the expenditure equation, where the variance of the measurement errors is now similar in magnitude to the income equation. This is not an artefact of the restriction imposed in Equation (3); varying the restriction or applying it to GDP(P), rather than GDP(E), leaves the value of σ_E^2 largely unchanged. In contrast, the variance of the measurement errors in the production equation remains around the same size as for the estimated GDP innovations.

The covariances between the measurement errors are positive and generally statistically significant. This is consistent with the fact that information from some surveys feeds into more than one measure of GDP. In contrast, covariances between innovations to GDP and the measurement errors are generally negative and statistically significant. This suggests that the characteristics of measurement errors vary over the business cycle, perhaps because the types of challenges the ABS faces in measuring GDP growth vary across the business cycle. In general, the fact that the covariances of innovations to GDP and measurement errors are statistically significant highlights the importance of controlling for these correlations when evaluating the pace of economic growth.

Figure 2 shows the plot of GDP derived from Model 2, GDP(M2). This measure is considerably smoother than GDP(A). This reflects the fact that when we allow for correlation between shocks, some large changes in multiple measures are attributed to measurement error, rather than treated as signal.

4.3 Model 3: Unemployment

In Model 3, we include the quarterly change in the unemployment rate as an additional observable variable. Table 3 presents the parameter estimates. The estimated mean parameter for the GDP process is similar to the previous models, although GDP growth has more persistence than in Models 1 and 2. The coefficients in the unemployment equation suggest that a one percentage point increase in the rate of quarterly GDP growth lowers the unemployment rate by around 0.6 percentage points, which is slightly above existing Okun's law estimates for Australia (Borland 2011).

The parameter estimates for the shock processes differ from the previous models' in two respects. First, the variance of GDP innovations is much smaller when we include the unemployment rate as an observable variable in the model. Second, the negative correlation between GDP innovations and measurement errors in the GDP(M2) measurement equations largely disappears. However, the covariances between the measurement errors remain positive and statistically significant.

Figure 3 compares this model's estimate of GDP growth to the published figures. Overall, the results for Model 3 are similar to those of Model 2 as, once again, our measure of GDP is smoother than GDP(A). The greatest difference lies in the recessions of the early 1980s and 1990s and the slowdown associated with the Global Financial Crisis in the late 2000s. The inclusion of the unemployment rate, which increased in all three episodes, lowers Model 3's estimate of GDP growth relative to the estimates in Models 1 and 2.

4.4 How Do Our Measures Compare with the Published Trend Measure of Gross Domestic Product?

Our methodology provides measures of GDP growth that incorporate information about the degree of noise generated by measurement error in the published estimates. The result is a smoother measure. The ABS also produces a smoother measure of output growth, constructed by applying a Henderson moving average to GDP(A). The ABS publishes the resulting measure, known as 'trend' GDP, at a quarterly frequency. Figure 4 compares the ABS' trend GDP with the measures introduced in this article.

The histories of the series are generally quite similar, which is encouraging. Trend GDP(A) has a disadvantage relative to our method, however, in that it suffers from end-point problems. The Henderson trends used by the ABS apply moving averages to past and future observations in a series. As the series approaches its end-point, there are fewer observations upon which to calculate these averages. While the ABS takes steps to ameliorate this issue, recent trend GDP data remain subject to substantial revision as new data are received.[15] In Section 5, we demonstrate that the techniques presented in this article appear to be less affected by end-point problems and so should provide users with a better indication of output growth in real time.

4.5 What Are the Relative Contributions of GDP(E), GDP(I) and GDP(P)?

At its core, our methodology represents an alternative way of combining the information in the ABS' three existing measures of GDP growth. One might wonder how our model weights each of the three measures and the extent to which this differs from the simple average used to construct GDP(A). To answer this question, we examine the Kalman gains, which govern the extent to which our models adjust their estimates of the rate of GDP growth in light of new observations of GDP(E), GDP(I) or GDP(P).

For each draw from the posterior distribution of model parameters, we can recover an estimate of the Kalman gain for each observable variable. Figure 5 summarises these Kalman gains for Model 3.[16] In the graph, each light-coloured dot shows the Kalman gains of two measures of GDP for an individual draw from the posterior distribution. The dark-coloured dot and circle represent the posterior median and 90 per cent probability interval for each pair. Intuitively, if most dots lie to the left of the dashed 45 degree line, then the Kalman gain for the observed GDP measure on the vertical axis is greater than that of the measure on the horizontal axis and vice versa. A mass of dots surrounding the dashed 45 degree line indicates that the model puts roughly equal weight on the two observed measures of GDP.

Figure 5 confirms that the model places more weight on GDP(P) than on the other two measures. It also places roughly equal weight on GDP(E) and GDP(I). This is consistent with the fact that the estimated measurement errors in the production equation are considerably smaller than those in the expenditure and income equations.

4.6 Gross Domestic Product Behaviour during Slowdowns

Although our measures of GDP exhibit similar cycles to GDP(A), the quarter-to-quarter growth rates differ. These differences are most relevant around business cycle turning points, when distinguishing signal from noise in GDP growth is of greatest importance.

In this section, we discuss the behaviour of our models during the Australian economy's two most recent slowdowns, which occurred in 2000–01 and 2008–09.

In both of these episodes, GDP(A) indicates that the Australian economy experienced a large contraction in economic activity, followed by a strong recovery in the subsequent quarter. In the earlier episode, the economy returned rapidly to trend growth. In contrast, in the third quarter of the 2008–09 episode, GDP growth slowed again, with GDP(A) expanding by a mere 0.1 per cent in the June quarter of 2009.

In the presence of measurement error, large changes in economic activity make policy-making difficult. Did the strong GDP growth recorded in the March quarters of 2001 and 2009 accurately signal that the economy had recovered from the declines of previous quarters? Or was it merely statistical noise that concealed ongoing economic weakness?

Our models suggest that neither the slowdown of 2000–01 nor the subsequent recovery was as dramatic as the GDP(A) outcome suggests (Figure 6). Models 2 and 3 suggest that the economy experienced a period of two-to-three quarters of substantially below-average growth, but did not actually contract. Model 1 displays a similar quarterly pattern to GDP(A), but with less extreme movements. All of the models suggest that, by early 2001, growth in economic activity had begun to recover.

In contrast, according to our models, the slowdown of 2008–09 was more prolonged than indicated by GDP(A) (Figure 7). Model 2 suggests that the economy experienced at least three quarters of growth substantially below average. Also, Model 3 records two consecutive quarters of negative growth in the December quarter of 2008 and March quarter of 2009. This is consistent with the beliefs of policy-makers at the time that the Australian economy was in recession in early 2009 (Stevens 2009). All three measures assign a large proportion of the recovery in GDP(A) growth in the March quarter of 2009 to measurement error. This is consistent with the fact that the increase in GDP growth in that quarter was primarily observable in GDP(E) and GDP(I), to which the models apply relatively less weight.

4.7 Is Australian Gross Domestic Product Measurement Different?

A natural benchmark against which to compare our results is Aruoba et al. (2013), who conduct a similar exercise using US data. Our results differ from theirs in two respects.

First, across a number of specifications, Aruoba et al. (2013) find that the average size of measurement errors in US GDP(I) is smaller than in US GDP(E). Consequently, their model places more weight on income than expenditure in constructing a measure of US GDP growth. In contrast, we find that in the Australian data, measurement errors on the income side of the accounts tend to be larger than on the expenditure side.

Second, Aruoba et al. (2013) find that innovations to US GDP are, on average, larger than measurement errors. In contrast, we find larger relative measurement errors in Australian GDP data.

We argue that differences in the structure of the US and Australian economies could explain the different pattern of GDP and measurement errors in the two economies. Relative to the United States, Australia is a smaller and more open economy and commodity exports are relatively more important. Given that commodity prices are typically more volatile than manufacturing or services prices, commodity exporters tend to experience greater volatility in nominal GDP than other economies. This nominal volatility makes real GDP measurement on the income side of the national accounts challenging because of the need to determine appropriate deflators to apply to volatile nominal GDP flows. Similar challenges apply when measuring expenditure; in particular, export and import volumes. To the extent that commodity prices and exchange rates are observable, it should be possible to deflate export and import values accurately. However, if prices and exchange rates are volatile, imposing appropriate deflators is more difficult, creating the possibility of additional measurement error. The volatility of Australian export prices could go some way to explaining the relatively large measurement errors that we report for Australian GDP(E) and GDP(I).

5. Comparison with GDP(A)

It is natural to compare the performance of our models against GDP(A). We first describe the statistical properties of our GDP measures and we then examine whether our GDP measures are better able to explain and forecast unemployment and inflation.

5.1 How Volatile Is Gross Domestic Product Growth?

Visual inspection suggested that our measures of GDP growth smooth out some of the volatility in the published ABS series. A statistical analysis of the alternative GDP measures confirms this conjecture.

Table 4 compares moments of the published GDP series to those of our models. The mean of our models is similar to those of the published series. However, other moments of the distributions differ. All three of our constructed measures are considerably less volatile than the ABS series, with the standard deviation of GDP growth being around one-third lower in our series than in GDP(A).

Our measures of GDP growth are also more persistent, with the correlation coefficients on our measures of GDP growth far larger than on the ABS series. As a consequence, our measures of GDP growth are also more predictable: an estimated AR(1) model of our constructed GDP series produces a far closer fit than it does for standard measures of GDP.

5.2 Real-Time Performance

In order to produce timely estimates, statistical agencies publish GDP before all information sources are available. They then revise these preliminary estimates as more information comes to light.[17] Bishop, Gill and Lancaster (2013) find that initial estimates of Australian GDP often differ substantially from later, more informed estimates. Knowing this, users may prefer measures that are less subject to revision, as long as those measures are close approximations to 'true' output growth.

We use real-time estimates of GDP(E), GDP(I) and GDP(P) to construct a history of real-time model estimates from 2001Q1 to 2013Q2.[18] We evaluate the real-time performance of our models using two common metrics: 'mean absolute revision' and 'mean revision'. Mean absolute revision measures the average size of revisions, regardless of sign. Mean revision is the average of revisions and can be interpreted as a tendency for GDP to be revised in a particular direction; that is, whether it is biased. Table 5 presents these statistics for GDP(A) and our models over the period 2001Q1–2009Q3. Final GDP is defined as the estimate after 4 years, consistent with Bishop, Gill and Lancaster (2013).

Our models are more reliable than GDP(A) in real time, with the mean absolute revisions statistically smaller than for GDP(A) at the 5 per cent level. These differences are economically meaningful as well: revisions to our models are around one-third smaller than revisions to GDP(A). Consistent with Bishop, Gill and Lancaster (2013), over this sample, there has been a slight upward tendency to revisions and this is evident across the measures, although less so for the model estimates.

In addition, our models' GDP growth estimates are also easier to forecast in real time. The root mean squared errors for out-of-sample AR(1) forecasts are 0.39 percentage points for Model 1 and 0.35 percentage points for Models 2 and 3, compared with 0.52 percentage points for GDP(A). This suggests that contemporaneous estimates from our models may provide a better indication of future out-turns than GDP(A).

Table 6 shows that our models also converge to their final values more quickly than GDP(A). The performance of Model 1 is particularly noteworthy as this model is highly correlated with GDP(A).

5.3 Explaining Unemployment

In this section, we examine whether our measures of GDP display a closer relationship with unemployment than GDP(A).[19] Macroeconomic theories typically predict a close relationship between output growth and unemployment, a relationship known as 'Okun's Law'. Figure 8 illustrates the Okun's Law relationship for Australia for GDP(A) and Models 1 and 2.[20] As theory would suggest, for all three measures, lower output growth is associated with an increase in the unemployment rate. But, the relationship between changes in unemployment and changes in GDP appears stronger for Models 1 and 2 than for GDP(A), represented by steeper fitted lines and higher adjusted R^2-values.

We confirm this result more formally by examining the in-sample fit of our models and GDP(A). To do this, we estimate the specification:

$$\Delta U_t = \alpha + \gamma U_{t-1} + \sum_{i=0}^{2} \beta_i \Delta Y_{t-i} + \varepsilon_t \qquad (5)$$

where ΔY_{t-i} is the quarterly growth rate of a measure of GDP in quarter $t - i$ and ΔU_{t-i} is the change in the unemployment rate in quarter $t - i$. The long-run response of unemployment to changes in output, known as 'Okun's coefficient', can be approximated by:

$$C = \sum_{i=0}^{2} \beta_i / (1 - \gamma)$$

Table 7 presents our results. The estimated coefficients are mostly statistically significant and are of the expected sign. The regressions, including Models 1 and 2, appear to fit the data better than those including GDP(A), as shown by the adjusted R^2-values, although the difference is not statistically significant. The coefficients on the lagged changes in unemployment are smaller in the regressions with our model measures, suggesting that our measures contribute more information than GDP(A) or a random walk. Finally, the coefficients on contemporaneous and lagged values of output growth are larger for our models than for GDP(A), reflected in larger Okun's coefficients. This may indicate attenuation bias in the regressions, including GDP(A), caused by the presence of measurement error.

6. Conclusion

In this article, we have used state-space methods to extract new indicators of underlying economic activity from the noisy published measures of expenditure, income and production. Although our measures are highly correlated with published GDP growth, they are noticeably less volatile and easier to forecast. Moreover, they explain variations in unemployment as well as or slightly better than the published GDP growth measures. Our measures also perform well in real time.

First version received September 2014;
final version accepted December 2014 (Eds).

References

An, S. and Schorfheide, F. 2007, 'Bayesian analysis of DSGE models', *Econometric Reviews*, vol. 26, pp. 113–72.

Aruoba, S. B., Diebold, F. X., Nalewaik, J., Schorfheide, F. and Song, D. 2013, 'Improving GDP measurement: A measurement-error perspective', National Bureau of Economic Research Working Paper no. 18954, Cambridge, Massachusetts.

Aspden, C. 1990, 'Which is the best short-term measure of gross domestic product? A statistical analysis of the short-term movements of the three measures of gross domestic product and their average', in *Australian National Accounts: National Income and Expenditure, June Quarter 1990*, Cat. no. 5206.0, Australian Bureau of Statistics, Canberra.

Australian Bureau of Statistics 2007, *Quality Dimensions of the Australian National Accounts*, Cat. no. 5216.0.55.002, ABS, Canberra.

Australian Bureau of Statistics 2011, *Spotlight on National Accounts: Which Is the Best Short-Term Measure of Gross Domestic Product?, May 2011*, Cat. no. 5202.0, ABS, Canberra.

Australian Bureau of Statistics 2012, *Australian System of National Accounts: Concepts, Sources and Methods, Edition 3*, Cat. no. 5216.0, ABS, Canberra.

Bishop, J., Gill, T. and Lancaster, D. 2013, 'GDP revisions: Measurement and implications', *Reserve Bank of Australia Bulletin*, March, pp. 11–22.

Borland, J. 2011, 'The Australian labour market in the 2000s: The quiet decade', in *The Australian Economy in the 2000s: Proceedings of a Conference*, eds H. Gerard and J. Kearns, Reserve Bank of Australia, Sydney.

Durbin, J. and Koopman, S. J. 2012, *Time Series Analysis by State Space Methods*, 2nd edn, Oxford University Press, Oxford.

Fixler, D. and Grimm, B. 2006, 'GDP estimates: Rationality tests and turning point performance', *Journal of Productivity Analysis*, vol. 25, pp. 213–29.

Laplace, P. S. 1818, *Deuxième Supplément a la Théorie Analytique des Probabilités*, Courcier, Paris.

Nalewaik, J. 2010, ‘The income- and expenditure-side measures of output growth’, *Brookings Papers on Economic Activity*, vol. 41, pp. 71–127.

Nalewaik, J. 2011, ‘The income- and expenditure-side estimates of U.S. output growth—An update to 2011Q2’, *Brookings Papers on Economic Activity*, vol. 42, pp. 385–402.

Rees, D., Lancaster, D. and Finlay, R. 2014, ‘A state-space approach to Australian GDP measurement’, Reserve Bank of Australia Research Discussion Paper no. 2014-12, Sydney.

Scutella, R. 1996, ‘Obtaining an efficient estimate of the true level of GDP’, Honours thesis, The University of Melbourne.

Stevens, G. 2009, ‘The road to recovery’, *Reserve Bank of Australia Bulletin*, May, pp. 16–23.

Timmermann, A. 2006, ‘Forecast combinations’, in *Handbook of Economic Forecasting*, eds G. Elliott, C. Granger and A. Timmermann, North-Holland, Amsterdam.

White, H. 1980, ‘A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity’, *Econometrica*, vol. 48, pp. 817–38.

Figure 1 Gross Domestic Product (GDP) Growth: Comparison with GDP(M1)

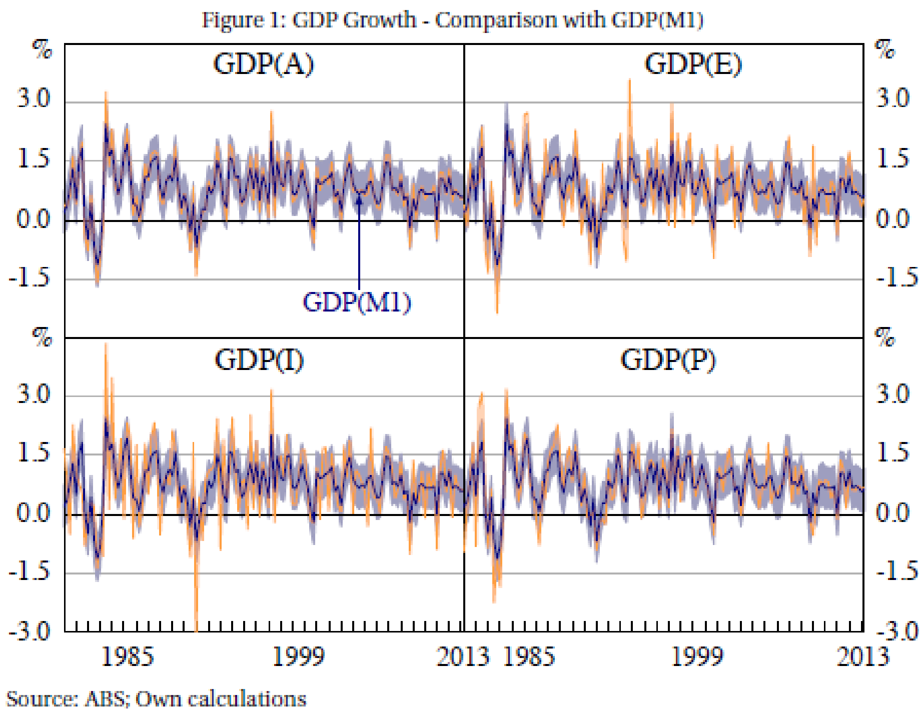

Sources: Australian Bureau of Statistics; own calculations.

Figure 2 Gross Domestic Product (GDP) Growth: Comparison with GDP(M2)

Sources: Australian Bureau of Statistics; own calculations.

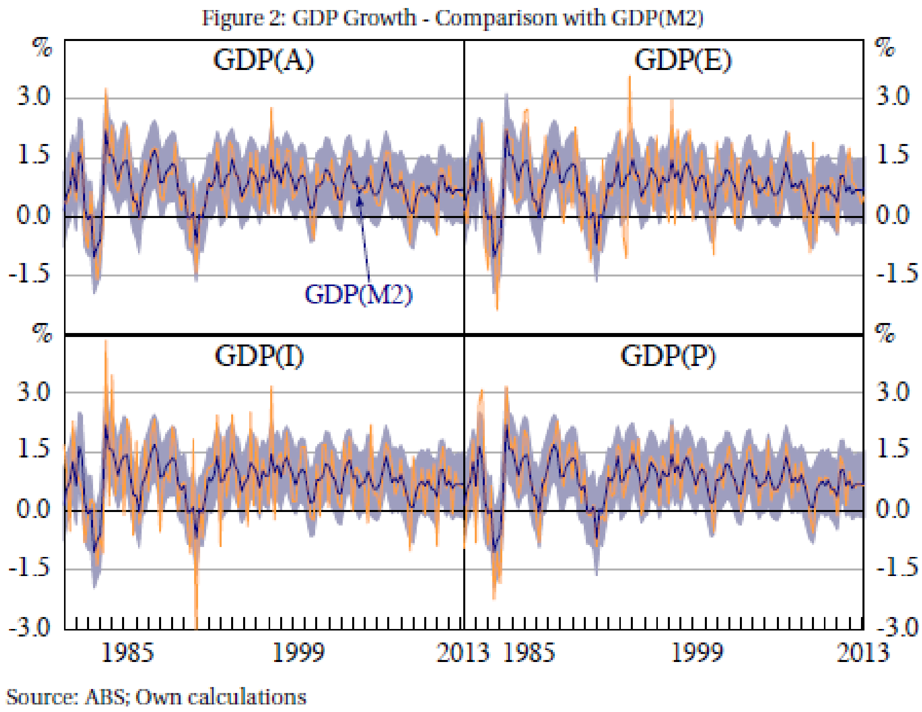

Figure 3 Gross Domestic Product (GDP) Growth: Comparison with GDP(M3)

Sources: Australian Bureau of Statistics; own calculations.

