ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Lettau, Martin; Ludvigson, Sydney; Barczi, Nathan

Working Paper

A primer on the economics and time series econometrics of wealth effects: A comment

Staff Report, No. 131

Provided in Cooperation with:
Federal Reserve Bank of New York

Suggested Citation: Lettau, Martin; Ludvigson, Sydney; Barczi, Nathan (2001) : A primer on the economics and time series econometrics of wealth effects: A comment, Staff Report, No. 131, Federal Reserve Bank of New York, New York, NY

This Version is available at:
https://hdl.handle.net/10419/60545

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

A Primer on the Economics and Time Series Econometrics of Wealth Effects: A Comment *

Martin Lettau, Sydney Ludvigson and Nathan Barczi
Federal Reserve Bank of New York

May 23, 2001

*Address: Research Department, Federal Reserve Bank of New York, 33 Liberty St., New York, New York 10045; e-mail: martin.lettau@ny.frb.org, sydney.ludvigson@ny.frb.org, and nathan.barczi@ny.frb.org. Lettau is also affiliated with the CEPR. We are grateful to Simon Potter for many helpful comments. The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Abstract

In a recent paper ("A Primer on the Economics and Time Series Econometrics of Wealth Effects," 2001), Davis and Palumbo investigate the empirical relation between three cointegrated variables: aggregate consumption, asset wealth, and labor income. Although cointegration implies that an equilibrium relation ties these variables together in the long run, the authors focus on the following structural question about the short-run dynamics: "How quickly does consumption adjust to changes in income and wealth? Is the adjustment rapid, occurring within a quarter, or more sluggish, taking place over many quarters?" The authors claim that their findings answer this question, and imply that spending adjusts only gradually after gains or losses in income or wealth have been realized. We argue here, however, that a statistical methodology different from that used by Davis and Palumbo is required to address these questions, and that once it has been employed, the resulting empirical evidence weighs considerably against their interpretation of the data.

1 Introduction

In a recent paper Davis and Palumbo ("A Primer on the Economics and Time Series Econometrics of Wealth Effects," 2001), investigate the empirical relation between three cointegrated variables: aggregate consumption, asset wealth, and labor income. Although cointegration implies that an equilibrium relation ties these variables together in the long run, the authors focus on the following structural question about the short-run dynamics: "How quickly does consumption adjust to changes in income and wealth? Is the adjustment rapid, occurring within a quarter, or more sluggish, taking place over many quarters?" The authors claim that their findings answer this question, and imply that spending adjusts only gradually after gains or losses in income or wealth have been realized. We argue here, however, that a statistical methodology different from that used by Davis and Palumbo is required to address these questions, and that once it has been employed, the resulting empirical evidence weighs considerably against their interpretation of the data. Nevertheless, their results may be of some interest in identifying the observable indicators that *forecast* quarterly personal consumer expenditure growth, as we briefly discuss at the end of this comment.

Our assessment is based on two broad observations about their empirical approach. First, Davis and Palumbo employ a single equation regression for consumption growth to form inferences about the speed with which consumption adjusts to last period's equilibrium (cointegrating) error. If last period's cointegrating error causes a "correction" to this period's consumption, the reasoning goes, the evidence suggests that spending adjusts sluggishly to movements in wealth or income, and the estimated adjustment, or "error-correction," parameter from this regression reveals how quickly that adjustment takes place. The problem with this reasoning is that it ignores empirical evidence that it is not consumption, but *wealth*, that does most of the error-correction subsequent to a shock that causes consumption, asset wealth, and labor income to deviate from their long-run equilibrium relation. Consequently, their estimates of the adjustment parameter governing the degree to which consumption participates in this error-correction–even if correct–would not by themselves reveal the length of time it takes spending to adapt to fluctuations in income or wealth. The difficulty lies with the single equation approach: to make such inferences about adjustment time, it is

necessary to take into account the dynamic response of *all* of the interrelated variables in the cointegrated system (including wealth and income) to shocks that create a cointegrating error.

Second, as we demonstrate in this comment, the estimates reported in Davis and Palumbo (2001) of the adjustment parameter of interest (what Davis and Palumbo term the *error-correction speed* for consumption) are in fact incorrect, because they are obtained by altering a single equation error-correction representation for consumption to include conditioning variables that are not weakly exogenous for the parameter they seek to estimate. We show how these conditioning variables may be incorporated into the analysis without violating weak exogeneity, so that the adjustment parameter of interest may be recovered. Once this adjustment parameter for consumption is estimated correctly, our results suggest that it is not a large negative number, as Davis and Palumbo report, but is instead small in absolute value, indeed close to zero. These results suggest that consumer spending adjusts to movements in wealth or income not over periods of many quarters, as Davis and Palumbo conclude, but rather within the span of about one quarter.

This comment discusses each of these broad observations in detail and addresses a number of other issues at the close. In addition, we show that if the adjustment parameter for consumption were truly as large as Davis and Palumbo argue it is, the cointegrating error for consumption, asset wealth and labor income should have long-horizon forecasting power for consumption growth, a phenomenon that is absent in aggregate spending data.

2 Single Equation versus System Estimation

Davis and Palumbo begin their paper by focusing on the long run properties of aggregate consumption, c_t, aggregate household net worth, a_t, and aggregate labor earnings, y_t. The data suggest that these variables have a single cointegrating relation, or long-run equilibrium (Ludvigson and Steindel (1999); Lettau and Ludvigson (2001a)). Once this long-run equilibrium has been identified, one can move on to address questions about the short-run dynamics such as those that concern which variables adjust to restore a cointegrating equilibrium subsequent to an equilibrium-distorting shock, and how long this adjustment

typically takes. Davis and Palumbo claim that their evidence implies that consumption does the subsequent adjusting and that this adjustment process may take several quarters or even years. It is with this claim that we take issue. In the next section, we argue that Davis and Palumbo's estimates of the adjustment parameter for consumption are incorrect, and that when this parameter is estimated correctly, it is much smaller (in absolute value) than they find. In this section, we show that, even if their estimate of this parameter were correct, their inferences concerning the length of time it would take consumption to adjust to an equilibrium-distorting shock do not follow because they have ignored the error-correction of other variables in the cointegrated system.

To discuss the estimation of what Davis and Palumbo call the error-correction speed for consumption, it will be useful to begin with a brief review of the vector error correction representation (VECM) for the three variables, c_t, a_t, and y_t. (A more extensive treatment of the long and short-run relations between these three variables can be found in Ludvigson and Steindel (1999), Lettau and Ludvigson (2001a), and Lettau and Ludvigson (2001d).) Because these three variables appear to share a single common trend, the parameters associated with the short-run dynamics are given by the following vector error-correction specification:

$$\Delta \mathbf{x}_t = \boldsymbol{\upsilon} + \boldsymbol{\gamma}\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1} + \boldsymbol{\Gamma}(L)\boldsymbol{\Delta}\mathbf{x}_{t-1} + \mathbf{e}_t, \tag{1}$$

where $\Delta \mathbf{x}_t$ is the vector of log first differences, $(\Delta c_t, \Delta a_t, \Delta y_t)'$, $\boldsymbol{\upsilon}$, and $\boldsymbol{\gamma} \equiv (\gamma_c, \gamma_a, \gamma_y)'$ are (3×1) vectors, $\boldsymbol{\Gamma}(L)$ is a finite order distributed lag operator, and $\widehat{\boldsymbol{\alpha}} \equiv (1, -\widehat{\alpha}_a, -\widehat{\alpha}_y)'$ is the (3×1) vector of previously estimated cointegrating coefficients. Throughout this comment, we use "hats" to denote the estimated values of parameters.

The term $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$ gives last period's equilibrium error, or cointegrating residual; $\boldsymbol{\gamma}$ is the vector of "adjustment" coefficients that tells us which variables participate to restore the long-run equilibrium when a deviation occurs. The Granger Representation Theorem states that, if a vector $\mathbf{x}_t$ is cointegrated, at least one of the adjustment parameters, γ_c, γ_a, or γ_y must be nonzero in the error-correction representation (1). Thus if x_j does at least some of the adjusting needed to restore the long-run equilibrium subsequent to a shock that distorts this equilibrium, γ_j should be different from zero in the equation for Δx_j of the error-correction representation (1).

Davis and Palumbo focus on estimating the element of $\boldsymbol{\gamma}$ that corresponds to the consumption equation in (1), γ_c. They call γ_c the error-correction speed for consumption. Some of their estimates, the ones they emphasize, suggest that this parameter is between -0.12 and -0.15. Davis and Palumbo argue that these estimates can be used to infer that the speed with which consumption adjusts subsequent to an equilibrium-distorting shock by constructing the monotonic sequence $1-(1+\gamma_c)^t$. This sequence is said to give the proportion of the disequilibrium that has been closed after t quarters by the sluggish adjustment in consumption. However, the sequence given above is only a valid measure of the proportion of the disequilibrium that has been closed after t quarters by the sluggish adjustment in consumption if consumption does *all* of the adjusting and wealth and labor income do *none* of it. If either wealth or labor income do even part of the adjusting, the adjustment by consumption will occur more quickly than implied by the sequence $1-(1+\gamma_c)^t$.