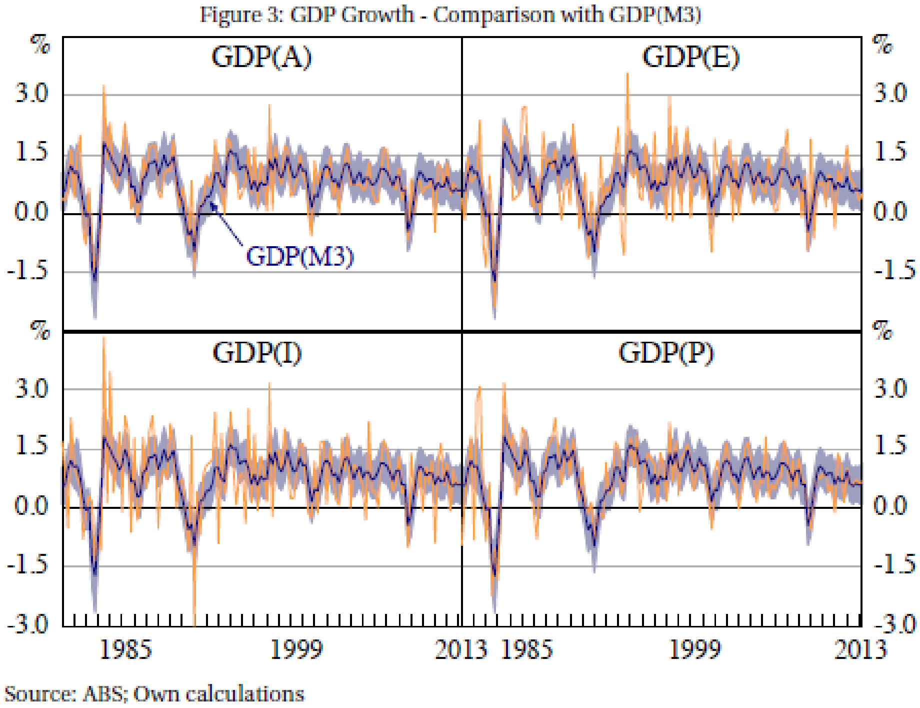

Figure 4 Comparison of Models with Trend Gross Domestic Product (GDP) Growth

Sources: Australian Bureau of Statistics; own calculations.

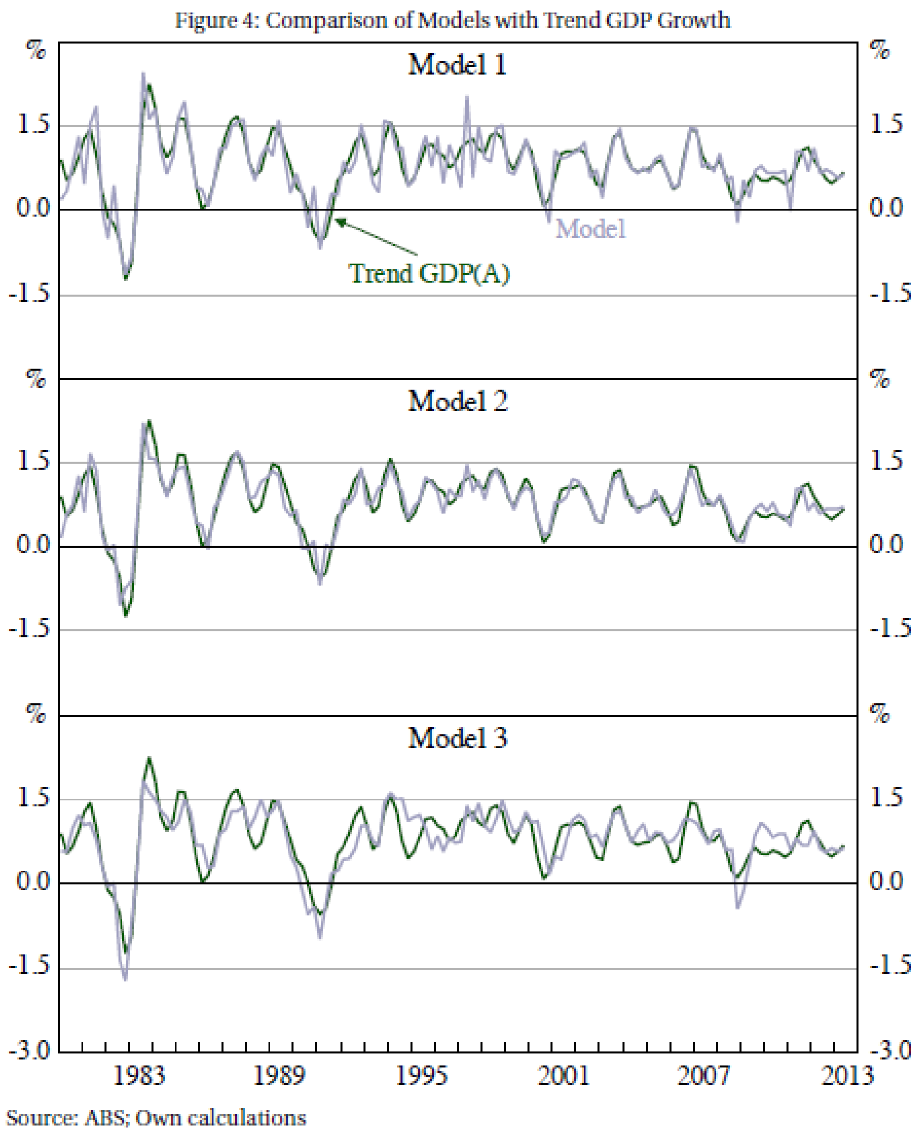

Figure 5 Kalman Gain (KG) Pairs

Source: Own calculations.

Figure 5: Kalman Gain Pairs

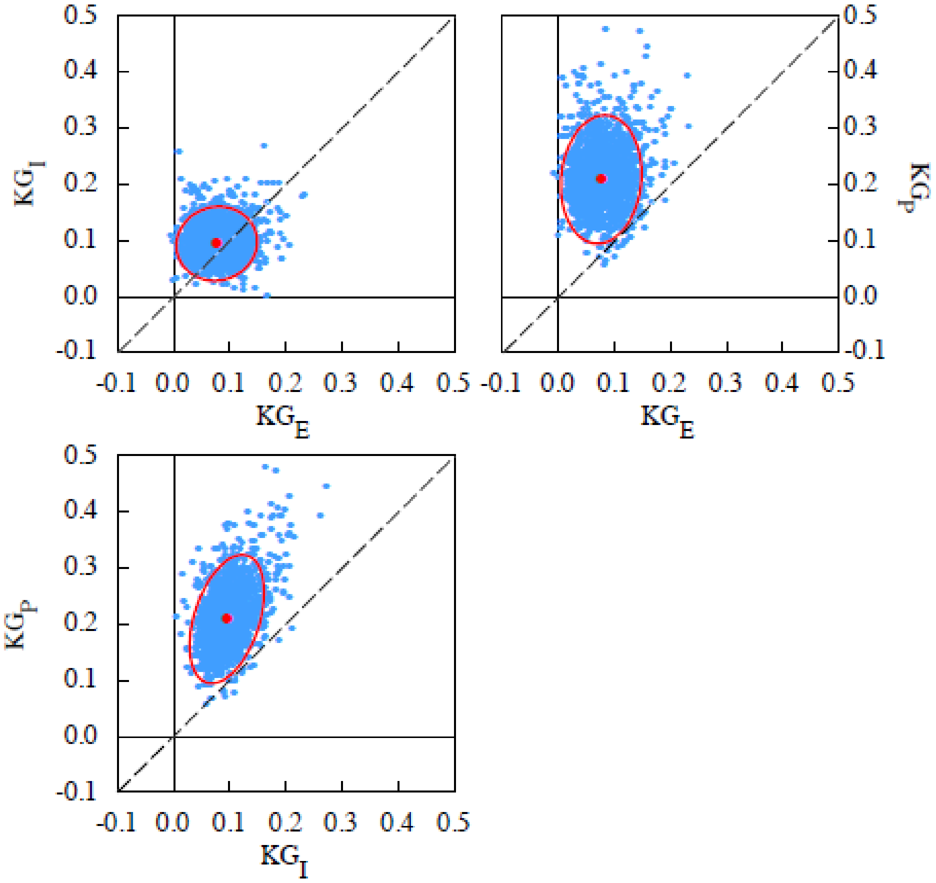

Source: Own calculations

Figure 6 Gross Domestic Product (GDP) Growth: December Quarter 2000

Sources: Australian Bureau of Statistics; own calculations.

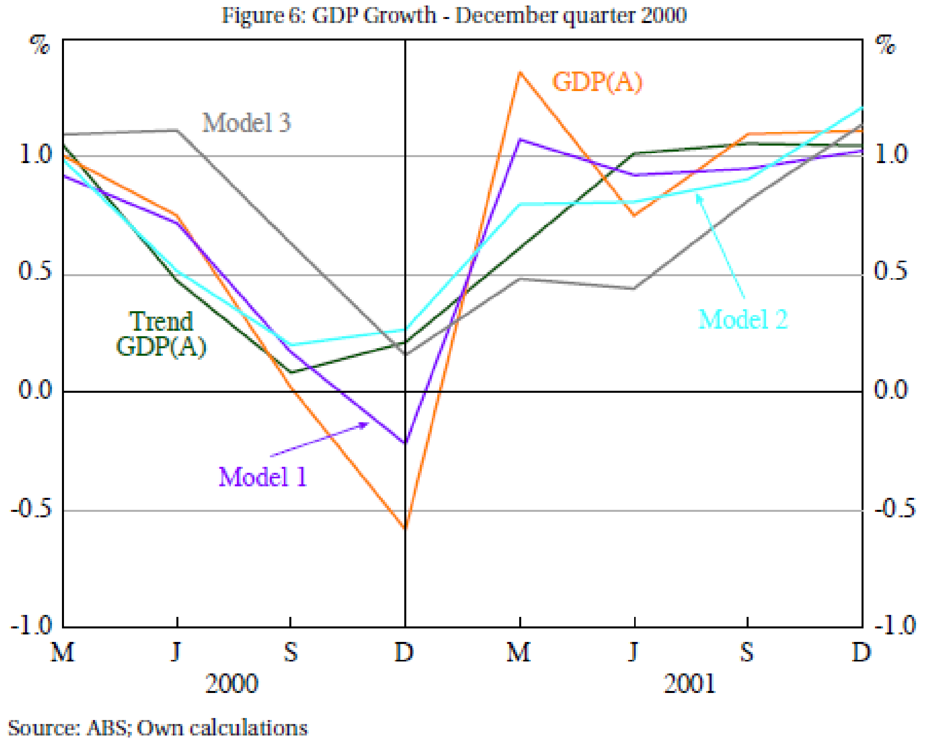

Figure 7 Gross Domestic Product (GDP) Growth: December Quarter 2008

Sources: Australian Bureau of Statistics; own calculations.

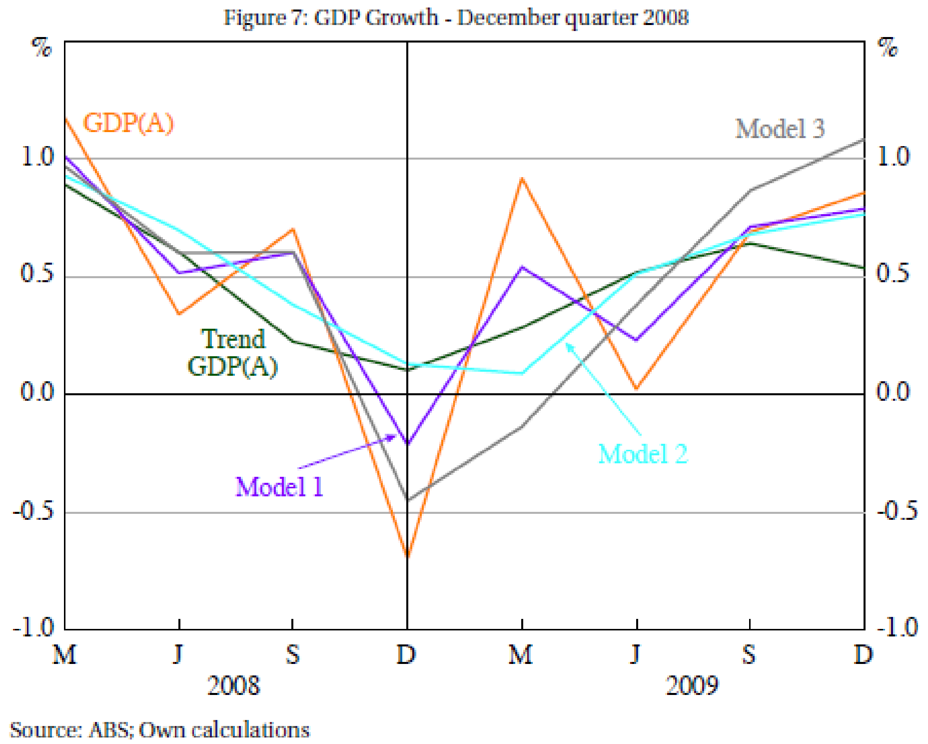

Figure 8 Okun's Law

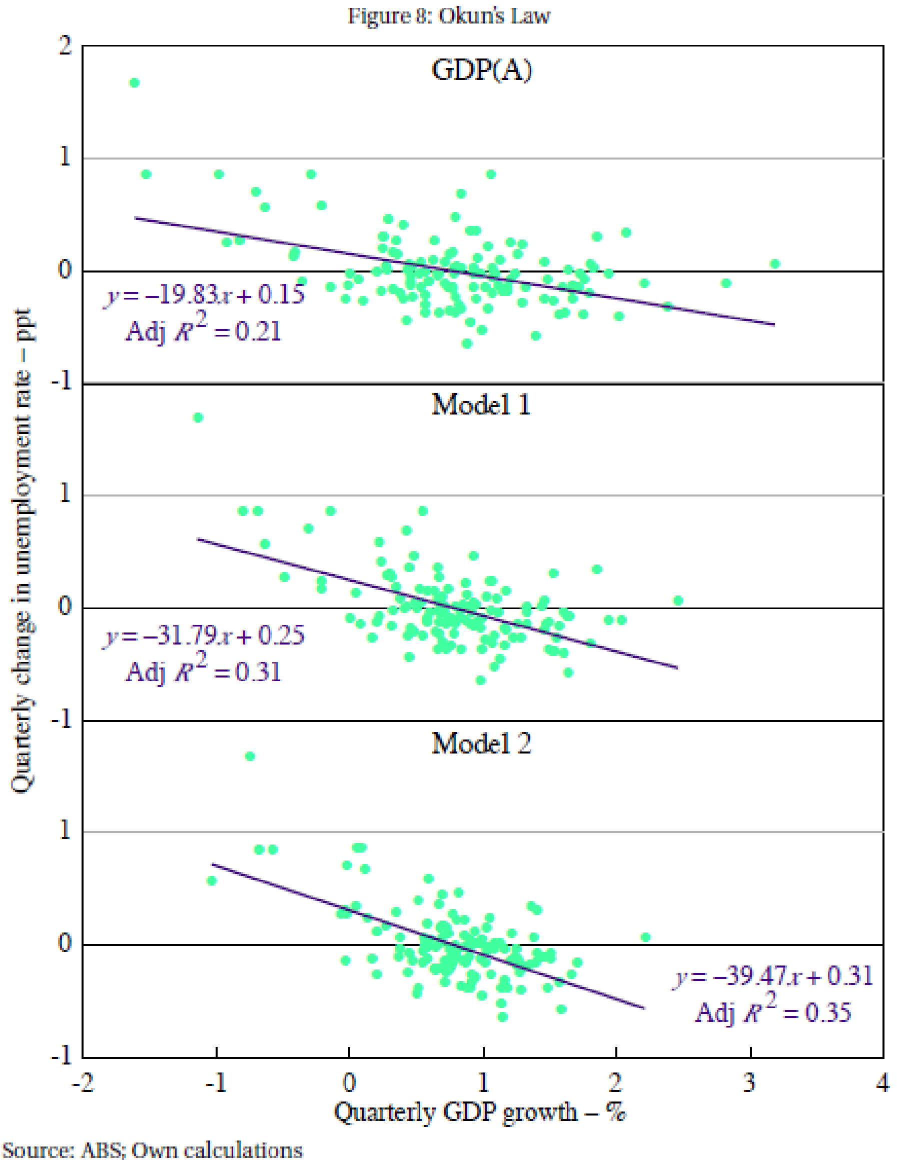

Note: GDP denotes gross domestic product and ppt denotes percentage points.

Source: Own calculations.

Table 1 Prior and Posterior Distributions: Model 1

Parameter	*Prior*			*Posterior*			
	Distribution	*Mean*	*Standard*	*Mode*	*Median*	*5% CI*[a]	*95% CI*

			deviation				
GDP[b] equation							
μ	Normal	0.80	10.0	0.79	0.78	0.62	0.95
ρ	Beta	0.50	0.2	0.37	0.40	0.24	0.56
Exogenous processes							
σ_G^2	Inverse gamma	2.00	4.0	0.38	0.39	0.29	0.52
σ_E^2	Inverse gamma	2.00	4.0	0.43	0.44	0.34	0.57
σ_I^2	Inverse gamma	2.00	4.0	0.68	0.71	0.56	0.89
σ_P^2	Inverse gamma	2.00	4.0	0.29	0.31	0.24	0.42
Marginal data density	516.61						

Notes: (a) CI denotes confidence interval.

(b) GDP denotes gross domestic product.

Table 2 Prior and Posterior Distributions: Model 2

Parameter	*Prior*			*Posterior*			
	Distribution	*Mean*	*Standard deviation*	*Mode*	*Median*	*5% CI*[a]	*95% CI*
GDP[b] equation							
μ	Normal	0.80	10.0	0.79	0.79	0.63	0.94
ρ	Beta	0.50	0.2	0.47	0.46	0.26	0.67
Exogenous							

processes							
σ_G^2	Inverse gamma	2.00	4.0	0.39	0.43	0.31	0.56
σ_{GE}	Normal	0.00	10.0	–0.25	–0.27	–0.48	–0.09
σ_{GI}	Normal	0.00	10.0	–0.17	–0.18	–0.46	0.10
σ_{GP}	Normal	0.00	10.0	–0.13	–0.16	–0.35	–0.02
σ_E^2	Inverse gamma	2.00	4.0	0.84	0.91	0.56	1.35
σ_I^2	Inverse gamma	2.00	4.0	0.92	0.98	0.58	1.46
σ_P^2	Inverse gamma	2.00	4.0	0.45	0.52	0.29	0.88
σ_{EI}	Normal	0.00	10.0	0.37	0.39	0.10	0.74
σ_{EP}	Normal	0.00	10.0	0.31	0.36	0.11	0.66
σ_{IP}	Normal	0.00	10.0	0.17	0.20	–0.07	0.52
Marginal data density	514.53						

Notes: (a) CI denotes confidence interval.

(b) GDP denotes gross domestic product.

Table 3 Prior and Posterior Distributions: Model 3

Parameter	*Prior*			*Posterior*			
	Distribution	*Mean*	*Standard deviation*	*Mode*	*Median*	*5% CI*[a]	*95% CI*
GDP[b] equation							
μ	Normal	0.80	10.0	0.79	0.79	0.59	1.01
ρ	Beta	0.50	0.2	0.65	0.62	0.46	0.77

GDP equation							
κ	Normal	0.00	10.0	0.54	0.51	0.38	0.82
λ	Normal	–0.50	10.0	–0.69	–0.64	–1.03	–0.50
Exogenous processes							
σ_G^2	Inverse gamma	2.00	4.0	0.24	0.28	0.19	0.46
σ_{GE}	Normal	0.00	10.0	0.00	–0.01	–0.11	0.05
σ_{GI}	Normal	0.00	10.0	–0.03	–0.04	–0.14	0.03
σ_{GP}	Normal	0.00	10.0	–0.04	–0.05	–0.14	0.01
σ_E^2	Inverse gamma	2.00	4.0	0.63	0.70	0.54	0.88
σ_I^2	Inverse gamma	2.00	4.0	0.80	0.87	0.69	1.10
σ_P^2	Inverse gamma	2.00	4.0	0.42	0.47	0.36	0.60
σ_U^2	Inverse gamma	0.30	4.0	0.08	0.08	0.05	0.21
σ_{EI}	Normal	0.00	10.0	0.25	0.23	0.10	0.40
σ_{EP}	Normal	0.00	10.0	0.19	0.22	0.11	0.34
σ_{IP}	Normal	0.00	10.0	0.09	0.11	0.00	0.24
σ_{GU}	Normal	0.00	10.0	0.12	0.12	0.07	0.21
Marginal data density	510.43						

Notes: (a) CI denotes confidence interval.