Davis and Palumbo implicitly assume that consumption does all of the adjusting and wealth and income do none of it; yet empirical evidence documented elsewhere suggests just the opposite, namely that wealth does all of the adjusting and consumption and labor income do virtually none of it. Davis and Palumbo seem to dismiss the possibility that wealth participates in the error-correction because, as they put it, such a phenomenon would "run counter to simple macroeconomic intuition" (page 26). We discuss the issue of whether such evidence makes economic sense below. The point here is an empirical one: because Davis and Palumbo focus on estimating a single equation for consumption, they miss strong evidence that wealth participates in the error-correction to an equilibrium-distorting shock. This is not the place to review all of that evidence; instead we simply provide the main elements here and note that the details can be found in Lettau and Ludvigson (2001a), and Lettau and Ludvigson (2001d).

These main elements are provided in Table 1, which reports estimates from the error-correction representations in (1) for Δc_t and Δa_t taken from Lettau and Ludvigson (2001d). The table reports that the adjustment parameter for the wealth equation, γ_a, is estimated to be large and strongly statistically different from zero at conventional significant levels (point estimate equal to 0.38; t-statistic equal to 3.24). By contrast, the estimated adjustment parameter for consumption, γ_c, is small and statistically insignificant at conventional signif-

icant levels (point estimate equal to -0.02; t-statistic equal to -0.66). This point estimate is considerably smaller (in absolute value) than that reported in most of Davis and Palumbo's regressions and the next section explains why that is so. The estimated adjustment parameter for the income, γ_y, (not shown) is also small and statistically indistinguishable from zero. Taken together, these results indicate that, when log consumption deviates from its habitual ratio with log labor income and log wealth, it is wealth, rather than consumption or labor income, that is forecast to adjust until the equilibrating relationship is restored.

To infer the speed of adjustment in one variable subsequent to an equilibrium-distorting shock, it is necessary to take into account the adjustment of all the variables in the cointegrated system. The importance of considering the entire system of cointegrated variables as opposed to a single equation of the VECM can be illustrated using a simple bivariate example. This bivariate example is a special case of (1) where there are two variables in $\mathbf{x}_t$ (which we denote x_{1t} and x_{2t}), one cointegrating relation (with cointegrating vector given by $\alpha = (1, -1)'$), and where, for simplicity, the parameters in $\mathbf{\Gamma}(L)$ are set to zero. (Whether the parameters in $\mathbf{\Gamma}(L)$ are zero or not will not matter for the basic argument given below.) For the purposes of this exercise, it matters not what type of shock is identified, only that the shock identified be one that causes a deviation from the long-run equilibrium relation given in this example by $x_{1t} = x_{2t}$. As long as the shock creates such a deviation, it is straightforward to calculate each variable's dynamic response to the equilibrium-distorting shock, under various assumptions about whether x_{1t} or x_{2t} or both participate in the adjustment needed to restore the long-run equilibrium.

Consider the case in which only x_{1t} participates in the adjustment needed to restore the long-run equilibrium. The VECM representation for this system might given by the following set of equations:

$$\begin{aligned} \Delta x_{1t} &= -0.5(x_{1t-1} - x_{2t-1}) + e_{1t} \\ \Delta x_{2t} &= e_{2t}, \end{aligned} \tag{2}$$

where the adjustment parameter for x_{1t} is equal to -0.5 and the adjustment parameter for x_{2t} is zero. We refer to this cointegrated system as example 1. Figure 1 plots the impulse responses of x_{1t} and x_{2t} to a unit shock in e_{2t}. The figure shows that the shock distorts

the long-run equilibrium $x_{1t} = x_{2t}$ by causing x_{2t} to increase on impact while x_{1t} remains unaffected. The long-run equilibrium is eventually restored as x_{1t} sluggishly returns to the value given by x_{2t}, equal to unity after the shock to e_{2t}. In this example, the proportion of the disequilibrium that has been closed after t periods is given by the sequence Davis and Palumbo use to calculate adjustment speeds, in this case, $1 - (1 - 0.5)^t$. After one quarter, 50 percent of the disequilibrium gap created by the shock has been closed by the adjustment in x_{1t}.

Compare these results with those from a system where *both* variables participate in the adjustment needed to restore the long-run equilibrium subsequent to an e_{2t} shock (example 2):

$$\begin{aligned} \Delta x_{1t} &= -0.5(x_{1t-1} - x_{2t-1}) + e_{1t} \\ \Delta x_{2t} &= 0.5(x_{1t-1} - x_{2t-1}) + e_{2t}. \end{aligned} \tag{3}$$

In this example, the adjustment parameters for x_{1t} and x_{2t} are equal to -0.5 and 0.5, respectively. We set the innovation in $e_{2t} = 2$ so that the long-run impact of the shock is the same as in example 1. The impulse responses to such a shock in e_{2t} for this system are given in Figure 2. The figure shows that the shock distorts the long-run equilibrium $x_{1t} = x_{2t}$ by causing x_{2t} to again increase on impact while x_{1t} initially remains unaffected. However, in this example, x_{2t} over-shoots its long-run value and the equilibrium is eventually restored as both x_{1t} and x_{2t} sluggishly adjust to their common value given by $x_{1t} = x_{2t} = 1$. Note that precisely the same adjustment parameter for x_{1t} appears in (2) and (3), and in both cases, precisely the same long-run value is reached by each variable. Yet in the first example, only 50 percent of the equilibrium error has been eliminated after one period, while in the second example, all of the equilibrium error has been eliminated after one period. Thus, $1 - (1 - 0.5)^t$ no longer reveals the speed of adjustment in x_{1t}, which occurs more quickly than implied by this sequence. Similarly, Davis and Palumbo's estimates of the sluggishness in the adjustment of consumption to an equilibrium distorting shock are likely to be inflated because they are only valid under the counterfactual assumption that consumption does all of the adjusting in the trivariate system with wealth and labor income.

One intuitively appealing way to investigate which variables in a cointegrated system

participate in the error-correction subsequent to an equilibrium distorting shock is to consider the *long-horizon forecasting* power of the cointegrating residual for the growth rates in each variable of the system. For example, if consumption sluggishly adjusts subsequent to a shock in labor income or wealth, the cointegrating residual, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$, should have long-horizon forecasting power for Δc_t. This follows because a shock to a_t or y_t will not be immediately accompanied by a full adjustment in c_t to its long-run value; these shocks will therefore initially create a cointegrating disequilibrium, or deviation in $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$. If the disequilibrium is subsequently closed by the sluggish adjustment of consumption to such shocks, the cointegrating error should forecast future movements in consumption. If, instead, wealth adjusts slowly over time, the cointegrating error should predict future movements in wealth.

These points may also be easily demonstrated by example. Consider a simple system for consumption, wealth and labor income:

$$\begin{aligned}
\Delta c_t &= \phi_{cc}\Delta c_{t-1} + \gamma_c(c_{t-1} - \widehat{\alpha}_a a_{t-1} - \widehat{\alpha}_y y_{t-1}) + u_{ct} \qquad (4)\\
\Delta y_t &= \phi_{cy}\Delta c_{t-1} + u_{yt} + \rho_{cy}u_{ct}\\
\Delta a_t &= \phi_{ca}\Delta c_{t-1} + u_{at} + \rho_{ya}u_{yt} + \rho_{ca}u_{ct} + \gamma_a(c_{t-1} - \widehat{\alpha}_a a_{t-1} - \widehat{\alpha}_y y_{t-1}),
\end{aligned}$$

where ϕ_{ij} and ρ_{ij}, $i, j = c, y, a$, are constants, γ_a corresponds to the second element of $\boldsymbol{\gamma}$ in (1), $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1} = (c_{t-1} - \widehat{\alpha}_a a_{t-1} - \widehat{\alpha}_y y_{t-1})$ and u_{ct}, u_{yt}, u_{at} are three shocks. Lettau and Ludvigson (2001d) find that the basic empirical pattern of the short-run dynamics for c, a, and y can be well captured by the structural VECM representation in (4) with $\gamma_c = 0$ and γ_a set equal to the estimated value reported in Table 1, about 0.38. Note that the adjustment parameter for income, γ_y, is set to zero. Setting the parameter $\phi_{cc} \geqslant 0$ captures the modest first-order serial correlation present in quarterly spending growth.[1] We simulate the system in (4) under various assumptions about the value of γ_c, and compute long-horizon forecasting regressions for the growth in consumption and the growth in wealth.

[1] Lettau and Ludvigson (2001d) discuss possible explanations for this serial correlation, including the time-averaging of aggregate consumption data. The other parameters in this system are set according to $\phi_{cy} = 0$, $\phi_{ca} = 0.2$, $\rho_{cy} = 0.5$, $\rho_{ca} = 0.1$; $\rho_{ya} = 0$, and $\gamma_a = 0.38$. A Monte Carlo simulation is run. The sample size of the simulation is 50,000 and the Monte Carlo is run 100 times. Table 2 gives the average statistics over 100 simulations.