(b) GDP denotes gross domestic product.

Table 4 Descriptive Statistics[a]

	Australian Bureau of Statistics series				*GDP(M) series*		
	GDP(A)	*GDP(E)*	*GDP(I)*	*GDP(P)*	*Model1*	*Model2*	*Model3*
Moments							
Mean	0.79	0.80	0.78	0.80	0.79	0.79	0.79
σ[b]	0.76	0.93	1.03	0.84	0.57	0.56	0.49
ρ_1[c]	0.21	–0.03	–0.19	0.31	0.47	0.79	0.68
Results from an AR(1)[d] regression							
RSE[e]	0.74	0.93	1.02	0.80	0.51	0.34	0.36
R^2-value	0.04	0.00	0.04	0.10	0.22	0.63	0.47

Notes: (a) The sample period is 1980Q1–2013Q2. Model-based statistics are for the posterior median estimate of true gross domestic product (GDP). (b) σ denotes standard deviation. (c) ρ_1 denotes the first-order correlation coefficient. (d) AR denotes autoregressive. (e) RSE denotes residual standard error from a fitted AR(1) model.

Table 5 Revisions to Gross Domestic Product (GDP)

(percentage points)

Measure	*GDP(A)*	*Model1*	*Model2*	*Model3*
Mean absolute revision	0.29	0.20	0.19	0.18
Mean revision	0.13	0.11	0.10	0.05

Notes: The sample period is 2001Q1–2009Q3. Revisions are calculated as the difference between each measure's growth estimate after 4 years and its initial growth estimate.

Table 6 Error Relative to 'Final' Estimate[a]

(mean absolute error, percentage points)

Measure	*GDP(A)*[b]	*Model1*	*Model2*	*Model3*
Initial	0.29	0.20	0.19	0.18
1 year	0.25	0.17	0.16	0.11
2 years	0.22	0.15	0.12	0.08

3years	0.17	0.10	0.09	0.06

Notes: (a) The sample period is 2001Q1–2009Q3. Errors are calculated as the difference between each measure's growth estimate after 4 years and its growth estimate at the specified horizon. (b) GDP denotes gross domestic product.

Table 7 Unemployment Rate: Okun's Law[a]

Parameter	*GDP(A)*[b]	*Model1*	*Model2*
α	0.23**[c]	0.32**	0.34**
ΔU_{t-1}	0.36**	0.29**	0.25**
ΔY_t	–0.12**	–0.17**	–0.11*
ΔY_{t-1}	–0.13**	–0.14**	–0.27**
ΔY_{t-2}	–0.04	–0.08*	–0.05
Implied Okun's coefficient	–0.45	–0.56	–0.59
Adjusted R^2-value	0.54	0.57	0.60

Notes: (a) The sample period is 1980Q4–2013Q2. The models were estimated using robust (White 1980) standard errors.
(b) GDP denotes gross domestic product.
(c) * and ** represent significance at the 5 and 1 per cent level, respectively.

Endnotes

1. GDP(E) is calculated as the sum of all expenditure by resident households, businesses and governments on final production, plus exports and the change in inventories, less imports. GDP(I) measures the income received for providing labour and capital services as inputs to production, adjusted for indirect taxes and subsidies. GDP(P) measures the value of production in the economy as the difference between the value of outputs and the value of intermediate inputs consumed in production. For more detail on the data construction methods, see ABS (2007, 2011, 2012). The Australian Bureau of Statistics (ABS) is one of only a few statistical agencies in the world to compile and publish all three measures of GDP.
2. See, for example, Bishop, Gill and Lancaster (2013) for a recent discussion of measurement error associated with the various GDP estimates.
3. In many other countries, a single measure of GDP is typically used.
4. See, for example, Timmermann (2006) for an overview of the literature or Laplace (1818) for an early application of model averaging.
5. In unpublished work using Australian national accounts data, Scutella (1996) also explored the possibility of extracting underlying economic growth from the noisy expenditure, income and production measures.
6. The model is unidentified in the sense that, with an unrestricted Σ, different model parameters can give rise to identical distributions for the observable quantities.
7. Appendix A of the working paper version of this article (Rees, Lancaster and Finlay 2014) contains a proof that the model is identified with this parametric restriction.
8. We experimented with alternative values of greater than or equal to between 0.5 and 1.1 and with applying the restriction to GDP(P) instead: all produced very similar results.

9. See An and Schorfheide (2007) for a description of these techniques.
10. Our estimation procedure assumes that the trend growth rate of GDP has been constant over our sample. To test whether this assumption is reasonable, we ran Bai–Perron tests for a break in the mean growth rate of GDP(A), using an AR(1) model over the sample 1980Q1–2013Q2. These tests did not point to any evidence of a break in the mean growth rate of GDP(A) over our sample.
11. Imposing a flat prior with a mean of zero produces almost identical results.
12. We constructed the quarterly unemployment rate as the average of the unemployment rates for each month in a quarter.
13. Table 4 contains descriptive statistics for all of the estimated GDP series.
14. Figures 1–3 of the working paper version of this article (Rees, Lancaster and Finlay 2014) also include 95 per cent credible intervals of our estimates of quarterly GDP growth.
15. Of course, seasonally adjusted series may also feature end-point problems if there are changes in seasonal patterns over time.
16. The distributions for the other models are similar.
17. See Bishop, Gill and Lancaster (2013) for a discussion of the revisions process.
18. Due to the time required for estimation, we re-estimate the model every four quarters using real-time data. We use these parameter estimates, combined with real-time national accounts data, to produce estimates for the subsequent three quarters.
19. The working paper version of this article (Rees, Lancaster and Finlay 2014)also examines the relationship between our statistical measures of GDP and inflation.
20. We exclude Model 3 because it is identified using the unemployment rate.

Conclusion

My research while studying for a PhD has resulted in three published papers, which together form the body of this thesis; in addition, I completed two papers while studying for a Master of Economics, which are included in an appendix to this thesis.

Reflecting my interests as well as those of my employer, the Reserve Bank of Australia, four of the five papers concern current, policy-relevant topics of interest to Australian economic policy-makers.

The first paper examines housing wealth effects using the HILDA Survey. We find that young homeowners respond more to higher house prices than do older homeowners, and the young renters also have a positive consumption response to higher house prices, albeit less so than that for young owners. This suggests that housing wealth effects as observed in Australia arise from an easing of collateral constrains and a common association between house prices and another factor such as income expectations, rather than through a 'traditional wealth effects' channel.

The second paper examines the drivers of the rise in Australian household saving over the 2000s using the Household Expenditure Survey of the ABS. We find that highly educated households, as well as those with high wealth and/or debt, increased their propensity to save the most over this period. Given the widely acknowledged link between education and lifetime income, we interpret these results as suggesting that higher saving was driven by a reduction in permanent income expectations following the GFC, as well as a desire to pay down debt and rebuild assets.

The third paper constructs new estimates of Australian GDP growth from the published national accounts estimates of expenditure, income and production. Across a range of specifications, our measures are substantially less volatile than headline GDP growth, while also having superior real-time properties and roughly equal utility in forecasting models. We conclude that much of the quarter-to-quarter volatility in Australian GDP growth reflects measurement error rather than true shifts in the level of economic activity, with our smoother measure potentially useful for policy-makers looking to abstract from quarter-to-quarter noise.

The fourth paper, included in an appendix, examines the effect of credit supply shocks on key macroeconomic variables during the GFC via a sign-restricted VAR model. We find that negative credit-supply shocks explain one-third to one-half of the fall in credit growth seen over the GFC, and around one-sixth of the fall in GDP growth. This suggests that credit supply shocks played an important but not dominant role in economic outcomes over the period, with other identified shocks having a larger impact overall.

Finally, the fifth paper, included in the appendix, represents a departure from the policy-focused empirical studies above, and is instead a theoretical work in the field of econometrics. The paper constructs, via carefully chosen sums of i.i.d. random variables, a new class of highly flexible (in terms of both permissible marginal distribution and permissible correlation structure) non-Gaussian random field, for possible use in economic and/or financial modelling. Although of little immediate policy relevance, the paper extends our basic understanding of stochastic processes, which are the building blocks of economic and financial models, and in so doing makes a contribution to the board economic literature.

Appendix A – Publication A.I

First published in Finlay, R. and Jaaskela, J. (2014). Credit supply shocks and the global financial crisis in three small open economies. *Journal of Macroeconomics,* **40**, 270-276.

Credit supply shocks and the global financial crisis in three small open economies

Richard Finlay[*,a,b], Jarkko P. Jääskelä[a]

[a] *Reserve Bank of Australia, 65 Martin Place, NSW 2000, Australia*
[b] *School of Economics, University of New South Wales, NSW 2052, Australia*

Abstract

We investigate the impact of domestic and foreign credit supply shocks on a number of key macroeconomic variables for three small open economies: Australia, Canada and the UK. We find that negative domestic and foreign credit supply shocks together explain, on average, one-third to one-half of the fall in business credit and rise in spreads seen in the three countries during the financial crisis; other identified non-credit-supply shocks explain the rest. Credit supply shocks also explain around one-sixth of the fall in output in the three countries, and one-quarter of the fall initially seen in UK inflation. This suggests that credit supply shocks played an important role in the financial crisis, but not a dominant one.

Key words: Sign-restricted VAR, Credit supply, Small open economy

JEL: E32, E51

1. Introduction

Business credit growth depends on the demand for and the supply of credit, both of which interact with the macroeconomy. During periods of high growth, businesses tend to increase their use of credit, while in the wake of economic

*Corresponding author. Tel.: +61 2 9551 8887; fax: +61 2 9551 8833
Email addresses: finlayr@rba.gov.au (Richard Finlay), jaaskelaj@rba.gov.au (Jarkko P. Jääskelä)

downturns borrowers can become conservative and cut back their demand for credit; lenders, too, can become more cautious, charging higher spreads or rationing credit (Figure 1). There exist episodes in history where the supply and demand drivers are reasonably clear. Financial deregulation, which occurred in many advanced economies during the 1980s, is an example where a supply constraint was lifted, resulting in strong credit growth. The merger and acquisition boom of the early 2000s is an example of demand-led growth. For the most part, however, econometric identification schemes are needed to disentangle the two factors.

Figure 1: Business Credit and Real GDP
Year-ended growth

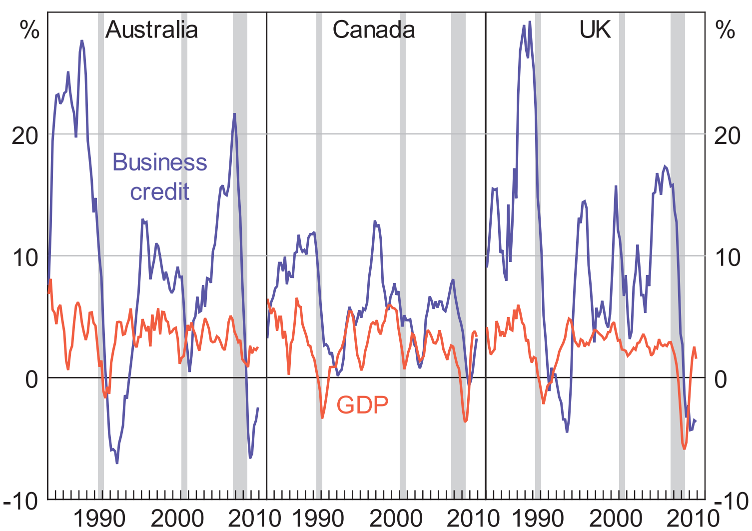

Note: shaded areas refer to US recessions as dated by the NBER.

A number of authors have analyzed the importance of credit shocks (see for example Helbling et al., 2011 and Gilchrist et al., 2009 for two widely cited papers). Our analysis is closely related to this growing literature but distinct in its approach. In particular, we use a sign-restricted VAR with domestic and foreign blocks to identify credit supply (and other) shocks in three small open economies (Australia, Canada and the UK), and their impact on credit, credit spreads, GDP and inflation, among other key variables. Three other closely related papers, Helbling et al. (2011), Meeks (2012) and Fornari and Stracca (2013), also use sign restrictions to identify credit shocks, but they differ from us in a few important dimensions. First, we identify a rich set of macroeconomic shocks based on a simple argument regarding the effect that various demand, supply and other shocks will have on observed quantities and prices (Helbling et al. identify only credit supply and productivity shocks; Meeks identifies only a credit shock; Fornari and Stracca identify aggregate demand, financial and monetary policy shocks). Our identification strategy allows us to gauge the importance of not just domestic and foreign credit supply shocks, but also domestic and foreign credit demand shocks, as well as other standard shocks that are left largely unidentified in the other papers. Moreover, we use the quantity of credit and credit spreads to identify different types of credit shocks. Helbling et al. also use the quantity of credit, along with credit spreads and default rates, to identify credit shocks, whereas Meeks uses only the (US) credit spread and default rate; Fornari and Stracca follow an alternative strategy by identifying a financial shock that has an impact on the quantity of credit and the relative share price of the financial sector. Finally, we focus on small open economies that are affected by exogenous foreign (US) shocks. Meeks focuses only on the US, while Helbling et al. and Fornari and Stracca consider a larger class of countries that includes both small open economies and large, relatively closed economies without distinguishing between them.[1]

[1]The small open economy assumption is important. For example we find that foreign shocks explain around 50 per cent of output variation in Australia, in line with the estimate of Dungey and Pagan (2000, Table 4); for Canada we find a figure of 67 per cent, a little higher than the 60 per cent found by Bhuiyan (2012, Table 4); while for the UK we find a figure of 55 per cent, in line with the estimate of Spencer and Liu (2010, Figure 6b).

We find that, on average, negative domestic and foreign credit supply shocks together explain one-third to one-half of the fall in business credit and rise in spreads seen in the three countries during the financial crisis, with identified non-credit-supply shocks (i.e., shocks that would be likely to affect credit demand) explaining the rest. Credit supply shocks explain around one-sixth of the fall in output in the three countries, and one-quarter of the fall initially seen in UK inflation. This suggests that credit supply shocks played an important role in the financial crisis, but not a dominant one.

Although not the focus of our paper, we note that credit supply and credit demand shocks together explain around one-tenth of US output variation over our sample, of which two-thirds is due to credit supply shocks. This is in line with Helbling et al. who find that credit shocks explain around one-tenth of variation in US and global GDP, as well as Fornari and Stracca who find that financial shocks explain around one-tenth of the variation in their pooled GDP variable, but less than Meeks who finds that credit shocks explain around one-fifth of the variation in US GDP. Regarding just the recession of 2007-2009, we find that credit shocks explain around 20 per cent of the fall in US output; this sits between the estimates of Helbling et al. (around 10 per cent) and Meeks (three-fifths).

2. Model and data

To extract the various shocks underlying movements in the data we estimate the following sign-restricted VAR:

$$\begin{bmatrix} w_t \\ d_t \end{bmatrix} = \alpha x_t + \sum_{i=1}^{p} A_i \begin{bmatrix} w_{t-i} \\ d_{t-i} \end{bmatrix} + A_0 \begin{bmatrix} \epsilon_t^w \\ \epsilon_t^d \end{bmatrix} \tag{1}$$

where w_t and d_t are vectors of endogenous variables, x_t is a vector of exogenous variables, and the matrix A_0 is the contemporaneous impact matrix of the vectors of mutually uncorrelated disturbances. There are eleven variables in the model and they can be divided into two groups. The first five variables $w_t = (y_t^w, \pi_t^w, r_t^w, sp_t^w, cr_t^w)'$ capture the world economy, proxied by the US: y_t^w is quarterly US non-farm GDP growth; π_t^w is quarterly US core inflation; r_t^w is the quarter-average Fed Funds rate; sp_t^w is the quarter-average BAA corporate bond spread to US treasuries; and cr_t^w is quarterly growth in US business credit. With the addition of q_t as the quarterly change in the real domestic/USD exchange rate, the second group of variables $d_t = (y_t^d, \pi_t^d, r_t^d, sp_t^d, cr_t^d, q_t)'$ are similarly defined but related to the domestic economy, being either Australia, Canada or the UK.[2]

Note that with the exception of Australia there is a slight mismatch between the price of credit series that we use, which are based on bond spreads, and the quantity of credit series, which refer to credit extended by banks. This is standard for cross-country studies in the literature – see for example Helbling et al. (2011) – as broader measures of business credit that include non-intermediated debt are generally not available over long sample periods outside of the US. When measuring credit spreads, the literature has tended to favour bond market measures rather than bank measures, since bond market data are both more likely to reflect market developments quickly and again are generally more readily available over long sample periods (Australia is a notable exception, with the price of bank credit to businesses, but not bond spreads, available over a long sample period).

[2]Due to data limitations some definitions are different between countries: GDP (farm and non-farm) is used for Canada and the UK; RPIX inflation is used for the UK; and the spread between large business variable rates and 3-month money-market rates is used for Australia. See data appendix for more details on the data series used.

We identify structural shocks by placing restrictions on the direction that variables in d_t and w_t move in response to different shocks. VAR models identified using this technique are known as sign-restricted VARs; they have been popularized by Faust (1998), Canova and De Nicoló (2002), Peersman (2005) and Uhlig (2005), among others. Six shocks are identified:

- The first three shocks – demand, supply and monetary policy shocks – are standard. An aggregate demand shock moves inflation, policy rates and output in the same direction; a monetary policy shock moves policy rates in the opposite direction to output and inflation; while an aggregate supply shock moves inflation and output in opposite directions (the effect on policy rates is left unrestricted). Each of these aggregate shocks would typically shift the quantity of credit, and spreads, but the direction of the responses is not specified *a priori*.

- Credit shocks are assumed not to have an immediate impact on aggregate macroeconomic variables as they originate in the financial sector and take time to filter through to product markets. We identify a credit *supply* shock as one that leads to an opposite movement between spreads and credit; a credit *demand* shock moves spreads and credit in the same direction. Both types of credit shock are assumed to impact output in the following period.

- In the domestic block a real exchange rate shock is identified. It has no contemporaneous impact on any of the other variables.

- Domestic shocks do not affect the foreign variables at any horizon, while the response of the domestic variables to the foreign shocks are left unrestricted. This is the small open economy assumption.

Note that we impose zero restrictions on impact for the credit and exchange rate shocks only. This assumption is important as it guarantees that all shocks are uniquely identified and therefore that the model does not suffer from the 'multiple shocks' problem discussed in Fry and Pagan (2011). If one chose *not* to impose these restrictions, shocks would not be uniquely identified, in which case interpretation of results becomes somewhat problematic.

The sign restrictions are given in Table 1. The restrictions are imposed for two periods following a shock (except for the exchange rate shock and zero restrictions for the initial impact of the credit shocks, which are on impact only).[3]

Table 1: Restrictions on impact

Shock :	y^w	π^w	r^w	sp^w	cr^w	y^d	π^d	r^d	sp^d	cr^d	q
W demand	↑	↑	↑	-	-	-	-	-	-	-	-
W supply	↑	↓	-	-	-	-	-	-	-	-	-
W monetary policy	↓	↓	↑	-	-	-	-	-	-	-	-
W credit supply	0,↑	0	0	↓	↑	-	-	-	-	-	-
W credit demand	0,↑	0	0	↑	↑	-	-	-	-	-	-
D demand	**0**	**0**	**0**	**0**	**0**	↑	↑	↑	-	-	-
D supply	**0**	**0**	**0**	**0**	**0**	↑	↓	-	-	-	-
D monetary policy	**0**	**0**	**0**	**0**	**0**	↓	↓	↑	-	-	-
D credit supply	**0**	**0**	**0**	**0**	**0**	0,↑	0	0	↓	↑	-
D credit demand	**0**	**0**	**0**	**0**	**0**	0,↑	0	0	↑	↑	-
D exchange rate	**0**	**0**	**0**	**0**	**0**	0	0	0	0	0	↑

Note: ↑ (↓) positive (negative) response of the variables in columns to shocks in rows. **0** no response (the small open economy assumption). 0,↑ no contemporaneous response followed by positive response. - no restriction imposed on the response.