Table 2 presents the results of using simulated data from the system (4) to generate forecasts of consumption growth and wealth growth at various horizons, H, i.e., forecasts of $\sum_{h=1}^{H} \Delta c_{t+h}$ and $\sum_{h=1}^{H} \Delta a_{t+h}$ using the cointegrating error, $(c_{t-1} - \widehat{\alpha}_a a_{t-1} - \widehat{\alpha}_y y_{t-1})$, as the single forecasting variable in each case. In panel A, $\gamma_c = 0$ so the cointegrating residual has no forecasting power for consumption growth at any horizon; both the R^2 statistic and the coefficient estimate are zero. By contrast, the cointegrating residual has substantial forecasting power for wealth growth since γ_a is nonzero and equal to 0.38. The R^2 for the asset wealth predictive regression peaks at about 0.67 four quarters out and the coefficient on $(c_{t-1} - \widehat{\alpha}_a a_{t-1} - \widehat{\alpha}_y y_{t-1})$ is strongly statistically significant.

What would the findings reported in Davis and Palumbo (2001) imply for these long-horizon predictive regressions? In Table 11 specification (C), the authors report an estimated value for $\gamma_c = -0.12$, and throughout the paper they maintain the assumption, implicitly in some places and explicitly in others, that wealth and income do none of the error-correction (i.e., $\gamma_a = \gamma_y = 0$). Panel B of Table 2 presents the results of using simulated data from the system (4) with the parameter $\gamma_c = -0.12$ and $\gamma_a = \gamma_y = 0$. Not surprisingly, we get the opposite forecasting pattern from that in Panel A: long-horizon consumption growth is now strongly forecastable by the cointegrating error (with the R^2 statistic peaking at 0.42 over a 12 quarter horizon), while long-horizon wealth growth is not at all predictable. Panel C shows that if $\gamma_c = -0.12$ *and* $\gamma_a = 0.38$, the cointegrating error has forecasting power for both $\sum_{h=1}^{H} \Delta c_{t+h}$ and $\sum_{h=1}^{H} \Delta a_{t+h}$: the coefficient estimates for $(c_{t-1} - \widehat{\alpha}_a a_{t-1} - \widehat{\alpha}_y y_{t-1})$ are strongly statistically significant in both forecasting equations and the R^2 statistics (which are hump-shaped in the horizon, H) are not negligible in either regression. What these simulated forecasting regressions demonstrate, is that, if the adjustment parameter is nonzero for one of the left-hand-side variables in the VECM (1), the cointegrating residual must have long-horizon forecasting power for the dependent variable of that equation.

For which variables does the cointegrating residual have long-horizon forecasting power in actual data? Table 3, reproduced from Lettau and Ludvigson (2001d), provides an answer. Panel A shows the long-horizon forecasts of consumption growth, controlling not only for the cointegrating error, but also for lags of Δc_t, Δa_t, and Δy_t. The coefficients on lagged Δc_t in the predictive regressions are small but statistically significant up to four quarters out,

reflecting the first-order serial correlation in quarterly spending growth. But, these findings provide no support for the proposition that consumption adjusts sluggishly to movements in income or wealth because the coefficient on the cointegrating error, $\widehat{\alpha}\mathbf{x}_{t-1} = (c_{t-1} - \widehat{\alpha}_a a_{t-1} - \widehat{\alpha}_y y_{t-1})$, is small and never statistically different from zero at any horizon in the future. By contrast, Panel B of Table 3 shows that the estimated cointegrating residual is a strong long-horizon forecaster of Δa_t; this variable is statistically significant and economically important at horizons ranging from one to 20 quarters, and the $\overline{R}^2$ statistic peaks at 0.37 at a 12 quarter forecast horizon. These results yield the same answers as those reported in Lettau and Ludvigson (2001a) from univariate forecasting regressions for Δc_t and Δa_t using the lagged cointegrating error as the sole predictive variable: those regressions show that the cointegrating residual has no forecasting power for future consumption growth, but substantial forecasting power for future wealth growth.

In summary, these results suggest that, if the adjustment parameter for consumption, γ_c, were nearly as large as that reported by Davis and Palumbo, the cointegrating residual for log consumption, asset wealth and labor income should tell us something about the future path of consumption. In fact, Table 3 shows that this residual has no forecasting power for consumption growth at any future horizon, but instead has strong forecasting power for the future growth in asset values. Accordingly, deviations from the common trend in c, a, and y, appear to be eliminated not by subsequent movements in consumption, but by subsequent movements in asset values.

But if consumption does not typically participate in the error-correction, how can Davis and Palumbo find a large (in absolute terms) and statistically significant value for the adjustment parameter, γ_c? One answer, we argue next, is that their estimates of this coefficient, based on reduced form equations for consumption growth, do not reveal the value of the true adjustment parameter, γ_c.

3 Estimating the Error-Correction Parameter for Consumption

The parameters in $\boldsymbol{\gamma}$ are what Davis and Palumbo call the error-correction speeds. We will instead refer to these coefficients simply as adjustment parameters, since, as discussed above, the value of γ_j tells us little about speed with which it takes x_j to adjust to an equilibrium-distorting shock when other variables in the cointegrating relation also do some of the adjusting.

The parameters in the error-correction representation, including those in $\boldsymbol{\gamma}$, may be consistently estimated by ordinary least squares estimation of (1) equation-by-equation (Engle and Granger (1987); Stock (1987)). This is not how Davis and Palumbo estimate γ_c, however, the element corresponding to the adjustment parameter for consumption. Instead, they estimate a single equation regression for consumption growth, adding conditioning variables to the set of explanatory variables that are part of the error-correction representation (1). Specifically, their single equation regression for consumption growth takes the form

$$\Delta c_t = \upsilon_c + \overline{\gamma}_c \widehat{\boldsymbol{\alpha}}' \mathbf{x}_{t-1} + \overline{\boldsymbol{\Gamma}}(L) \boldsymbol{\Delta} \mathbf{x}_{t-1} + \mathbf{C}(L) \mathbf{z}_t + u_t, \tag{5}$$

where $\mathbf{z}_t$ is a vector of additional predetermined regressors that are not part of the error-correction specification (1), where $\overline{\boldsymbol{\Gamma}}(L)$ and $\mathbf{C}(L)$ are lagged polynomial vectors, and $\overline{\gamma}_c$ is a coefficient, all of which Davis and Palumbo estimate by running ordinary least squares on (5).

Given that consistent estimates of the adjustment parameters in $\boldsymbol{\gamma}$ can be obtained by ordinary least squares estimation of (1) equation-by-equation, it is reasonable to ask why one would augment that specification by including the additional regressors $\mathbf{C}(L)\mathbf{z}_t$? One reason for doing so is that, in finite samples, efficiency gains can be made by including additional variables if they are in fact important short-run determinants of consumption growth. Both the estimated variance of the residuals as well as the standard errors of the coefficient estimates will be smaller in (5) than in (1) if the additional variables are truly relevant for Δc_t. Indeed, efficiency gains seem to be the main motivation behind Davis and Palumbo's estimation of the specification (5), since they argue on pages 39-40 that such

specifications "provide a better fit to the historical data on consumption growth" than does the standard error-correction specification which excludes the additional regressors.

The key question now concerns whether the adjustment parameter, γ_c, from (1) can be revealed by estimation of a specification like (5) that includes the additional regressors $\mathbf{C}(L)\mathbf{z}_t$. It turns out that they can, only under very special circumstances. In the likely event that the additional explanatory variables are not orthogonal to the cointegrating error, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$, estimation of (5) will not uncover $\boldsymbol{\gamma}$, the adjustment parameters in (1). In the language of Engle, Hendry, and Richard (1983), if the variables in $\mathbf{z}_t$ are not weakly exogenous for the vector $\boldsymbol{\gamma}$, the adjustment parameter for consumption, γ_c, cannot be recovered from estimation of a single equation specification like (5). This point has been emphasized by Johansen (1992) and Seo (1998). The reason for this is that the adjustment parameters in $\boldsymbol{\gamma}$ give the influence of the cointegrating error, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$, on each variable in $\Delta\mathbf{x}_t$, controlling for a specific set of variables–namely those captured by $\boldsymbol{\Gamma}(L)\boldsymbol{\Delta}\mathbf{x}_{t-1}$. Once this specification is altered to include the additional explanatory variables that are not weakly exogenous, estimates of the coefficient $\overline{\gamma}_c$ in (5) will no longer reveal an estimate of γ_c, the adjustment parameter of interest in (1). We underscore this point by denoting the parameters of the short-run dynamics in (5) as $\overline{\gamma}_c$ and $\overline{\boldsymbol{\Gamma}}$ to signify that they are in general distinct from the parameters γ_c and $\boldsymbol{\Gamma}$ in the error-correction specification (1).

If the researcher is interested in determining which variables typically participate in the error-correction subsequent to an equilibrium-distorting shock, it is not the coefficient $\overline{\gamma}_c$ in (5), but the adjustment parameters in $\boldsymbol{\gamma}$ from (1), that are of interest. The general idea is that the fitted coefficient γ_c in (1) captures the covariation between the last period's cointegrating error, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$, and this period's growth in consumption that cannot be captured by variation in $\boldsymbol{\Gamma}(L)\boldsymbol{\Delta}\mathbf{x}_{t-1}$ alone–that is independent only of past movements by variables involved in the cointegrating relation itself. In estimating this adjustment parameter for consumption, one would not want to *remove* variation in the cointegrating error, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$, that is correlated with the predetermined variables in $\mathbf{z}_t$, for such movements may be among the most quantitatively important sources of independent variation in $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$. If the regressors in $\mathbf{z}_t$ and their lags are correlated with $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$, the estimate of $\overline{\gamma}_c$ that will be obtained from (5) will tell us something about how consumption adapts to a cointegrating disequilibrium that is *not* associated with

variation in $\mathbf{z}_t$ (and its lags), but nothing about how consumption adapts to variation in $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$ that *is* associated with variation in $\mathbf{z}_t$. If most of the variation in $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$ is in fact associated with variation in $\mathbf{z}_t$, such an estimate will tell us virtually nothing about how consumption adapts to most disequilibrium shocks.