The sample period runs from the March quarter 1984 to the December quarter 2010, and is selected to include an earlier period of weak credit growth and recession in the early 1990s, as shown in Figure 1. In order to capture the introduction of inflation targeting and the structural decline in inflation and interest

[3]Note also that we place parametric restrictions on $A_i, i = 0, 1, ..., p$ such that domestic shocks (ϵ_t^d) and domestic variables (d_t) cannot influence world variables (w_t) at any horizon, maintaining the small open economy assumption.

rates that occurred in the early 1990s in many advanced economies, a constant and dummy variable are included in the x_t matrix. The dummy variable is equal to 1 from 1993 onwards, and 0 otherwise.

We estimate the VAR using Bayesian techniques, with the prior and posterior distributions of the reduced-form VAR being Normal-Wishart. To construct impulse response functions we first draw a random realization of the variance-covariance matrix from Equation (1), $\Sigma = A_0 A_0'$, from the posterior distribution. We then recover a candidate A_0 matrix by calculating the Choleski factor of Σ and multiplying this by an orthogonal Givens rotation matrix that maintains the small open economy assumption of domestic shocks not affecting foreign variables, where the rotation angles are drawn from the uniform distribution (for any orthogonal matrix Q such that $QQ' = I$, if $\tilde{A}_0 \tilde{A}_0' = \Sigma$ then $(\tilde{A}_0 Q)(\tilde{A}_0 Q)' = \tilde{A}_0 Q Q' \tilde{A}_0' = \Sigma$ also, so that if $\tilde{A}_0$ is a candidate A_0 matrix then so is $\tilde{A}_0 Q$). Finally, we compute the implied impulse response; if the response satisfies the restrictions outlined in Table 1 it is kept, otherwise it is discarded. Our estimation algorithm (although not our model) follows Section 2.3 of Peersman (2005) and the appendix of Peersman and Straub (2009) closely, to which we refer the reader for further details (see also Appendix B of Uhlig, 2005, for discussion and details on the Normal-Wishart distribution in the sign-restricted VAR context).

The system is estimated in growth rates (with the exception of interest rate variables which appear in levels), given the lack of evidence for cointegration. Working with growth rates and log levels are both common approaches in the literature: for example Peersman (2005), Gilchrist et al. (2009) and Helbling et al. (2011) express their data in growth rates, while Uhlig (2005), Meeks (2012) and Fornari and Stracca (2013) express their data in log levels. Tests suggest that all variables used are stationary, and the lag length $p = 2$ was selected using likelihood ratio tests.

3. Results

Figures 2 and 3 show selected impulse response functions from the models, based on 1,000 accepted draws for each. For Australia, Canada and the UK, Figure 2 shows the impulse responses of credit growth, credit spreads, output growth and inflation to a *domestic* credit supply shock that increases quarterly business credit growth by 1 percentage point. The solid blue lines plot the median impulse responses; the dashed blue lines represent the 16th and 84th percentiles of the responses; and the red line shows the 'Median Target' measure.[4][5] Results are broadly similar across countries: the shocks to business credit and credit spreads are persistent, with credit growth remaining elevated and spreads falling 25-50 basis points on impact and remaining depressed for a number of periods; there is a small and short-lived output response for Australia and Canada and a slightly more persistent response for the UK; and there is a minimal inflation response, except in Australia where a small but persistent response is present. Mirroring this inflation response, for Australia there is also a policy rate tightening which reaches roughly 50 basis points after six quarters before falling towards zero (not shown).[6]

Figure 3 shows the response of the same domestic variables to a *foreign* credit supply shock that increases foreign business credit growth by 1 percentage

[4]Following Sims and Zha (1999), 16-84 percentile bands have become convention in the sign-restricted VAR literature; they correspond to a band one standard deviation wide.

[5]Fry and Pagan (2011) criticize the practice of using the median response as a measure of central tendency because it mixes the responses of different candidate models. They suggest selecting a single model and hence a unique set of impulse responses from the Monte Carlo draws that is closest to the set of median responses.

[6]Although not the focus of this paper, we note that impulse responses for the US generally accord with those reported above: a positive US credit supply shock (which corresponds to a foreign credit supply shock for the small open economies) leads to a persistent increase in credit growth and fall in credit spreads, a small and short lived positive output response, and a small and mildly persistent increase in policy rates.

Figure 2: Impulse response of domestic variables to a positive domestic credit supply shock

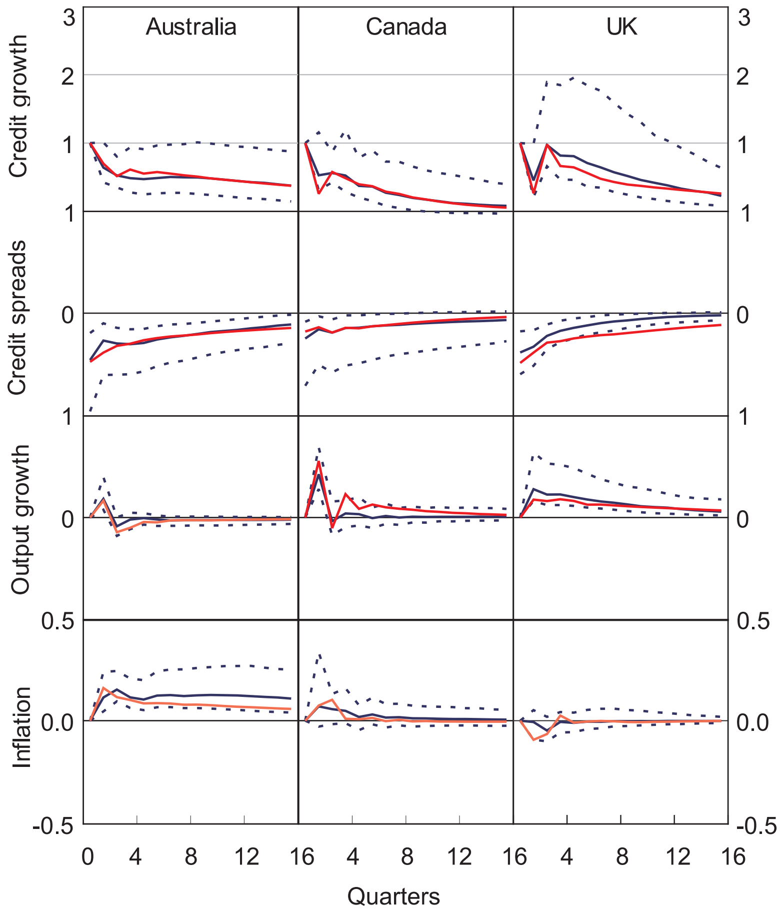

Note: solid blue lines plot median impulse response; dashed blue lines plot 16th and 84th percentiles; red lines plot median target.

point. The 16 and 84 percentile responses generally cover zero, although there are positive credit growth responses in all countries, peaking around the 1 year horizon at 50–150 basis points; initial negative credit spread responses of 25–75 basis points; and a positive inflation response in the UK (and a matching Bank rate response, not shown), indicating that expansions of foreign (US) credit are inflationary for the UK economy. This is likely to reflect the large and

Figure 3: Impulse response of domestic variables to a positive foreign credit supply shock

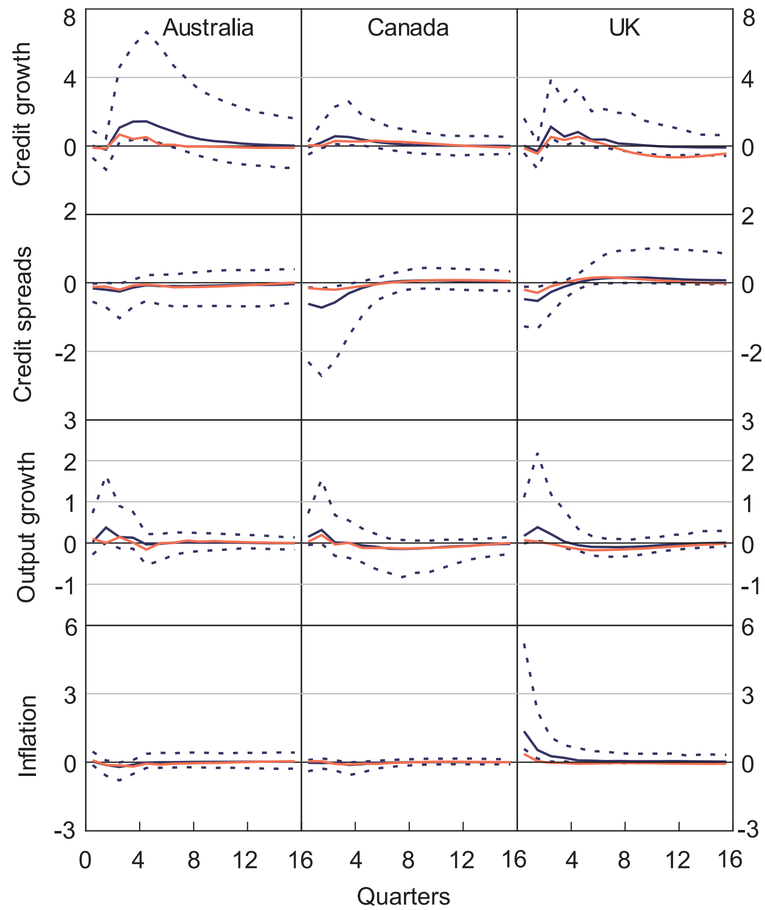

Note: solid blue lines plot median impulse response; dashed blue lines plot 16th and 84th percentiles; red lines plot median target.

internationally connected nature of the UK financial sector.

The graphs in Figure 4 show the median deviation from trend of credit growth, credit spreads, output growth and inflation caused by the identified domestic and foreign credit supply shocks in the three countries (top row of each block of graphs), and the median deviation from trend caused by all other identified shocks (bottom row of each block of graphs), all since 2000, where the shaded periods correspond to US recessions as dated by the NBER.

Figure 4: Historical decompositions

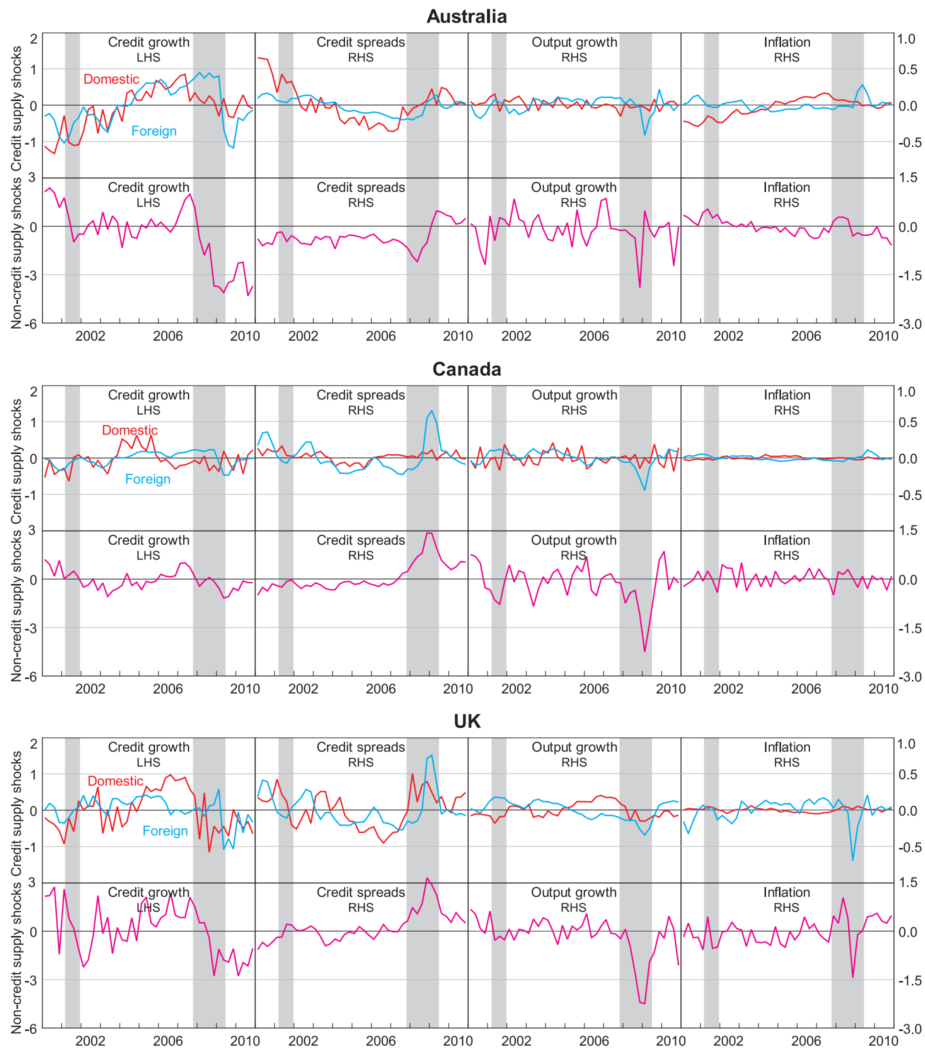

Note: lines plot median deviation from trend caused by identified shocks; shaded areas refer to US recessions as dated by the NBER.

Considering credit supply shocks first, the results suggest that positive credit supply shocks boosted credit growth in the years leading up to the financial crisis, before negative credit supply shocks drove credit growth sharply lower (although in Canada the boost came earlier, and the fall during the crisis was less sharp). Foreign credit supply shocks dominated in the case of Australia, while foreign and domestic shocks were of broadly similar magnitude for Canada and the UK around the time of the crisis. The reverse is the case with spreads – positive credit supply shocks first depressed spreads before negative shocks sent them sharply higher as the crisis commenced, with foreign shocks relatively more important Canada and the UK compared with Australia. The results suggest that there was a marginal positive impact on output from credit supply shocks in the years before the financial crisis, and that negative credit supply shocks (mostly foreign) contributed to the sharp decline in output growth seen during the crisis. Negative foreign credit supply shocks drove UK inflation sharply lower during the crisis, but had a much more muted effect on inflation in Australia and Canada. This is unsurprising given the UK's larger and more globally integrated financial system – negative credit supply shocks, in particular foreign, and the associated contraction in financial sector activity, are likely to have affected the UK economy more forcefully and more directly than Australia or Canada. The real exchange rate was also driven 5–10 per cent lower by negative foreign credit supply shocks in all three countries in late 2008 (not shown).

Considering all other identified shocks aggregated together, the results suggest that credit growth was again pushed lower, and spreads higher, in the wake of the financial crisis in all countries, with shock magnitudes one to two times larger than those caused by credit supply shocks. Within the other identified shocks, foreign aggregate demand shocks were generally the main driver, although foreign aggregate supply shocks also played a role, as did foreign credit demand shocks in Canada and the UK, and domestic credit demand shocks in Australia. Non-credit-supply shocks had a negative impact on output growth roughly five times larger than the impact caused by credit supply shocks, with foreign aggregate demand shocks being the single largest driver in all countries (foreign aggregate supply shocks also played a significant role). For the UK, foreign credit demand shocks were additionally important, while towards the end of our sample negative domestic aggregate supply shocks drove output lower. The inflation responses in Australia and Canada to non-credit-supply shocks were fairly small, while in the UK inflation was initially driven lower by non-credit-supply shocks (mainly foreign aggregate demand shocks), before being pushed higher (by negative domestic aggregate supply shocks).

4. Conclusions

Business credit fell substantially during the financial crisis in a number of advanced economies, while spreads increased and output contracted. Our analysis, based on sign-restricted VAR models estimated for Australia, Canada and the UK, suggests that the fall in business credit and rise in spreads was caused by both a reduction in credit supply originating from foreign and domestic factors and an independent reduction in credit demand (taken as the effect on credit from all non-credit-supply shocks), with the magnitudes of the non-credit-supply shocks one to two times larger than those of the credit supply shocks. Negative credit supply shocks did depress output in the three countries studied, but the effects of non-credit-supply shocks were around five times larger. For inflation the story is more mixed, with credit supply shocks having a relatively small effect on inflation in Australia and Canada, but a large negative impact in the UK. Given the UK's larger and more globally integrated financial sector it is unsurprising that negative credit supply shocks (and the contraction in financial activity that they entail) would have a larger impact there. Non-credit-supply shocks also had an initially more negative, then more positive, impact on inflation in the UK than in Australia or Canada.

Data appendix

GDP growth, credit growth and inflation are seasonally adjusted.

For the US, treasury and BAA corporate bond yields are from H.15 on the website of the Board of Governors of the Federal Reserve System, while business credit is defined as commercial and industrial loans from commercial banks in the US plus commercial real estate loans, available from H.8.

For Australia, spreads are taken as the spread between large business variable rates and 3-month money-market rates from Statistical Tables F1 and F5 on the Reserve Bank of Australia's website. Business credit is available from Statistical Table D1.

For Canada, spreads are calculated as the yield on the BofA Merrill Lynch 10+ Year Canada Corporate Index (code: F9C0), less the yield on the Government of Canada benchmark bond (code: V122544; available from the Bank of Canada's website); prior to 2001 we splice a series that is calculated as the long term corporate bond less long term government bond yield, as was previously published on the Bank of Canada's website at http://www.bankofcanada.ca/wp-content/uploads/2010/09/annual_page14.pdf (this publication has since been removed; data are available from the authors upon request). Total business credit is used (DataStream code CNB169).

For the UK, spreads are taken as the yield on the BofA Merrill Lynch Sterling Corporate Non-Financial Index (URNF) less the yield on the BofA Merrill Lynch 5-10 year UK Gilt Index (G6L0); prior to 1996 we splice equivalent series from Global Financial Data (codes INGBRW and IGGBR10D). UK business credit is taken as monetary financial institutions' sterling net lending excluding securitisations to private non-financial corporations (available on the Bank of

England's website under the code LPQVUGP).

Acknowledgment

We would like to thank James Hansen, Adrian Pagan, editors of the journal William Lastrapes and Douglas McMillin, and two anonymous referees for constructive and useful comments and suggestions. Responsibility for any remaining errors rests with the authors. The views expressed in this paper are those of the authors and are not necessarily those of the Reserve Bank of Australia.

References

Bhuiyan, R., 2012. Monetary transmission mechanisms in a small open economy: a Bayesian structural VAR approach. Canadian Journal of Economics 45. 1037–1061.

Canova, F., De Nicoló, G., 2002. Monetary disturbances matter for business fluctuations in the G-7. Journal of Monetary Economics 49, 1131–1159.

Dungey, M., Pagan, A., 2000. A structural VAR model of the Australian economy. The Economic Record 76, 321–342.

Faust, J., 1998. The robustness of identified VAR conclusions about money. Carnegie-Rochester Conference Series on Public Policy 49, 207–244.

Fornari, F., Stracca, L., 2013. What does a financial shock do? First international evidence. ECB working paper 1522.

Fry, R., Pagan, A., 2011. Sign restrictions in structural vector autoregressions: a critical review. Journal of Economic Literature 49, 938-60.

Gilchrist, S., Yankov, V., Zakrajsek, E., 2009. Credit market shocks and economic fluctuations: evidence from corporate bond and stock markets. Journal of Monetary Economics 56, 471–493.

Helbling, T., Huidrom, R., Kose, M., Otrok C., 2011. Do credit shocks matter? A global perspective. European Economic Review 55, 340–353.

Meeks, R., 2012. Do credit market shocks drive output fluctuations? Evidence from corporate spreads and defaults. Journal of Economic Dynamics & Control 36, 568–584.

Peersman, G., 2005. What caused the early millennium slowdown? Evidence based on vector autoregressions. Journal of Applied Econometrics 20, 185–207.

Peersman, G., Straub, R., 2009. Technology shocks and robust sign restrictions in a euro area SVAR. International Economic Review 50, 727–750.

Sims C., Zha, T., 1999. Error bands for impulse responses. Econometrica 67, 1113–1156.

Spencer P., Liu, Z., 2010. An open-economy macro-finance model of international interdependence: the OECD, US and the UK. Journal of Banking & Finance 34, 667–680.

Uhlig, H., 2005. What are the effects of monetary policy on output? Results from an agnostic identification procedure. Journal of Monetary Economics 52, 381–419.

Appendix A – Publication A.II

First published in Finlay, R. and Seneta, E. (2014). Random fields with Pólya correlation structure. *Journal of Applied Probability*, **51**, 1037-1050.

RANDOM FIELDS WITH PÓLYA CORRELATION STRUCTURE

RICHARD FINLAY * ** AND

EUGENE SENETA,* *** *University of Sydney*

Abstract

We construct random fields with Pólya-type autocorrelation function and dampened Pólya cross-correlation function. The marginal distribution of the random fields may be taken as any infinitely divisible distribution with finite variance, and the random fields are fully characterized in terms of their joint characteristic function. This makes available a new class of non-Gaussian random field with flexible correlation structure for use in modeling and estimation.