To understand how consumption adapts to most disequilibrium shocks, one must estimate the parameters of interest, $\boldsymbol{\gamma}$, which can be achieved by estimating (1) using ordinary least squares equation-by-equation. In large samples, such an estimation strategy will be a good one. In small samples, however, we would forego possible efficiency gains that might be achieved by including additional explanatory variables for the growth rates of the cointegrated variables. Thus, we need a procedure that decouples achieving these efficiency gains from the inclusion of additional regressors that are not weakly exogenous for the adjustment parameters of interest.

Such a decoupling is straightforward using a simple two-step procedure. In the first step, an orthogonal measure of $\mathbf{z}_t$ is obtained by regressing each element of $\mathbf{z}_t$ on all of the right-hand-side variables in (1) and saving the residuals, $\widetilde{\mathbf{z}}_t$. In the second step, estimation of (5) is carried out by including $\widetilde{\mathbf{z}}_t$ in place of $\mathbf{z}_t$ as additional explanatory variables. Note that efficiency gains will be preserved in the replacement of $\mathbf{z}_t$ with $\widetilde{\mathbf{z}}_t$ if there is information in $\mathbf{z}_t$ that is not already captured by movements in $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$ and the terms in $\boldsymbol{\Gamma}(L)\boldsymbol{\Delta}\mathbf{x}_{t-1}$. At the same time, the procedure allows the researcher to uncover an estimate of γ_c since the regressors in $\widetilde{\mathbf{z}}_t$ will be weakly exogenous for the adjustment parameters of interest, $\boldsymbol{\gamma}$.

To illustrate this point, we first obtained the data used by Davis and Palumbo and reproduced the regression results reported in their Table 11, specification C. These results are reproduced by estimating their specification of (5). In this specification, they include the following variables in $\mathbf{z}_t$: a fitted value for current income growth, which we will denote $\Delta\widehat{y}_t \equiv z_{1t}$, a lag of the change in the unemployment rate, $DUNEMP_{t-1} \equiv z_{2t}$, a lag of an inflation adjusted federal-funds-rate, $RFF_{t-1} \equiv z_{3t}$, and the current level of the University of Michigan Unemployment Expectations index, $UNEXP_t \equiv z_{4t}$.[2] This specification sets $\mathbf{z}_t =$

[2]The paper states that the Michigan Consumer Sentiment Index was used as an explanatory variable, but communication with the authors revealed that they instead used the Michigan Unemployment Expectations Index.

$(z_{1t}, z_{2t}, z_{3t}, z_{4t})' \equiv (\Delta\widehat{y}_t, DUNEMP_{t-1}, RFF_{t-1}, UNEXP_t)'$. Davis and Palumbo appear to set the lag order in (5) equal to one, except that they exclude lags of income growth.[3] To reproduce their results, we use their measure of $\Delta\widehat{y}_t$, which they provided for us based on the fitted values from a regression of Δy_t on lags of income growth, wealth growth, consumption growth, the unemployment rate and the real federal-funds-rate. In addition, we use their timing convention for wealth (measuring it as end-of-period), their measure of consumption growth (total personal consumption expenditures) and their sample period: 1960:Q1 2000:Q1.

In the two-step procedure, described above, the orthogonalized regressors, $\widetilde{\mathbf{z}}_t$, are included in place of $\mathbf{z}_t$ as explanatory variables in a regression for consumption growth. The second-stage regression takes the form

$$\Delta c_t = \upsilon_c + \gamma_c \widehat{\boldsymbol{\alpha}}' \mathbf{x}_{t-1} + \boldsymbol{\Gamma}(L)\Delta c_{t-1} + \widetilde{\mathbf{C}}(L)\widetilde{\mathbf{z}}_t + \widetilde{u}_t, \tag{6}$$

where $\widetilde{\mathbf{z}}_t$ is the vector of residuals from a regression of each variable in $\mathbf{z}_t = (z_{1t}, z_{2t}, z_{3t}, z_{4t})' \equiv (\Delta\widehat{y}_t, DUNEMP_{t-1}, RFF_{t-1}, UNEXP_t)'$ on all of the right-hand-side variables of a second-order specification of the standard error-correction representation (1).[4] Note that (6) is directly analogous to (5).