Keywords: Random field; Infinitely divisible distribution; Pólya autocorrelation

2010 Mathematics Subject Classification: Primary 60G60

Secondary 60G10;60E07

1. Introduction

Our primary object of study in this paper is a class of random field indexed in the temporal domain over $\mathbb{R}$ and in the spatial domain over $\{1, 2, \ldots d\}$, $d \in \mathbb{N}^+$, that is, a multivariate stochastic process. We denote a random field $\{\mathbf{Z}(t)\}$, $t \in \mathbb{R}$ with spatial dimension defined on $\{1, 2, \ldots, d\}$, $d \in \mathbb{N}^+$ by $\{\mathbf{Z}(t)\} = \{Z_1(t), \ldots, Z_d(t)\}$. If all $\{Z_h(t)\}$ have second-order moments then the covariance matrix function of $\{\mathbf{Z}(t)\}$ is given by $C(t, t+s) = \mathbb{C}\mathrm{ov}(\mathbf{Z}(t), \mathbf{Z}(t+s)) = \mathbb{E}(\mathbf{Z}(t) - \mathbb{E}\mathbf{Z}(t))(\mathbf{Z}(t+s) - \mathbb{E}\mathbf{Z}(t+s))'$. The hth diagonal entry of $C(t, t+s)$ corresponds to the autocovariance (or direct covariance) between $Z_h(t)$ and $Z_h(t+s)$, while the g, hth off-diagonal entry corresponds to the so-called cross-covariance between $Z_g(t)$ and $Z_h(t+s)$, $g \neq h$. Thus the diagonal entries of $C(t, t+s)$ are autocovariance functions and the off-diagonal entries are cross-covariance functions. If both $C(t, t+s)$ and $\mathbb{E}\mathbf{Z}(t)$ are independent of t, then $\{\mathbf{Z}(t)\}$ is said to be second-order stationary.

The class of admissible autocovariance and cross-covariance functions for Gaussian second-order stationary random fields, as well as other closely related elliptically contoured random fields, is well-

* Postal address: School of Mathematics and Statistics F07, University of Sydney, NSW 2006, Australia.

** Email address: richardf@maths.usyd.edu.au

*** Email address: eseneta@maths.usyd.edu.au

known. Here we consider elliptically contoured random fields constructed as Gaussian random fields multiplied by non-negative random variables; the marginal distribution of the resulting random field is altered by the multiplication but the correlation structure is not. In this case the covariance matrix function may be taken as any function that satisfies $C(t, t+s) = C(t+s, t)'$ and

$$\sum_{i=1}^{n}\sum_{j=1}^{n} a_i' C(t_i, t_j) a_j \geq 0 \tag{1}$$

for all $n \in \mathbb{N}^+$, $t_k \in \mathbb{R}$ and $a_k \in \mathbb{R}^d$ for $k = 1, \dots, n$ (see for example Cramér and Leadbetter (1967) and Gikhman and Skorokhod (1969), as well as Ma (2009, 2011a, 2011b, 2011c, 2011d) and the references therein). Du and Ma (2013), for example, construct an elliptically contoured random field that may take any matrix function satisfying Equation (1) as its covariance matrix function.

For random fields that are non-Gaussian, however, Equation (1) is in general a necessary but not a sufficient condition for the covariance structure, and the range of admissible covariance structures must be investigated on a case-by-case basis. For example, for a log-Gaussian random field, Equation (1) is a necessary but not sufficient condition for its covariance structure.

In this article we construct second-order stationary random fields in both continuous and discrete time: Section 3 constructs a random field in continuous time with Pólya-type autocorrelation function and dampened Pólya cross-correlation function; Section 4 constructs a random field in discrete time, the more practically useful setting, with Young-type autocorrelation function and dampened Young cross-correlation function; while Section 5 presents a number of extensions (Pólya- and Young-type autocorrelation functions are defined in Section 2 below). Importantly, the marginal distribution of the random fields may be taken as any infinitely divisible distribution with finite variance, where by marginal distribution we mean the distribution of the random variable $Z_h(t)$ for fixed t. This extends results from Finlay and Seneta (2007) and Finlay, Fung and Seneta (2011) to the multivariate setting, and makes available a new class of non-Gaussian second-order stationary random field with flexible correlation structure for use in modeling and estimation.

Other authors have constructed non-Gaussian random fields. In addition to those papers already cited, Marfè (2012, 2013) for example constructs a multivariate Lévy process that can accommodate a flexible range of linear and non-linear dependencies across the spatial dimension and for which the marginal distribution may approximate any Lévy type. Our construction has a number of advantages over alternatives in the literature, however. The marginal distribution of our random fields may be taken as any infinitely divisible distribution with finite variance, whereas, for example, the marginal distributions of the elliptically contoured random fields discussed above are restricted to be of normal variance-mixing type and so exclude any non-symmetric distribution or any distribution that does not have support on $(-\infty, \infty)$, such as a distribution on the positive half-line. Further, since the elliptically contoured random fields are constructed as Gaussian random fields multiplied by a non-negative random variable, given a realization of that random variable they revert to being Gaussian. Our random fields can also be endowed with a rich and dynamic correlation structure across both the spatial and temporal dimensions. Although endowed with a rich dependence structure along the spatial dimension, the Lévy process constructed by Marfè has independent increments, so the increments lack a dependence structure along the time domain (it is the stationary increments of Marfè's process, rather than the process itself, that is most closely related to the processes that we construct). Finally, our method of construction, based on sums of independent and identically distributed (iid) random variables, lends itself particularly easily to numerical simulation, while the random fields are fully characterized in terms of their joint characteristic function, allowing for efficient estimation.

2. Pólya- and Young-type autocorrelation functions

Pólya (1949) provides a simple sufficient condition for the admissibility of a continuous time autocorrelation function of a univariate Gaussian process, being essentially that a function $\rho(s)$ is admissible if it is real-valued, continuous and symmetric about the origin, with $\rho(0) = 1$, $\rho(s)$ convex for $s > 0$ and $\rho(s) \to 0$ as $s \to \infty$ (see also Lukacs (1960), Theorem 4.3.1, as well as Chung (2001) and Christakos (1984)). In fact the condition was originally stated in terms of characteristic functions, but a function is a real-valued characteristic function if and only if it is also an admissible autocorrelation function (see for example Finlay, Fung and Seneta (2011)).

This Pólya condition is useful in the univariate setting as it is reasonably flexible and importantly is easy to check in practice. There is a more general necessary and sufficient condition, being that $\rho(s)$ satisfy

$$\sum_{i=1}^{n}\sum_{j=1}^{n}\rho(t_i - t_j)a_i\bar{a}_j \geq 0$$

for all $n \in \mathbb{N}^+$, $t_k \in \mathbb{R}$ and $a_k \in \mathbb{C}$ for $k = 1,\ldots,n$ (see for example Feller (1966), Section XIX.3), but its practical use is limited as for a given $\rho(s)$ it can be difficult to check.

A related theorem from Young (1913) gives an analogous result for the discrete time setting. For $\rho(s)$, $s \in \mathbb{N}$, Young's theorem essentially states that $\rho(s)$ is an admissible discrete time autocorrelation function if it is real-valued and symmetric on $\{0 \pm 1, \pm 2, \ldots\}$, with $\rho(0) = 1$, $\rho(s) \to 0$ as $s \to \infty$, and $\rho(s) \geq 0$, $\rho(s+1) - \rho(s) \leq 0$, $\rho(s+2) - 2\rho(s+1) + \rho(s) \geq 0$ for $s = 0, 1, 2, \ldots$ (see also Zygmund (1968), Chapter V, as well as Kolmogoroff (1923)). Similar to the Pólya condition, the result was originally stated in the context of Fourier series, but the Fourier series can be interpreted as a symmetric probability density function on $(-\pi, \pi)$ and, inverting the Fourier series, the $\rho(s)$ for $s \in \mathbb{N}$ (the Fourier coefficients) can be interpreted as the characteristic function of this probability density function evaluated at the integers. Being the characteristic function of a symmetric density function, and so real-valued, $\rho(s)$, $s \in \mathbb{N}$ is also an admissible discrete time autocorrelation function.

These Pólya and Young sufficient conditions turn out to define the set of autocorrelation and cross-correlation functions possible using the method that we employ below; our method essentially involves constructing random fields via carefully chosen sums of iid random variables, and the Pólya (in continuous time) and Young (in discrete time) conditions ensure that all sums that we consider are non-negative.

3. A random field in continuous time

Assumption 1. *$\rho(s)$, $s \in \mathbb{R}$ is a continuous function symmetric about $s = 0$ satisfying $\rho(0) = 1$, $\rho(s) \to 0$ as $s \to \infty$, and for $s \in [0, \infty)$ satisfying $\rho(s) \geq 0$, $\rho'(s) \leq 0$ and $\rho''(s) \geq 0$.*

Note that Assumption 1 implies $\rho'(s) \to 0$ as $s \to \infty$.

Assumption 2. *$\kappa(s)$, $s \in \mathbb{R}$ is a continuous function satisfying $0 \leq \kappa(s) \leq 1$.*

Theorem 1. *If $\rho(s)$ is a function satisfying Assumption 1 and $\kappa(s)$ is a function satisfying Assumption 2, then there exists a second-order stationary random field $\{\mathbf{V}(t)\} = \{V_1(t), \ldots, V_d(t)\}$, $t \in \mathbb{R}$, $d \in \mathbb{N}^+$ such that $\mathbb{C}or(V_h(t),\ V_h(t+s)) = \rho(s)$ and $\mathbb{C}or(V_g(t),\ V_h(t+s)) = \int_0^\infty \int_0^\infty \kappa(s+u+v)\rho''(s+u+v)\mathrm{d}v\mathrm{d}u$ for $s > 0$, $g, h = 1, \ldots, d$, $g \neq h$. The marginal distribution of $V_h(t)$ can be taken as any infinitely divisible distribution with finite variance.*

Corollary 1. *If $\kappa(s) = K$ for $K \in [0, 1]$ a constant, then $\mathbb{C}or(V_g(t),\ V_h(t+s))$ reduces to $K\rho(s)$.*

The rest of this section, and in particular Lemmas 1 to 3 below, constitute the proof of Theorem 1. Lemmas 1 to 3 generalize Lemmas 2 and 3 in Finlay, Fung and Seneta (2011), where the result was proved for the univariate case (see also Finlay and Seneta (2007), where the result was proved in the discrete time univariate case for processes with gamma marginal distribution).

Let D_1 denote a given infinitely divisible distribution with finite variance, and $D_{1/n}$ the distribution of the n iid random variables whose sum has distribution D_1. Fix $n \in \mathbb{N}^+$ and set $Y^n_{i,j,h} \stackrel{\mathcal{D}}{=} D_{1/n}$, $i = 1, \ldots, n$, $j = 0, \pm 1, \pm 2, \ldots$, $h = 0, 1, \ldots, d$ with all the $Y^n_{i,j,h}$ mutually independent, where '$\stackrel{\mathcal{D}}{=}$' denotes equality in distribution. Let $[x]$ denote the integer part, and to simplify notation define $\rho_n(x) = [n\rho(x/n)]$ and $f_n(x) = \rho_n(x) - \rho_n(x + 1)$. For $\kappa(s)$ any continuous function such that $0 \leq \kappa(s) \leq 1$, for each j and for $h = 1, \ldots, d$, define a new set of random variables $\tilde{Y}^n_{i,j,h}$ such that $\tilde{Y}^n_{i,j,h} = Y^n_{i,j,0}$ for $i = f_n(k+1)+1, \ldots, f^\kappa_n(k+1)$, and $\tilde{Y}^n_{i,j,h} = Y^n_{i,j,h}$ for $i = f^\kappa_n(k+1)+1, \ldots, f_n(k)$ for $k = 0, 1, 2 \ldots$ until we have assigned a value to all the $\tilde{Y}^n_{i,j,h}$ for $i = 1, \ldots, n - \rho_n(1)$, where we define $f^\kappa_n(x+1) = f_n(x+1) + [\kappa(x/n)(\rho_n(x) - 2\rho_n(x+1) + \rho_n(x+2))]$. That is, the $\tilde{Y}^n_{i,j,h}$ are constructed such that for each j and h, for i between $\rho_n(k+1) - \rho_n(k+2) + 1$ and $\rho_n(k) - \rho_n(k+1)$ a fraction $\kappa(k/n)$ of the $\tilde{Y}^n_{i,j,h}$ are drawn from the $Y^n_{i,j,0}$ and the remaining fraction $1 - \kappa(k/n)$ are drawn from the $Y^n_{i,j,h}$.

Now using the convention that $\sum_{i=m+1}^{m} x_i = 0$ for any $m \geq 0$, define $V^n_h(t)$ for each $h = 1, \ldots, d$ by

$$V^n_h(t) = \sum_{j=-\infty}^{[nt]} \left(\sum_{k=[nt]-j}^{\infty} \left(\sum_{i=f_n(k+1)+1}^{f_n^\kappa(k+1)} Y^n_{i,j,0} + \sum_{i=f_n^\kappa(k+1)+1}^{f_n(k)} Y^n_{i,j,h} \right) \right) \tag{2}$$
$$= \sum_{j=-\infty}^{[nt]} \left(\sum_{k=[nt]-j}^{\infty} \left(\sum_{i=\rho_n(k+1)-\rho_n(k+2)+1}^{\rho_n(k)-\rho_n(k+1)} \tilde{Y}^n_{i,j,h} \right) \right)$$
$$= \sum_{j=-\infty}^{[nt]} \left(\sum_{i=1}^{\rho_n([nt]-j)-\rho_n([nt]-j+1)} \tilde{Y}^n_{i,j,h} \right). \tag{3}$$

$\{\mathbf{V}^n(t)\}$ is defined so that $V^n_h(t)$ for each h and t has marginal distribution D_1. This follows since $\rho(s) \to 0$ as $s \to \infty$ by Assumption 1, so that $\sum_{j=-\infty}^{[nt]} \rho_n([nt]-j) - \rho_n([nt]-j+1) = \rho_n(0) = n$ of the $\tilde{Y}^n_{i,j,h} \overset{\mathcal{D}}{=} D_{1/n}$ are summed in Equation (3), ensuring that $V^n_h(t) \overset{\mathcal{D}}{=} D_1$. Further, for each h, t and s the number of $\tilde{Y}^n_{i,j,h}$ common to $V^n_h(t)$ and $V^n_h(t+s)$ is such that $\mathbb{C}\mathrm{or}(V^n_h(t),\ V^n_h(t+s)) \to \rho(s)$ as $n \to \infty$, and similarly $\mathbb{C}\mathrm{or}(V^n_g(t),\ V^n_h(t+s)) \to \int_0^\infty \int_0^\infty \kappa(s+u+v)\rho''(s+u+v)\mathrm{d}v\mathrm{d}u$ as $n \to \infty$ for each $g \neq h$, as shown in Lemmas 1 and 2 (correlation between $V^n_g(t)$ and $V^n_h(t+s)$ is created via the $Y^n_{i,j,0}$, which from Equation (2) are common to the $V^n_h(t)$ for each $h = 1, 2, \ldots, d$). Note that although Equations (2) and (3) appear to involve infinite sums, for any fixed n all summands for k greater than some finite number and/or j less than some finite number are zero (for both Equations (2) and (3) at most n summands are non-zero since a total of n of the $Y^n_{i,j,0}$, $Y^n_{i,j,h}$ or $\tilde{Y}^n_{i,j,h}$ are summed, and $\rho_n(x) - \rho_n(x+1) \in \mathbb{N}$ decreases and becomes zero as x becomes large).

Lemma 1. *Under Assumption 1, for any $t \in \mathbb{R}$ and $s \geq 0$, $\mathbb{C}or(V^n_h(t),\ V^n_h(t+s)) \to \rho(s)$ as $n \to \infty$.*

Proof. Using Equation (3), consider any $V^n_h(t)$ and $V^n_h(t+s)$ for $t \in \mathbb{R}$, $s \geq 0$. Then for any $j \leq [nt]$, $V^n_h(t)$ contains the first $\rho_n([nt]-j) - \rho_n([nt]-j+1)$ of the $\tilde{Y}^n_{i,j,h}$, while $V^n_h(t+s)$ contains the first $\rho_n([nt+ns]-j) - \rho_n([nt+ns]-j+1)$ of the same $\tilde{Y}^n_{i,j,h}$. But $s > 0$ so by Assumption 1 $\rho_n([nt+ns]-j) - \rho_n([nt+ns]-j+1) \leq \rho_n([nt]-j) - \rho_n([nt]-j+1)$ for large n, so the number of $\tilde{Y}^n_{i,j,h}$ common to both $V^n_h(t)$ and $V^n_h(t+s)$ is simply $\rho_n([nt+ns]-j) - \rho_n([nt+ns]-j+1)$ (recall that $\rho_n(x)$ is defined as $[n\rho(x/n)]$). For $j > [nt]$, $V^n_h(t)$ contains none of the $\tilde{Y}^n_{i,j,h}$, so the total number of $\tilde{Y}^n_{i,j,h}$ common to both $V^n_h(t)$ and $V^n_h(t+s)$ is

$$\sum_{j=-\infty}^{[nt]} \rho_n([nt+ns]-j) - \rho_n([nt+ns]-j+1) = \rho_n([nt+ns]-[nt]) = [n\rho(([nt+ns]-[nt])/n)]$$

and $\mathbb{C}\mathrm{or}(V^n_h(t),\ V^n_h(t+s)) = [n\rho(([nt+ns]-[nt])/n)]/n \to \rho(s)$ as $n \to \infty$. This last step follows from Assumption 1 (see the conclusion of Lemma 2 of Finlay, Fung and Seneta (2011), p. 260).

Lemma 2. *Under Assumptions 1 and 2, for any time $t \in \mathbb{R}$, $g \neq h$ and $s \geq 0$, $\mathbb{C}or(V_g^n(t),\ V_h^n(t+s)) \to \int_0^\infty \int_0^\infty \kappa(s+u+v)\rho''(s+u+v)\mathrm{d}v\mathrm{d}u$ as $n \to \infty$.*

Proof. Using Equation (2), consider any $V_g^n(t)$ and $V_h^n(t+s)$ for $g \neq h$, $s \geq 0$. Then for any $j \leq [nt]$ and $k \geq [nt+ns]-j$, $V_g^n(t)$ and $V_h^n(t+s)$ both contain $[\kappa(k/n)(\rho_n(k) - 2\rho_n(k+1) + \rho_n(k+2))]$ of the same $Y_{i,j,0}^n$. For $k < [nt+ns]-j$, $V_h^n(t+s)$ contains none of the $Y_{i,j,0}^n$, while for $j > [nt]$, $V_h^g(t)$ contains none of the $Y_{i,j,0}^n$. As such the number of $Y_{i,j,0}^n$ common to both $V_g^n(t)$ and $V_h^n(t+s)$ is $\sum_{j=-\infty}^{[nt]} \sum_{k=[nt+ns]-j}^{\infty} [\kappa(k/n)(\rho_n(k) - 2\rho_n(k+1) + \rho_n(k+2))]$, and, ignoring rounding issues associated with taking the integer part, $\mathbb{C}\mathrm{or}(V_g^n(t),\ V_h^n(t+s))$ is given by

$$
\begin{aligned}
&\frac{1}{n}\sum_{j=-\infty}^{nt}\ \sum_{k=nt+ns-j}^{\infty} \kappa(k/n)\left(\rho_n(k) - 2\rho_n(k+1) + \rho_n(k+2)\right)\\
=&\frac{1}{n^2}\sum_{j=-\infty}^{nt}\ \sum_{k=nt+ns-j}^{\infty} \kappa(k/n)\left(\frac{\rho(k/n) - 2\rho(k/n+1/n) + \rho(k/n+2/n)}{1/n^2}\right)\\
=&\frac{1}{n^2}\sum_{u=0}^{\infty}\sum_{v=0}^{\infty} \kappa(s+u/n+v/n)\left(\frac{\rho(s+u/n+v/n) - 2\rho(s+u/n+v/n+1/n) + \rho(s+u/n+v/n+2/n)}{1/n^2}\right)
\end{aligned}
\tag{4}
$$

where Equation (4) follows by making the change of variable $u = nt - j$ and $v = k - u - ns$. Equation (4) converges to

$$
\mathbb{C}\mathrm{or}(V_g^n(t),\ V_h^n(t+s)) = \int_0^\infty \int_0^\infty \kappa(s+u+v)\rho''(s+u+v)\mathrm{d}v\mathrm{d}u \tag{5}
$$

as $n \to \infty$ since $n^2(\rho(x) - 2\rho(x+1/n) + \rho(x+2/n)) \to \rho''(x)$ as $n \to \infty$. By noting that $\sum_{u=0}^\infty \sum_{v=0}^\infty f(s+u+v) = \sum_{u=0}^\infty (u+1) f(s+u)$ for any function f, one can also show that Equation (4) converges to an expression equivalent to Equation (5) given by

$$
\int_0^\infty u\kappa(s+u)\rho''(s+u)\mathrm{d}u.
$$

Note that if $\kappa(s) = K$ for $K \in [0,1]$ a constant, then Equation (5) reduces to $K\rho(s)$.

Lemma 3. *Under Assumptions 1 and 2 there exists a process $\{\mathbf{V}(t)\}$, $t \in \mathbb{R}$ with finite dimensional distributions (and therefore marginal distribution and correlation structure) as implied by Equations (2) and (3) as $n \to \infty$.*

Proof. First we show that the finite dimensional distributions of $\{\mathbf{V}^n(t)\}$, $t \in \mathbb{R}$, converge and define a proper set of random variables as $n \to \infty$.

Fix $p \in \mathbb{N}^+$ and let $a_{1,1}, \ldots, a_{1,p}, \ldots, a_{d,1}, \ldots, a_{d,p} \in \mathbb{R}$ and $-\infty < s_1 < s_2 < \cdots < s_p$, all in $\mathbb{R}$. To ease notation set $g(t,j) = \rho_n([ns_t] - j) - \rho_n([ns_t] - j + 1)$. Then starting from Equation (3), one can show that $\sum_{h=1}^{d}\sum_{t=1}^{p} a_{h,t}V_h^n(s_t)$ is given by

$$\sum_{h=1}^{d}\left(\sum_{k=1}^{p}\left(\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\left(\sum_{i=1}^{g(p,j)}\left(\left(\sum_{t=k}^{p} a_{h,t}\right)\tilde{Y}_{i,j,h}^{n}\right)\right)\right.\right.$$
$$\left.\left.+\sum_{k=1}^{p-1}\left(\sum_{l=1}^{p-k}\left(\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\left(\sum_{i=g(p-l+1,j)+1}^{g(p-l,j)}\left(\left(\sum_{t=k}^{p-l} a_{h,t}\right)\tilde{Y}_{i,j,h}^{n}\right)\right)\right)\right)\right) \tag{6}$$

where we define $s_0 = -\infty$.

The above expression reorders the summation of the $\tilde{Y}_{i,j,h}^n$ appearing in $\sum_{h=1}^{d}\sum_{t=1}^{p} a_{h,t}V_h^n(s_t)$ so that any $\tilde{Y}_{i,j,h}^n$ appearing more than once in the sum are grouped together. But the $\tilde{Y}_{i,j,h}^n$ are not iid since by construction they are drawn from a set of (common) $Y_{i,j,0}^n$ and (unique) $Y_{i,j,h}^n$. Refining Equation (6) to a grouping of all the $Y_{i,j,h}^n$ results in

$$\sum_{k=1}^{p}\left(\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\left(\sum_{m=-\infty}^{j}\left(\sum_{i=g(p,m-1)+1}^{g^{\kappa}(p,m-1)}\left(\left(\sum_{h=1}^{d}\sum_{t=k}^{p} a_{h,t}\right)Y_{i,j,0}^{n}\right)\right)\right)\right) \tag{7}$$
$$+\sum_{k=1}^{p-1}\left(\sum_{l=1}^{p-k}\left(\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\left(\sum_{m=1}^{[ns_{p-l+1}]-[ns_{p-l}]}\left(\sum_{i=g(p-l+1,j+m-1)+1}^{g^{\kappa}(p-l+1,j+m-1)}\left(\left(\sum_{h=1}^{d}\sum_{t=k}^{p-l} a_{h,t}\right)Y_{i,j,0}^{n}\right)\right)\right)\right)\right)$$
$$+\sum_{h=1}^{d}\left(\sum_{k=1}^{p}\left(\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\left(\sum_{m=-\infty}^{j}\left(\sum_{i=g^{\kappa}(p,m-1)+1}^{g(p,m)}\left(\left(\sum_{t=k}^{p} a_{h,t}\right)Y_{i,j,h}^{n}\right)\right)\right)\right)\right.$$
$$\left.+\sum_{k=1}^{p-1}\left(\sum_{l=1}^{p-k}\left(\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\left(\sum_{m=1}^{[ns_{p-l+1}]-[ns_{p-l}]}\left(\sum_{i=g^{\kappa}(p-l+1,j+m-1)+1}^{g(p-l+1,j+m)}\left(\left(\sum_{t=k}^{p-l} a_{h,t}\right)Y_{i,j,h}^{n}\right)\right)\right)\right)\right)\right)$$

where we define $g^{\kappa}(t,m-1) = g(t,m-1) + [\kappa(([ns_t]-m)/n)(g(t,m)-g(t,m-1))]$, so that $\sum_{i=g(t,m-1)+1}^{g^{\kappa}(t,m-1)}$ is the sum of the $\tilde{Y}_{i,j,h}^n$ between $g(t,m-1)$ and $g(t,m)$ which are drawn from the $Y_{i,j,0}^n$, of which there are $[\kappa(([ns_t]-m)/n)(g(t,m)-g(t,m-1))]$ in total, and $\sum_{i=g^{\kappa}(t,m-1)+1}^{g(t,m)}$ is the sum of the $\tilde{Y}_{i,j,h}^n$ between $g(t,m-1)$ and $g(t,m)$ which are drawn from the $Y_{i,j,h}^n$, of which there are $(g(t,m)-g(t,m-1)) - [\kappa(([ns_t]-m)/n)(g(t,m)-g(t,m-1))]$ in total.

Each $Y_{i,j,h}^n$ is iid $D_{1/n}$ distributed, with characteristic function $\phi_{1/n}^{D}(t)$ say, so the characteristic function of $(\mathbf{V}^n(s_1),\ldots,\mathbf{V}^n(s_p))$, defined as $\mathbb{E}\exp(\mathrm{i}\sum_{h=1}^{d}\sum_{t=1}^{p} a_{h,t}V_h^n(s_t))$, is given by

$$\Phi_p^n(a_{1,1},\ldots,a_{d,p}) = \left(\prod_{k=1}^{p}\left(\phi_{1/n}^D\left(\sum_{h=1}^{d}\sum_{t=k}^{p}a_{h,t}\right)\right)^{\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\sum_{m=-\infty}^{j}g^\kappa(p,m-1)-g(p,m-1)}\right) \tag{8}$$

$$\times\left(\prod_{k=1}^{p-1}\left(\prod_{l=1}^{p-k}\left(\phi_{1/n}^D\left(\sum_{h=1}^{d}\sum_{t=k}^{p-l}a_{h,t}\right)\right)^{\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\sum_{m=1}^{[ns_{p-l+1}]-[ns_{p-l}]}g^\kappa(p-l+1,j+m-1)-g(p-l+1,j+m-1)}\right)\right)$$

$$\times\left(\prod_{h=1}^{d}\left(\prod_{k=1}^{p}\left(\phi_{1/n}^D\left(\sum_{t=k}^{p}a_{h,t}\right)\right)^{\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\sum_{m=-\infty}^{j}g(p,m)-g^\kappa(p,m-1)}\right)\right)$$

$$\times\left(\prod_{h=1}^{d}\left(\prod_{k=1}^{p-1}\left(\prod_{l=1}^{p-k}\left(\phi_{1/n}^D\left(\sum_{t=k}^{p-l}a_{h,t}\right)\right)^{\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\sum_{m=1}^{[ns_{p-l+1}]-[ns_{p-l}]}g(p-l+1,j+m)-g^\kappa(p-l+1,j+m-1)}\right)\right)\right)$$

where we use the convention that $\prod_{i=m+1}^{m}x_i = 1$ for any $m \geq 0$.

As D_1 is infinitely divisible, $\phi_{1/n}^D(t) = (\phi_1^D(t))^{1/n}$. Now $1/n\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\sum_{m=-\infty}^{j}g^\kappa(p,m-1) - g(p,m-1) = 1/n\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\sum_{m=-\infty}^{j}[\kappa(([ns_p]-m)/n)(\rho_n([ns_p]-m)-2\rho_n([ns_p]-m+1)+\rho_n([ns_p]-m+2))] \to \int_{s_{k-1}}^{s_k}\int_{-\infty}^{y}\kappa(s_p-x)\rho''(s_p-x)\mathrm{d}x\mathrm{d}y$, while $1/n\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\sum_{m=1}^{[ns_{p-l+1}]-[ns_{p-l}]}g^\kappa(p-l+1,j+m-1) - g(p-l+1,j+m-1) = 1/n\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\sum_{m=1}^{[ns_{p-l+1}]-[ns_{p-l}]}[\kappa(([ns_{p-l+1}]-j-m)/n)(\rho_n([ns_{p-l+1}]-j-m)-2\rho_n([ns_{p-l+1}]-j-m+1)+\rho_n([ns_{p-l+1}]-j-m+2))] \to \int_{s_{k-1}}^{s_k}\int_0^{s_{p-l+1}-s_{p-l}}\kappa(s_{p-l+1}-x-y)\rho''(s_{p-l+1}-x-y)\mathrm{d}y\mathrm{d}x$. Similarly, $1/n\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\sum_{m=-\infty}^{j}g(p,m)-g^\kappa(p,m-1) \to \int_{s_{k-1}}^{s_k}\int_{-\infty}^{y}(1-\kappa(s_p-x))\rho''(s_p-x)\mathrm{d}x\mathrm{d}y$ and $1/n\sum_{j=[ns_{k-1}]+1}^{[ns_k]}\sum_{m=1}^{[ns_{p-l+1}]-[ns_{p-l}]}g(p-l+1,j+m)-g^\kappa(p-l+1,j+m-1) \to \int_{s_{k-1}}^{s_k}\int_0^{s_{p-l+1}-s_{p-l}}(1-\kappa(s_{p-l+1}-x-y))\rho''(s_{p-l+1}-x-y)\mathrm{d}y\mathrm{d}x$.

Hence $\Phi_p^n(a_{1,1},\ldots,a_{d,p})$ converges to a function $\Phi_p(a_{1,1},\ldots,a_{d,p})$ given by

$$\Phi_p(a_{1,1},\ldots,a_{d,p}) = \left(\prod_{k=1}^{p}\left(\phi_1^D\left(\sum_{h=1}^{d}\sum_{t=k}^{p}a_{h,t}\right)\right)^{\int_{s_{k-1}}^{s_k}\int_{-\infty}^{y}\kappa(s_p-x)\rho''(s_p-x)\mathrm{d}x\mathrm{d}y}\right) \tag{9}$$

$$\times\left(\prod_{k=1}^{p-1}\left(\prod_{l=1}^{p-k}\left(\phi_1^D\left(\sum_{h=1}^{d}\sum_{t=k}^{p-l}a_{h,t}\right)\right)^{\int_{s_{k-1}}^{s_k}\int_0^{s_{p-l+1}-s_{p-l}}\kappa(s_{p-l+1}-x-y)\rho''(s_{p-l+1}-x-y)\mathrm{d}y\mathrm{d}x}\right)\right)$$

$$\times\left(\prod_{h=1}^{d}\left(\prod_{k=1}^{p}\left(\phi_1^D\left(\sum_{t=k}^{p}a_{h,t}\right)\right)^{\int_{s_{k-1}}^{s_k}\int_{-\infty}^{y}(1-\kappa(s_p-x))\rho''(s_p-x)\mathrm{d}x\mathrm{d}y}\right)\right)$$

$$\times\left(\prod_{h=1}^{d}\left(\prod_{k=1}^{p-1}\left(\prod_{l=1}^{p-k}\left(\phi_1^D\left(\sum_{t=k}^{p-l}a_{h,t}\right)\right)^{\int_{s_{k-1}}^{s_k}\int_0^{s_{p-l+1}-s_{p-l}}(1-\kappa(s_{p-l+1}-x-y))\rho''(s_{p-l+1}-x-y)\mathrm{d}y\mathrm{d}x}\right)\right)\right)$$

which is continuous about the origin so long as $\phi_1^D(t)$ is. Weak convergence of the finite dimensional distributions of $\{\mathbf{V}^n(t)\}$ to proper random variables follows from Billingsley (1968), Theorem 7.6. (Noting that for $s_p \geq s_j > s_k$, $\int_{-\infty}^{s_k}\int_{-\infty}^{y}f(s_p-x)\mathrm{d}x\mathrm{d}y = \int_0^\infty\int_{s_p-s_j}^\infty f(s_j-s_k+x+y)\mathrm{d}x\mathrm{d}y$ and

$\int_{-\infty}^{s_k}\int_0^{s_p-s_j} f(s_p-x-y)\mathrm{d}x\mathrm{d}y = \int_0^{\infty}\int_0^{s_p-s_j} f(s_j-s_k+x+y)\mathrm{d}x\mathrm{d}y$, one can also use Equation (9) to verify that $V_h(t) \stackrel{\mathcal{D}}{=} D_1$ and that the desired correlation structure holds.)

Hence the finite dimensional distributions as $n\to\infty$ of $\{\mathbf{V}^n(t)\}$ as defined by Equation (3) are consistent, and so by Kolmogorov's Existence Theorem there exists a random field $\{\mathbf{V}(t)\}$, $t\in\mathbb{R}$ with these same finite dimensional distributions (see for example Khoshnevisan (2002)).

Kolmogorov's Continuity Theorem provides a sufficient condition for $\{\mathbf{V}(t)\}$ to have a modification with almost surely continuous sample paths, being that there exist constants $C>0$, $p>0$ and $\gamma>d$ such that $\mathbb{E}|\mathbf{V}(t)-\mathbf{V}(t+s)|^p \le C|s|^\gamma$ (here $\{\tilde{\mathbf{V}}(t)\}$ is said to be a modification of $\{\mathbf{V}(t)\}$ if $\mathrm{P}(\tilde{\mathbf{V}}(t)=\mathbf{V}(t))=1$ for all t; see for example Khoshnevisan (2002), Øksendal (2003)). Taking $p=2$ and $|\mathbf{x}|=\sqrt{x_1^2+\cdots+x_d^2}$ we have that $\mathbb{E}|\mathbf{V}(t)-\mathbf{V}(t+s)|^2 = \mathbb{E}(V_1(t)-V_1(t+s))^2+\cdots+\mathbb{E}(V_d(t)-V_d(t+s))^2 = d\mathbb{E}(V_1(t)-V_1(t+s))^2 = d\mathbb{V}\mathrm{ar}(V_1(t)-V_1(t+s))$. Now

$$V_1^n(t)-V_1^n(t+s) = \sum_{j=-\infty}^{[nt]}\left(\sum_{i=1}^{\rho_n([nt]-j)-\rho_n([nt]-j+1)}\tilde{Y}^n_{i,j,1}\right) - \sum_{j=-\infty}^{[nt+ns]}\left(\sum_{i=1}^{\rho_n([nt+ns]-j)-\rho_n([nt+ns]-j+1)}\tilde{Y}^n_{i,j,1}\right)$$
$$= \sum_{j=-\infty}^{[nt]}\left(\sum_{i=\rho_n([nt+ns]-j)-\rho_n([nt+ns]-j+1)+1}^{\rho_n([nt]-j)-\rho_n([nt]-j+1)}\tilde{Y}^n_{i,j,1}\right) - \sum_{j=[nt]+1}^{[nt+ns]}\left(\sum_{i=1}^{\rho_n([nt+ns]-j)-\rho_n([nt+ns]-j+1)}\tilde{Y}^n_{i,j,1}\right) \tag{10}$$

and since the $\tilde{Y}^n_{i,j,1}$ are iid, $\mathbb{V}\mathrm{ar}(V_1^n(t)-V_1^n(t+s))$ is given by the number of $\tilde{Y}^n_{i,j,1}$ included in the sums that constitute Equation (10), multiplied by $\mathbb{V}\mathrm{ar}(\tilde{Y}^n_{1,1,1})=\sigma^2/n$, where we define σ^2 as the variance of $V_1^n(t)\stackrel{\mathcal{D}}{=}D_1$. The number of $\tilde{Y}^n_{i,j,1}$ included in Equation (10) is given by $2(\rho_n(0)-\rho_n([nt+ns]-[nt])) \le 2n(1-\rho(s+2/n)+1/n)$, so that $\mathbb{V}\mathrm{ar}(V_1^n(t)-V_1^n(t+s)) \le 2(1-\rho(s+2/n)+1/n)\sigma^2 \to 2(1-\rho(s))\sigma^2$ as $n\to\infty$, and therefore $\mathbb{E}|\mathbf{V}(t)-\mathbf{V}(t+s)|^2 \le 2d\sigma^2(1-\rho(s))$. As such, if there exists a $C>0$ and $\gamma>d$ such that for any s, $1-\rho(s)\le C|s|^\gamma$, or equivalently, if $\rho(s)\ge 1-C|s|^\gamma$, then $\{\mathbf{V}(t)\}$ will have a continuous modification.

4. A random field in discrete time

Assumption 3. *$\rho(s)$, $s\in\mathbb{N}$ is a function symmetric about $s=0$ satisfying $\rho(0)=1$, $\rho(s)\to 0$ as $s\to\infty$, and for $s>0$ satisfying $\rho(s)\ge 0$, $\rho(s+1)-\rho(s)\le 0$, $\rho(s+2)-2\rho(s+1)+\rho(s)\ge 0$. This is the discrete time analogue of Assumption 1.*

Note that Assumption 3 implies $\rho(s+1)-\rho(s)\to 0$ as $s\to\infty$.

Assumption 4. *$\kappa(s)$, $s\in\mathbb{N}$ is such that $0\le\kappa(s)\le 1$.*

Theorem 2. *If $\rho(s)$ is a function satisfying Assumption 3 and $\kappa(s)$ is a function satisfying Assumption 4, then there exists a second-order stationary random field $\{\mathbf{X}(t)\}=\{X_1(t),\ldots,X_d(t)\}$, $t\in\mathbb{N}$, $d\in\mathbb{N}^+$ such that $\mathbb{C}or(X_h(t),\ X_h(t+s))=\rho(s)$ and $\mathbb{C}or(X_g(t),\ X_h(t+s))=\sum_{j=0}^{\infty}\sum_{k=0}^{\infty}\kappa(s+j+$*

$k)(\rho(s+j+k)-2\rho(s+j+k+1)+\rho(s+j+k+2))$ for $s\in\mathbb{N}^+$, $g,h=1,\dots,d$, $g\neq h$. The marginal distribution of $X_h(t)$ can be taken as any infinitely divisible distribution with finite variance.

Corollary 2. *If $\kappa(s)=K$ for $K\in[0,1]$ a constant, then $\mathbb{C}or(X_g(t),\ X_h(t+s))$ reduces to $K\rho(s)$.*

Proof. Redefining $f_n(x)$ as $[n\rho(x)]-[n\rho(x+1)]$ and $f_n^\kappa(x+1)=f_n(x+1)+[\kappa(x)([n\rho(x)]-2[n\rho(x+1)]+[n\rho(x+2)])]$, and defining

$$\begin{aligned} X_h^n(t) &= \sum_{j=-\infty}^{t}\left(\sum_{k=t-j}^{\infty}\left(\sum_{i=f_n(k+1)+1}^{f_n^\kappa(k+1)} Y_{i,j,0}^n+\sum_{i=f_n^\kappa(k+1)+1}^{f_n(k)} Y_{i,j,h}^n\right)\right)\\ &= \sum_{j=-\infty}^{t}\left(\sum_{k=t-j}^{\infty}\left(\sum_{i=[n\rho(k+1)]-[n\rho(k+2)]+1}^{[n\rho(k)]-[n\rho(k+1)]} \tilde{Y}_{i,j,h}^n\right)\right)\\ &= \sum_{j=-\infty}^{t}\Big(\sum_{i=1}^{[n\rho(t-j)]-[n\rho(t-j+1)]} \tilde{Y}_{i,j,h}^n\Big)\overset{\mathcal{D}}{=} D_1 \end{aligned}$$

one can show, via an almost identical argument to that used in Lemmas 1 and 2, that $\mathbb{C}\mathrm{or}(X_h^n(t),\ X_h^n(t+s))\to\rho(s)$ and $\mathbb{C}\mathrm{or}(X_g^n(t),\ X_h^n(t+s))\to\sum_{j=0}^{\infty}\sum_{k=0}^{\infty}\kappa(s+j+k)(\rho(s+j+k)-2\rho(s+j+k+1)+\rho(s+j+k+2))$ as $n\to\infty$. One can further show that $\sum_{h=1}^{d}\sum_{t=1}^{p}a_{h,t}X_h^n(s_t)$ is given by Equation (7) where we replace $[ns_k]$ wherever it appears by s_k, replace $\rho_n(x)$ by $[n\rho(x)]$, replace $\kappa(x/n)$ by $\kappa(x)$, redefine $g(t,j)$ as $g(t,j)=[n\rho(s_t-j)]-[n\rho(s_t-j+1)]$, and redefine $g^\kappa(t,m-1)$ as $g^\kappa(t,m-1)=g(t,m-1)+[\kappa(s_t-m)(g(t,m)-g(t,m-1))]$, where as before $\sum_{i=g(t,m-1)+1}^{g^\kappa(t,m-1)}$ is the sum of the $\tilde{Y}_{i,j,h}^n$ between $g(t,m-1)$ and $g(t,m)$ which are drawn from the $Y_{i,j,0}^n$, of which there are $[\kappa(s_t-m)(g(t,m)-g(t,m-1))]$ in total, and $\sum_{i=g^\kappa(t,m-1)+1}^{g(t,m)}$ is the sum of the $\tilde{Y}_{i,j,h}^n$ between $g(t,m-1)$ and $g(t,m)$ which are drawn from the $Y_{i,j,h}^n$, of which there are $(g(t,m)-g(t,m-1))-[\kappa(s_t-m)(g(t,m)-g(t,m-1))]$ in total. Using these same redefinitions, the characteristic function of $(\mathbf{X}^n(s_1),\dots,\mathbf{X}^n(s_p))$ is given by Equation (8). Equation (8) then converges to an expression similar to Equation (9) as $n\to\infty$, where we now replace $\int_{s_{k-1}}^{s_k}\int_{-\infty}^{y}\kappa(s_p-x)\rho''(s_p-x)\mathrm{d}x\mathrm{d}y$ with $\sum_{j=s_{k-1}+1}^{s_k}\sum_{m=-\infty}^{j}\kappa(s_p-m)(\rho(s_p-m)-2\rho(s_p-m+1)+\rho(s_p-m+2))$, replace $\int_{s_{k-1}}^{s_k}\int_0^{s_{p-l+1}-s_{p-l}}\kappa(s_{p-l+1}-x-y)\rho''(s_{p-l+1}-x-y)\mathrm{d}y\mathrm{d}x$ with $\sum_{j=s_{k-1}+1}^{s_k}\sum_{m=1}^{s_{p-l+1}-s_{p-l}}\kappa(s_{p-l+1}-j-m)(\rho(s_{p-l+1}-j-m)-2\rho(s_{p-l+1}-j-m+1)+\rho(s_{p-l+1}-j-m+2))$, and similarly for the expressions involving $1-\kappa$. Weak convergence of the finite dimensional distributions of $\{\mathbf{X}^n(t)\}$ follows from Billingsley (1968), Theorem 7.6 (see also the second paragraph on p. 30), which in the discrete time case is enough to prove that our process $\{\mathbf{X}^n(t)\}$ converges weakly to the limit process $\{\mathbf{X}(t)\}$.

5. Possible extensions

To keep the exposition as simple as possible we imposed a number of constraints on our construction which can be relaxed, as detailed below.

5.1. Extending the domain

We constructed $\{\mathbf{V}(t)\}$ as a random field where the time dimension was defined on $\mathbb{R}$ and the spatial dimension was defined on $\{1, 2, \ldots d\}$. By considering $\sum_{h=1}^{d}\sum_{t=1}^{p} a_{h,t}V_{\eta_h}^{n}(s_t)$ for $d, p \in \mathbb{N}^+$ and $\eta_h \in \mathbb{N}$ or $\eta_h \in \mathbb{R}$ in Lemma 3, instead of $\sum_{h=1}^{d}\sum_{t=1}^{p} a_{h,t}V_{h}^{n}(s_t)$, one can use the argument put forward in Lemma 3 to show that the finite dimensional distributions of $\{V_{\eta_1}^{n}(t), \ldots, V_{\eta_d}^{n}(t)\}$ are well-defined for any $\eta_h \in \mathbb{N}$ or $\eta_h \in \mathbb{R}$, and therefore that a limit random field with spatial dimension defined on $\mathbb{N}$ or $\mathbb{R}$, instead of just on $\{1, 2, \ldots, d\}$, exists. The characteristic function of the finite dimensional distributions of the new process, $\mathbb{E}\exp(\mathrm{i}\sum_{h=1}^{d}\sum_{t=1}^{p} a_{h,t}V_{\eta_h}^{n}(s_t))$, is unchanged from Equation (9).

5.2. Altering the marginal distribution

In our construction, each $\{V_h(t)\}$, $h = 1, \ldots, d$, has the same marginal distribution. This is not necessary – the marginal distribution of $\{V_h(t)\}$ for each h is determined by the number of $\tilde{Y}_{i,j,h}^{n}$ that are summed in Equations (2) and (3), and this can be varied. For example, by summing in Equation (3) from $i = 1$ to $0.7(\rho_n([nt] - j) - \rho_n([nt] - j + 1))$ for each j for $h = 1$, instead of from $i = 1$ to $\rho_n([nt] - j) - \rho_n([nt] - j + 1)$, $\{V_1(t)\}$ will have marginal distribution $D_{0.7}$, while $\{V_h(t)\}$, $h \neq 1$ will have marginal distribution D_1. Note that we still require that all $\{V_h(t)\}$ belong to the same class of infinitely divisible distribution.

5.3. Allowing $\rho(s) \to \delta > 0$ as $s \to \infty$

In our construction we assumed that $\rho(s) \to 0$ as $s \to \infty$, which ensures that a total of n of the $\tilde{Y}_{i,j,h}^{n}$ are summed in Equations (2) and (3). Let $\rho^*(s)$ satisfy all requirements of Assumption 1 except for $\rho^*(s) \to 0$ as $s \to \infty$, and instead let $\rho^*(s) \to \delta > 0$ as $s \to \infty$. We can construct a random field $\{\mathbf{V}^*(t)\}$ which has marginal distribution D_1 for D_1 any infinitely divisible distribution with finite variance, autocorrelation of $\mathbb{C}\mathrm{or}(V_h^*(t),\ V_h^*(t+s)) = \rho^*(s)$ and cross-correlation of $\mathbb{C}\mathrm{or}(V_g^*(t),\ V_h^*(t+s)) = \int_0^\infty\int_0^\infty \kappa(s+u+v)\rho^{*\prime\prime}(s+u+v)\mathrm{d}v\mathrm{d}u$ as follows. Define $\rho(s) = (\rho^*(s) - \delta)/(1 - \delta)$ so that $\rho(s)$ satisfies all requirements of Assumption 1, including that for $\rho(s) \to 0$ as $s \to \infty$, and construct $\{\mathbf{V}(t)\}$ as per Section 3 except taking $Y_{i,j,h}^{n} \overset{\mathcal{D}}{=} D_{(1-\delta)/n}$ in Equations (2) and (3) instead of $Y_{i,j,h}^{n} \overset{\mathcal{D}}{=} D_{1/n}$, so that $V_h(t) \overset{\mathcal{D}}{=} D_{1-\delta}$ for each h and t. Next define a set of random variables $\{\mathbf{V}^\delta\} = \{V_1^\delta, \ldots, V_d^\delta\}$

such that each $V_h^\delta \stackrel{\mathcal{D}}{=} D_\delta$ and is independent of $\{\mathbf{V}(t)\}$, and define $\{\mathbf{V}^*(t)\} = \{\mathbf{V}(t)\} + \{\mathbf{V}^\delta\}$. Then

$$\begin{aligned}\mathbb{C}\mathrm{or}(V_h^*(t),\ V_h^*(t+s)) &= \mathbb{C}\mathrm{or}(V_h(t)+V_h^\delta,\ V_h(t+s)+V_h^\delta)\\ &= (1-\delta)\mathbb{C}\mathrm{or}(V_h(t),\ V_h(t+s)) + \delta\mathbb{C}\mathrm{or}(V_h^\delta,\ V_h^\delta)\\ &= (1-\delta)\rho(s)+\delta = \rho^*(s)-\delta+\delta = \rho^*(s)\end{aligned}$$

while

$$\begin{aligned}\mathbb{C}\mathrm{or}(V_g^*(t),\ V_h^*(t+s)) &= \mathbb{C}\mathrm{or}(V_g(t)+V_g^\delta,\ V_h(t+s)+V_h^\delta)\\ &= (1-\delta)\mathbb{C}\mathrm{or}(V_g(t),\ V_h(t+s)) + \delta\mathbb{C}\mathrm{or}(V_g^\delta,\ V_h^\delta)\\ &= (1-\delta)\int_0^\infty\int_0^\infty \kappa(s+u+v)\rho''(s+u+v)\mathrm{d}v\mathrm{d}u + \delta\mathbb{C}\mathrm{or}(V_g^\delta,\ V_h^\delta)\\ &= \int_0^\infty\int_0^\infty \kappa(s+u+v)\rho^{*\prime\prime}(s+u+v)\mathrm{d}v\mathrm{d}u + \delta\mathbb{C}\mathrm{or}(V_g^\delta,\ V_h^\delta)\end{aligned}$$

since $\rho''(s) = \rho^{*\prime\prime}(s)/(1-\delta)$. Constructing $\{\mathbf{V}^\delta\}$ such that each V_h^δ is iid yields $\mathbb{C}\mathrm{or}(V_g^*(t),\ V_h^*(t+s)) = \int_0^\infty\int_0^\infty \kappa(s+u+v)\rho^{*\prime\prime}(s+u+v)\mathrm{d}v\mathrm{d}u$, but $\{\mathbf{V}^\delta\}$ may constructed so that $\mathbb{C}\mathrm{or}(V_g^\delta,\ V_h^\delta)$ takes any value between 0 and 1.

5.4. Allowing the cross-correlation function to vary

The cross-correlation between $\{V_g(t)\}$ and $\{V_h(t)\}$ is determined by the degree of overlap between the $\tilde{Y}^n_{i,j,g}$ and $\tilde{Y}^n_{i,j,h}$ in Equation (3). For simplicity, in our construction we choose to have all overlap occur via the $Y^n_{i,j,0}$, and to have the same degree of overlap, and thus the same cross-correlation function, between all g and h. This is not necessary. For example, create additional random variables $Y^n_{i,j,h} \stackrel{\mathcal{D}}{=} D_{1/n}$ for $h = -1, -2, \ldots$, where as before $i = 1, \ldots, n$ and $j = 0, \pm1, \pm2, \ldots$, again with all the $Y^n_{i,j,h}$ mutually independent. Now for each j, for $h = 1, \ldots, d$ and for $c(a)$ any function satisfying Assumption 3, define $\tilde{Y}^n_{i,j,h}$ such that for $k = 0, 1, 2\ldots$ and for $a = 0, 1, 2\ldots$, $\tilde{Y}^n_{i,j,h} = Y^n_{i,j,-(h+a)}$ for $i = e_n^\kappa(x+1,a)+1, \ldots, e_n^\kappa(x+1,a+1)$, and $\tilde{Y}^n_{i,j,h} = Y^n_{i,j,h}$ for $i = f_n^\kappa(k+1)+1, \ldots, f_n(k)$, where for a given j and h we adopt the convention that as k and a increase, if $\tilde{Y}^n_{i,j,h}$ for some i has already been assigned a value then we do not reassign it a new value, and where $e_n^\kappa(x+1,a) = \rho_n(x+1)-\rho_n(x+2)+[(1-c(a))\kappa(x/n)(\rho_n(x)-2\rho_n(x+1)+\rho_n(x+2))]$, $f_n(x) = \rho_n(x)-\rho_n(x+1)$ and $f_n^\kappa(x+1) = f_n(x+1)+[\kappa(x/n)(\rho_n(x)-2\rho_n(x+1)+\rho_n(x+2))]$. That is, the $\tilde{Y}^n_{i,j,h}$ are constructed such that for each j and h, for i between $\rho_n(k+1)-\rho_n(k+2)+1$ and $\rho_n(k)-\rho_n(k+1)$ a fraction $(c(a)-c(a+1))\kappa(k/n)$ of the $\tilde{Y}^n_{i,j,h}$ are drawn from the $Y^n_{i,j,-(h+a)}$ for $a = 0, 1, 2, \ldots$, and the remaining fraction $1-\kappa(k/n)$ are drawn from the $Y^n_{i,j,h}$.

In this case, again adopting the convention that $\sum_{i=m+1}^{m} x_i = 0$ for any $m \geq 0$, we have

$$
\begin{aligned}
V_h^n(t) &= \sum_{j=-\infty}^{[nt]} \left(\sum_{k=[nt]-j}^{\infty} \left(\sum_{a=0}^{\infty} \left(\sum_{i=e_n^{\kappa}(x+1,a)+1}^{e_n^{\kappa}(x+1,a+1)} Y_{i,j,-(h+a)}^n \right) + \sum_{i=f_n^{\kappa}(k+1)+1}^{f_n(k)} Y_{i,j,h}^n \right) \right) \\
&= \sum_{j=-\infty}^{[nt]} \left(\sum_{k=[nt]-j}^{\infty} \left(\sum_{i=\rho_n(k+1)-\rho_n(k+2)+1}^{\rho_n(k)-\rho_n(k+1)} \tilde{Y}_{i,j,h}^n \right) \right) \\
&= \sum_{j=-\infty}^{[nt]} \left(\sum_{i=1}^{\rho_n([nt]-j)-\rho_n([nt]-j+1)} \tilde{Y}_{i,j,h}^n \right).
\end{aligned}
$$

Ignoring rounding issues associated with taking the integer part, this ensures that for $g < h$ the number of $Y_{i,j,-a}^n$, $a = 1, 2, \ldots$ common to both the $\tilde{Y}_{i,j,g}^n$ and the $\tilde{Y}_{i,j,h}^n$ for i between $\rho_n(k+1) - \rho_n(k+2) + 1$ and $\rho_n(k) - \rho_n(k+1)$ is given by

$$
\begin{aligned}
&\sum_{a=h}^{\infty} \min \Big((c(a-h) - c(a-h+1))\kappa(k/n)(\rho_n(k) - 2\rho_n(k+1) + \rho_n(k+2)), \\
&\qquad\qquad (c(a-g) - c(a-g+1))\kappa(k/n)(\rho_n(k) - 2\rho_n(k+1) + \rho_n(k+2)) \Big) \\
&= \sum_{a=h}^{\infty} (c(a-g) - c(a-g+1))\kappa(k/n)(\rho_n(k) - 2\rho_n(k+1) + \rho_n(k+2)) \\
&= c(h-g)\kappa(k/n)(\rho_n(k) - 2\rho_n(k+1) + \rho_n(k+2)).
\end{aligned}
$$

Using an argument similar to that used in Lemma 2, one can show that for any $g \neq h$,

$$
\mathbb{C}\mathrm{or}(V_g^n(t),\ V_h^n(t+s)) \to c(|h-g|) \int_0^{\infty} \int_0^{\infty} \kappa(s+u+v)\rho''(s+u+v)\mathrm{d}v\mathrm{d}u
$$

as $n \to \infty$. That is, the cross-correlation function is now dampened by $c(|h-g|)$ and so falls as g and h move further apart.

The characteristic function of this new process is given by

$$
\begin{aligned}
\Phi_p(a_{1,1},\dots,a_{d,p}) = {} & \left(\prod_{k=1}^{p}\left(\phi_1^D\left(\sum_{h=1}^{d}\sum_{t=k}^{p}a_{h,t}\right)\right)^{c(d-1)\int_{s_{k-1}}^{s_k}\int_{-\infty}^{y}\kappa(s_p-x)\rho''(s_p-x)\mathrm{d}x\mathrm{d}y}\right) \qquad (11)\\
& \times\left(\prod_{j=2}^{d}\left(\prod_{k=1}^{p}\left(\phi_1^D\left(\sum_{h=j}^{d}\sum_{t=k}^{p}a_{h,t}\right)\right)^{-c'(d-j)\int_{s_{k-1}}^{s_k}\int_{-\infty}^{y}\kappa(s_p-x)\rho''(s_p-x)\mathrm{d}x\mathrm{d}y}\right)\right)\\
& \times\left(\prod_{j=2}^{d-1}\left(\prod_{i=2}^{j}\left(\prod_{k=1}^{p}\left(\phi_1^D\left(\sum_{h=i}^{j}\sum_{t=k}^{p}a_{h,t}\right)\right)^{c''(j-i)\int_{s_{k-1}}^{s_k}\int_{-\infty}^{y}\kappa(s_p-x)\rho''(s_p-x)\mathrm{d}x\mathrm{d}y}\right)\right)\right)\\
& \times\left(\prod_{j=1}^{d-1}\left(\prod_{k=1}^{p}\left(\phi_1^D\left(\sum_{h=1}^{j}\sum_{t=k}^{p}a_{h,t}\right)\right)^{-c'(j-1)\int_{s_{k-1}}^{s_k}\int_{-\infty}^{y}\kappa(s_p-x)\rho''(s_p-x)\mathrm{d}x\mathrm{d}y}\right)\right)\\
& \times\left(\prod_{k=1}^{p-1}\left(\prod_{l=1}^{p-k}\left(\phi_1^D\left(\sum_{h=1}^{d}\sum_{t=k}^{p-l}a_{h,t}\right)\right)^{c(d-1)\int_{s_{k-1}}^{s_k}\int_{0}^{s_{p-l+1}-s_{p-l}}\kappa(s_{p-l+1}-x-y)\rho''(s_{p-l+1}-x-y)\mathrm{d}y\mathrm{d}x}\right)\right)\\
& \times\left(\prod_{j=2}^{d}\left(\prod_{k=1}^{p-1}\left(\prod_{l=1}^{p-k}\left(\phi_1^D\left(\sum_{h=j}^{d}\sum_{t=k}^{p-l}a_{h,t}\right)\right)^{-c'(d-j)\int_{s_{k-1}}^{s_k}\int_{0}^{s_{p-l+1}-s_{p-l}}\kappa(s_{p-l+1}-x-y)\rho''(s_{p-l+1}-x-y)\mathrm{d}y\mathrm{d}x}\right)\right)\right)\\
& \times\left(\prod_{j=2}^{d-1}\left(\prod_{i=2}^{j}\left(\prod_{k=1}^{p-1}\left(\prod_{l=1}^{p-k}\left(\phi_1^D\left(\sum_{h=i}^{j}\sum_{t=k}^{p-l}a_{h,t}\right)\right)^{c''(j-i)\int_{s_{k-1}}^{s_k}\int_{0}^{s_{p-l+1}-s_{p-l}}\kappa(s_{p-l+1}-x-y)\rho''(s_{p-l+1}-x-y)\mathrm{d}y\mathrm{d}x}\right)\right)\right)\right)\\
& \times\left(\prod_{j=1}^{d-1}\left(\prod_{k=1}^{p-1}\left(\prod_{l=1}^{p-k}\left(\phi_1^D\left(\sum_{h=1}^{j}\sum_{t=k}^{p-l}a_{h,t}\right)\right)^{-c'(j-1)\int_{s_{k-1}}^{s_k}\int_{0}^{s_{p-l+1}-s_{p-l}}\kappa(s_{p-l+1}-x-y)\rho''(s_{p-l+1}-x-y)\mathrm{d}y\mathrm{d}x}\right)\right)\right)\\
& \times\left(\prod_{h=1}^{d}\left(\prod_{k=1}^{p}\left(\phi_1^D\left(\sum_{t=k}^{p}a_{h,t}\right)\right)^{\int_{s_{k-1}}^{s_k}\int_{-\infty}^{y}(1-\kappa(s_p-x))\rho''(s_p-x)\mathrm{d}x\mathrm{d}y}\right)\right)\\
& \times\left(\prod_{h=1}^{d}\left(\prod_{k=1}^{p-1}\left(\prod_{l=1}^{p-k}\left(\phi_1^D\left(\sum_{t=k}^{p-l}a_{h,t}\right)\right)^{\int_{s_{k-1}}^{s_k}\int_{0}^{s_{p-l+1}-s_{p-l}}(1-\kappa(s_{p-l+1}-x-y))\rho''(s_{p-l+1}-x-y)\mathrm{d}y\mathrm{d}x}\right)\right)\right)
\end{aligned}
$$

where in a slight abuse of notation we define $c'(s) = c(s+1)-c(s)$ and $c''(s) = c(s+2)-2c(s+1)+c(s)$.

To construct a random field with a varying cross-correlation where the spatial dimension is defined on $\mathbb{R}$ instead of $\{1, 2, \dots, d\}$, one can alter the argument used above to consider the spatial dimension in increments of $1/n$, instead of unit increments, and then let $n \to \infty$. This is essentially how Section 3 and Section 4 differ, with time implicitly considered in increments of $1/n$ in Section 3 as opposed to unit increments in Section 4. In this case the characteristic function of $\mathbb{E}\exp(\mathrm{i}\sum_{h=1}^{d}\sum_{t=1}^{p}a_{h,t}V_{\eta_h}^{n}(s_t))$ is as in Equation (11), but replacing $c(d-1)$ with $c(\eta_d-\eta_1)$, $-c'(d-j)$ with $c(\eta_d-\eta_j)-c(\eta_d-\eta_{j-1})$, $c''(j-i)$ with $c(\eta_j-\eta_i)-c(\eta_{j+1}-\eta_i)-c(\eta_j-\eta_{i-1})+c(\eta_{j+1}-\eta_{i-1})$, and $-c'(j-1)$ with $c(\eta_j-\eta_1)-c(\eta_{j+1}-\eta_1)$. Alternative cross-correlation structures are possible, with the only limit being the degree of overlap one can construct between the $\tilde{Y}^n_{i,j,h}$ for varying h.

The constraints outlined above can be relaxed either individually or jointly, and although we have couched this section in terms of the continuous time process $\{\mathbf{V}(t)\}$, similar points hold in the discrete time case for $\{\mathbf{X}(t)\}$ also.

6. Conclusion

We have constructed stationary random fields in discrete and continuous time which can have any desired infinitely divisible marginal distribution with finite variance, any autocorrelation function that is positive and convex, and a wide range of cross-correlation functions. This supplements earlier results on Gaussian and related random fields, and makes available non-Gaussian random fields with rich correlation structures which can be used directly in modeling and estimation.

Acknowledgements

The authors would like to thank the editor and an external referee for helpful comments that improved the paper.

References

[1] Billingsley, P. (1968). *Convergence of Probability Measures.* John Wiley, New York.

[2] Christakos G. (1984). On the problem of permissible covariance and variogram models. *Water Resour. Res.* **20**, 251–265.

[3] Chung, K.L. (2001). *A Course in Probability Theory*, 3rd edn. Academic Press, San Diego.

[4] Cramér, H. and Leadbetter, M.R. (1967). *Stationary and Related Stochastic Processes: Sample Function Properties and Their Applications.* Wiley, New York.

[5] Du, J. and Ma, C. (2013). Vector random fields with compactly supported covariance matrix functions. *J. Statist. Plann. Inference* **143**, 457–467.

[6] Feller, W. (1966). *An Introduction to Probability Theory and Its Applications*, vol. 2. Wiley, New York.

[7] Finlay, R. and Seneta, E. (2007). A gamma activity time process with noninteger parameter and self-similar limit. *J. Appl. Probab.* **44**, 950–959.

[8] Finlay, R., Fung, T. and Seneta, E. (2011). Autocorrelation Functions. *Int. Stat. Rev.* **79**, 255–271.

[9] Gikhman, I.I. and Skorokhod, A.V. (1969). *Introduction to the Theory of Random Processes.* W.B. Saunders Co., Philadelphia.

[10] Khoshnevisan, D. (2002). *Multiparameter processes: an introduction to random fields.* Springer, New York.

[11] Kolmogoroff, A. (1923). Sur l'ordre de grandeur des coefficients de la série de Fourier-Lebesgue. *Bull. Acad. Polon. Sci. Classe A*, 83–86.

[12] Lukacs, E. (1960). *Characteristic Functions*. Charles Griffin and Company, London.

[13] Ma, C. (2009). Construction of non-Gaussian random fields with any given correlation structure. *J. Statist. Plann. Inference* **139**, 780–787.

[14] Ma, C. (2011a). Vector random fields with second-order moments or second-order increments. *Stoch. Anal. Appl.* **29**, 197–215.

[15] Ma, C. (2011b). Covariance matrices for second-order vector random fields in space and time. *IEEE Trans. Signal Process.* **59**, 2160–2168.

[16] Ma, C. (2011c). Covariance matrix functions of χ^2 vector random fields in space and time. *IEEE Trans. Commun.* **59**, 2554–2561.

[17] Ma, C. (2011d). Vector random fields with long-range dependence. *Fractals* **19**, 249–258.

[18] Marfè, R. (2012). A multivariate pure-jump model with multi-factorial dependence structure. *Int. J. Theor. Appl. Finance* **15**.

[19] Marfè, R. (2013). Multivariate Lévy processes with dependent jump intensity. *Quant. Finance* **13**, in press.

[20] Pólya, G. (1949). Remarks on characteristic functions. In: Proceedings of the [First] Berkeley Symposium on Mathematical Statistics and Probability, University of California Press, Berkeley and Los Angeles, 115–123.

[21] Young, W.H. (1913). On the Fourier Series of bounded functions. *Proc. London Math. Soc.* **12**, 41–70.

[22] Zygmund, A. (1968). *Trigonometric Series*, 2nd edn., vol. 1. Cambridge University Press, Cambridge.

[23] Øksendal, B. (2003). *Stochastic differential equations: an introduction with applications*, 6th edn. Springer, New York.

Bibliography

An S, Schorfheide F (2007). Bayesian analysis of DSGE models. *Econometric Reviews*, vol. 26, pp. 113–72.

Aron J, Duca JV, Muellbauer J, Murata K, Murphy A (2012). Credit, Housing Collateral, and Consumption: Evidence from Japan, the U.K., and the U.S. *The Review of Income and Wealth*, 58(3), 397–423.

Aruoba SB, Diebold FX, Nalewaik J, Schorfheide F, Song D (2013). Improving GDP measurement: A measurement-error perspective. National Bureau of Economic Research Working Paper no. 18954, Cambridge, Massachusetts.

Aspden C (1990). Which is the best short-term measure of gross domestic product? A statistical analysis of the short-term movements of the three measures of gross domestic product and their average. in *Australian National Accounts: National Income and Expenditure, June Quarter 1990*, Cat. no. 5206.0, Australian Bureau of Statistics, Canberra.

Attanasio OP (1999). Consumption. In JB Taylor and M Woodford (eds), *Handbook of Macroeconomics*, *Volume 1B*, Handbooks in Economics 15, Elsevier Science, Amsterdam, pp 741–812.

Attanasio OP, Blow L, Hamilton R, Leicester A (2009). Booms and Busts: Consumption, House Prices and Expectations. *Economica*, 76(301), 20–50.

Attanasio OP, Weber G (1994). The UK Consumption Boom of the Late 1980s: Aggregate Implications of Microeconomic Evidence. *The Economic Journal*, 104(427), pp 1269–1302.

Attanasio OP, Weber G (2010). Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy. *Journal of Economic Literature*, 48(3), pp 693–751.

Australian Bureau of Statistics (2007). *Quality Dimensions of the Australian National Accounts*, Cat. no. 5216.0.55.002.

Australian Bureau of Statistics (2008). *Experimental Estimates of Imputed Rent Australia*, Information Paper, ABS Cat No 6525.0.

Australian Bureau of Statistics (2011). *Spotlight on National Accounts: Which Is the Best Short-Term Measure of Gross Domestic Product?, May 2011*, Cat. no. 5202.0.

Australian Bureau of Statistics (2012). *Australian System of National Accounts: Concepts, Sources and Methods, Edition 3*, Cat. no. 5216.0.

Benito A. (2006). Does Job Insecurity Affect Household Consumption. *Oxford Economic Papers*, 58, 157–181.

Berger-Thomson L, Chung E, McKibbin R (2009). Estimating Marginal Propensities to Consume in Australia Using Micro Data. RBA Research Discussion Paper No 2009-07.

Bhuiyan, R (2012). Monetary transmission mechanisms in a small open economy: a Bayesian structural VAR approach. *Canadian Journal of Economics*, 45, pp 1037–1061.

Billingsley P (1968). Convergence of Probability Measures. John Wiley, New York.

Bishop J, Gill T, Lancaster D (2013). GDP revisions: Measurement and implications. *Reserve Bank of Australia Bulletin*, March, pp. 11–22.

Blinder AS (1973). Wage Discrimination: Reduced Form and Structural Estimates. *The Journal of Human Resources*, 8(4), pp 436–455.

Borland J (2011). The Australian labour market in the 2000s: The quiet decade. in *The Australian Economy in the 2000s: Proceedings of a Conference*, eds H. Gerard and J. Kearns, Reserve Bank of Australia, Sydney.

Browning M, Crossley TF (2001). The Life-Cycle Model of Consumption and Saving', *The Journal of Economic Perspectives*, 15(3), pp 3–22.

Browning M, Gortz M, Leth-Petersen S (2013). Housing Wealth and Consumption: A Micro Panel Study. *The Economic Journal*, 123(May), 401–28.

Browning M, Lusardi A (1996). Household Saving: Micro Theories and Micro Facts. *Journal of Economic Literature*, 34(4), pp 1797–1855.

Brzozowski M, Crossley TF (2011). Viewpoint: Measuring the Well-Being of the Poor with Income or Consumption: A Canadian Perspective. *The Canadian Journal of Economics*, 44(1), pp 88–106.

Buiter WH (2010). Housing Wealth isn't Wealth. *Economics: The Open-Access, Open-Assessment E-Journal*, 4(2010-22).

Campbell JY, Cocco JF (2007). How do House Prices Affect Consumption? Evidence from Micro Data. *Journal of Monetary Economics*, 54(3), 591–621.

Canova F, De Nicolo G (2002). Monetary disturbances matter for business fluctuations in the G-7. *Journal of Monetary Economics*, 49, pp 1131–1159.

Carroll CD (1992). The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence. *Brookings Papers on Economic Activity*, 1992(2), pp 61–156.

Carroll CD, Dynan KE, Krane SD (2003). Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets. *The Review of Economics and Statistics*, 85(3), 586–604.

Carroll CD, Samwick AA (1997). The Nature of Precautionary Wealth. *Journal of Monetary Economics*, 40(1), pp 41–71.

Chamon MD, Prasad ES (2010). Why are Saving Rates of Urban Households in China Rising? *American Economic Journal: Macroeconomics*, 2(1), pp 93–130.

Christakos G (1984). On the problem of permissible covariance and variogram models. *Water Resour. Res.*, 20, pp 251–265.

Chung KL (2001). A Course in Probability Theory, 3rd edn. Academic Press, San Diego.

Contreras J, Nichols J (2010). Consumption Responses to Permanent and Transitory Shocks to House Appreciation. Board of Governors of the Federal Reserve System Finance and Economics Discussion Series, No 2010-32.

Cramer H, Leadbetter MR (1967). Stationary and Related Stochastic Processes: Sample Function Properties and Their Applications. Wiley, New York.

Cristini A, Sevilla A (2014). Do House Prices Affect Consumption? A Re-assessment of the Wealth Hypothesis. *Economica*, 81(323).

Deaton AS (1991). Saving and Liquidity Restraints. *Econometrica*, 59(5), pp 1221–1248.

Deaton AS, Paxson CH (1994). Saving, Growth, and Aging in Taiwan. In DA Wise (ed), *Studies in the Economics of Aging,* National Bureau of Economic Research Project Report, The University of Chicago Press, Chicago, pp 331–357.

Disney R, Bridges S, Gathergood J (2010). House Price Shocks and Household Indebtedness in the United Kingdom. *Economica*, 77(307), 472–96.

Du J, Ma C (2013). Vector random fields with compactly supported covariance matrix functions. *J. Statist. Plann. Inference,* 143, pp 457–467.

Duca J, Muellbauer J, Murphy A (2011). House Prices and Credit Constraints: Making Sense of the US Experience. *The Economic Journal*, 121(552), 533–51.

Dungey M, Pagan A (2000). A structural VAR model of the Australian economy. *The Economic Record*, 76, pp 321–342.

Durbin J, Koopman SJ (2012) *Time Series Analysis by State Space Methods*, 2nd edn, Oxford University Press, Oxford.

Dvornak N, Kohler, M (2003). Housing Wealth, Stock Market Wealth and Consumption: A Panel Analysis for Australia. RBA Research Discussion Paper No 2003-07.

Dvornak N, Kohler M (2007). Housing Wealth, Stock Market Wealth and Consumption: A Panel Analysis for Australia. *Economic Record,* 83, 117–130.

Dynan KE, Kohn DL (2007). The Rise in US Household Indebtedness: Causes and Consequences. In C Kent and J Lawson (eds), *The Structure and Resilience of the*

Financial System, Proceedings of a Conference, Reserve Bank of Australia, Sydney, pp 84–113.

Dynan KE, Skinner J, Zeldes SP (2004). Do the Rich Save More? *Journal of Political Economy*, 112(2), pp 397–444.

Faust J (1998). The robustness of identified VAR conclusions about money. Carnegie-Rochester Conference Series on Public Policy, 49, pp 207–244.

Gilchrist S, Yankov V, Zakrajsek E (2009). Credit market shocks and economic fluctuations: evidence from corporate bond and stock markets. *Journal of Monetary Economics*, 56, pp 471–493.

Feller W (1966). An Introduction to Probability Theory and Its Applications, vol. 2. Wiley, New York.

Finlay R, Fung T, Seneta E (2011). Autocorrelation Functions. *Int. Stat. Rev*., 79, pp 255–271.

Finlay R, Jääskelä JP (2014). Credit Supply Shocks and the Global Financial Crisis in Three Small Open Economies. *Journal of Macroeconomics*, 40, pp 270–276.

Finlay R, Seneta E (2007). A gamma activity time process with noninteger parameter and self-similar limit. *J. Appl. Probab*., 44, pp 950–959.

Fixler D, Grimm B (2006). GDP estimates: Rationality tests and turning point performance. *Journal of Productivity Analysis*, vol. 25, pp. 213–29.

Fornari F, Stracca L (2013). What does a financial shock do? First international evidence. ECB working paper 1522.

Fry R, Pagan A (2011). Sign restrictions in structural vector autoregressions: a critical review. *Journal of Economic Literature*, 49, pp 938–960.

Gikhman II, Skorokhod AV (1969). Introduction to the Theory of Random Processes. W.B. Saunders Co., Philadelphia.

Gourinchas PO, Parker JA (2001). The Empirical Evidence of Precautionary Saving. *The American Economic Review*, 91(2), 406–12.

Harris MN, Loundes J, Webster E (2002). Determinants of Household Saving in Australia. *Economic Record*, 78(241), pp 207–223.

Helbling T, Huidrom R, Kose M, Otrok C (2011). Do credit shocks matter? A global perspective. *European Economic Review*, 55, pp 340–353.

Iacoviello M (2004). Consumption, House Prices, and Collateral Constraints: A Structural Econometric Analysis. *Journal of Housing Economics*, 13(4), pp 304–320.

Islam A, Parasnis J, Fausten D (2013). Do Immigrants Save Less than Natives? Immigrant and Native Saving Behaviour in Australia. *Economic Record*, 89(284), pp 52–71.

Khoshnevisan D. (2002). Multiparameter processes: an introduction to random fields. Springer, New York.

Kolmogoroff A. (1923). Sur l'ordre de grandeur des coe_cients de la s_erie de Fourier-Lebesgue. *Bull. Acad. Polon. Sci. Classe A*, pp 83–86.

Laplace PS (1818). *Deuxième Supplément a la Théorie Analytique des Probabilités*, Courcier, Paris.

Lehnert A (2004). Housing, Consumption, and Credit Constraints. Board of Governors of the Federal Reserve System Finance and Economics Discussion Series, No 2004-63.

Lukacs E (1960). Characteristic Functions. Charles Griffin and Company, London.

Ma C. (2009). Construction of non-Gaussian random fields with any given correlation structure. *J. Statist. Plann. Inference*, 139, pp 780–787.

Ma C (2011a). Vector random fields with second-order moments or second-order increments. *Stoch. Anal. Appl*. 29, pp 197–215.

Ma C (2011b). Covariance matrices for second-order vector random fields in space and time. *IEEE Trans. Signal Process*., 59, pp 2160–2168.

Ma C (2011c). Covariance matrix functions of chi-squared vector random fields in space and time. *IEEE Trans. Commun,* 59, pp 2554–2561.

Ma C (2011d). Vector random fields with long-range dependence. *Fractals,* 19, pp 249–258.

Marfe R (2012). A multivariate pure-jump model with multi-factorial dependence structure. *Int. J. Theor. Appl. Finance,* 15.

Marfe R (2013). Multivariate Levy processes with dependent jump intensity. *Quant. Finance,* 13.

Meeks R (2012). Do credit market shocks drive output fluctuations? Evidence from corporate spreads and defaults. *Journal of Economic Dynamics & Control*, 36, pp 568–584.

Meyer BD, Sullivan JX (2011). Viewpoint: Further Results on Measuring the Well-Being of the Poor using Income and Consumption. *The Canadian Journal of Economics*, 44(1), pp 52–87.

Mishkin FS (2007). Housing and the Monetary Transmission Mechanism. NBER Working Paper, No 13518.

Muellbauer J (2007a). Housing, Credit and Consumer Expenditure, in Housing, Housing Finance, and Monetary Policy. A Symposium sponsored by the Federal Reserve Bank of Kansas City, Kansas City, 267–334.

Muellbauer J (2007b). Housing and Personal Wealth in a Global Context. WIDER (World Institute for Development Economics Research) Research Paper No 2007/27.

Muellbauer J (2009). Household Decisions, Credit Markets and the Macroeconomy: Implications for the Design of Central Bank Models. Paper prepared for the Eighth BIS Annual Conference. Financial System and Macroeconomic Resilience: Revisited, Basel, 25–26 June.

Nalewaik J (2010). The income- and expenditure-side measures of output growth. *Brookings Papers on Economic Activity*, vol. 41, pp. 71–127.

Nalewaik J (2011). The income- and expenditure-side estimates of U.S. output growth—An update to 2011Q2. *Brookings Papers on Economic Activity*, vol. 42, pp. 385–402.

Oaxaca R (1973). Male-Female Wage Differentials in Urban Labor Markets. *International Economic Review*, 14(3), pp 693–709.

Peersman G (2005). What caused the early millennium slowdown? Evidence based on vector autoregressions. *Journal of Applied Econometrics,* 20, pp 185–207.

Peersman G, Straub R (2009). Technology shocks and robust sign restrictions in a euro area SVAR. *International Economic Review*, 50, pp 727–750.

Polya G (1949). Remarks on characteristic functions. In: Proceedings of the [First] Berkeley Symposium on Mathematical Statistics and Probability, University of California Press, Berkeley and Los Angeles, pp 115–123.

Rees D, Lancaster D, Finlay R (2014). A state-space approach to Australian GDP measurement. Reserve Bank of Australia Research Discussion Paper no. 2014-12.

Reserve Bank of Australia (2008). Financial Stability Review, September.

Sabelhaus J, Groen JA (2000). Can Permanent-Income Theory Explain Cross-Sectional Consumption Patterns? *The Review of Economics and Statistics,* 82(3), pp 431–438.

Scutella R (1996). Obtaining an efficient estimate of the true level of GDP. Honours thesis, The University of Melbourne.

Sims C, Zha T (1999). Error bands for impulse responses. *Econometrica*, 67, pp 1113–1156.

Skinner J (1989). Housing Wealth and Aggregate Saving. *Regional Science & Urban Economics*, 19(2), 305–24.

Spencer P, Liu Z (2010). An open-economy macro-finance model of international interdependence: the OECD, US and the UK. *Journal of Banking & Finance*, 34, pp 667–680.

Stevens G (2009). The road to recovery. *Reserve Bank of Australia Bulletin*, May, pp. 16–23.

Thaler RH, Shefrin HM (1981). An Economic Theory of Self-Control. *Journal of Political Economy*, 89(2), pp 392–406.

Timmermann A (2006). Forecast combinations. in *Handbook of Economic Forecasting*, eds G. Elliott, C. Granger and A. Timmermann, North-Holland, Amsterdam.

Uhlig H (2005). What are the effects of monetary policy on output? Results from an agnostic identification procedure. *Journal of Monetary Economics,* 52, pp 381–419.

White H (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, vol. 48, pp. 817–38.

Windsor C, Jääskelä J, Finlay R (2013). Home Prices and Household Spending. *RBA Discussion Paper*, 2013-04.

Windsor C, Jääskelä JP, Finlay R (2015). Housing Wealth Effects: Evidence from an Australian Panel. *Economica*, vol. 82, pp 552–577.

Wooldridge JM (2002). Econometric Analysis of Cross Section and Panel Data. MIT Press, Cambridge.

Yates J, Whelan S (2009). Housing Wealth and Consumer Spending. AHURI Final Report No 132, Australian Housing and Urban Research Institute, Melbourne.

Young WH (1913). On the Fourier Series of bounded functions. *Proc. London Math. Soc*., 12, pp 41–70.

Zeldes SP (1989). Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy*, 97(2), pp 305–346.

Zygmund A (1968). Trigonometric Series, 2nd edn., vol. 1. Cambridge University Press, Cambridge.

Øksendal B (2003). Stochastic differential equations: an introduction with applications, 6th edn. Springer, New York.