Table 4 compares the results from our estimation of Davis and Palumbo's specification of (5) (given in Panel B) with the alternate specification (6) (given in Panel C). To provide a benchmark, Panel A of this table shows the results of estimating the consumption equation of the standard error-correction specification (1) over the same sample period and using the same data used in Davis and Palumbo (2001).

Estimation of the standard error-correction representation in Panel A suggests that the adjustment parameter for consumption growth, γ_c, is small (point estimate equal to -0.04) and statistically indistinguishable from zero (t-statistic equal to -0.99). Similar results are reported in Table 1 where only nondurables and services expenditure is used. By contrast,

[3] Presumably these lags are excluded on the grounds that they have little marginal predictive power once the fitted value of income growth, $\Delta\widehat{y}_t$, is included.

[4] This second-order lag length was chosen in accordance with the Akaike and Schwarz criteria.

Panel B shows that estimation of the specification considered by Davis and Palumbo (equation (5)) uncovers a value for $\overline{\gamma}_c$ that is much larger in absolute value (point estimate equal to -0.128). (We discuss the t-statistics for this regression below.) Note that this point estimate reproduces the result given in the third row of Table 11 of Davis and Palumbo (2001). Comparing the results reported in Panels A and B reveals the improvement in fit and efficiency that is obtained by including the additional variables $\mathbf{z}_t$ and its lags in the regression: the standard error of the regression is smaller in Panel B (equal to 0.005) than it is in Panel A (equal to 0.006).

Now consider the results from estimating the alternative specification (6) reported in Panel C. Comparing the regression output with that in Panel B demonstrates that there is no efficiency loss in moving from (5) to (6): the standard error of the estimate is precisely the same in Panel C as it is in the Davis-Palumbo regression specification of Panel B. More important, the estimate of γ_c reported in Panel C is much smaller (in absolute value) than the estimate of $\overline{\gamma}_c$ reported in Panel B; indeed the former is essentially zero (point estimate equal to -0.005, t-statistic equal to -0.18). The reason these estimates differ so greatly is that the cointegrating error, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$, is strongly marginally correlated with several of the variables contained in $\mathbf{z}_t$, including the real federal-funds-rate, z_{3t}, the change in the unemployment rate, z_{2t}, and the fitted income growth variable, z_{1t}. Thus, the additional covariates Davis and Palumbo include in the standard error-correction specification are not weakly exogenous for the parameter they seek to estimate and hence the true adjustment parameter for consumption cannot be recovered by estimating (5). Fortunately, they can be recovered by estimating (6), all the while preserving the efficiency gains in (5).

It is worth emphasizing that no loss of fit is incurred by moving from the specification employed by Davis and Palumbo in equation (5) to that in (6). This is important because Davis and Palumbo seem to favor their estimates of $\overline{\gamma}_c$ from the specification in (5) on the grounds that they are obtained from better fitting consumption growth regressions than are the other estimates they report in their Table 11. But since the empirical specification we suggest in (6) fits just as well as that in (5), such considerations provide no reason to favor one specification over the other. What *does* recommend the specification in (6) over that in (5) is that the additional covariates in the former are–unlike those in the latter–weakly exogenous

for both the cointegrating relation and the adjustment coefficient, γ_c. Accordingly, valid estimates of this adjustment parameter may be obtained from estimating the specification in (6), but not from (5).

The estimate of γ_c reported in Panel C is almost one-tenth of the size of the estimate of this parameter from a standard error-correction representation, reported in Panel A, although both are quite small and statistically indistinguishable from zero. Since the empirical model in Panel C improves efficiency over the model in Panel A, this result suggests that the true value of the adjustment parameter for consumption is even *smaller* than what would be suggested by an estimation of the error-correction representation (1), and is quite close to zero. Similarly, Lettau and Ludvigson (2001d) find that the estimate of the adjustment parameter for labor income, γ_y, is small and statistically indistinguishable from zero.[5] These findings strengthen the conclusion that almost all of the adjustment needed to restore the cointegrating equilibrium between c, a, and y, is done by wealth, while very little of it is done by consumption or labor earnings.

We now come back to the issue of the t-statistic for the estimate of $\overline{\gamma}_c$ in (6). The traditional OLS estimate of the standard error underlying the t-statistic for this parameter (reported in Panel B of Table 4) is unlikely to be correct, since the regression uses a generated regressor, $\Delta\widehat{y}_t$. Communication with the authors also revealed that no correction was made in Davis and Palumbo (2001) to the standard errors for the use of this generated regressor. Thus the standard errors they report in their Table 11, as well as those in Panel B of our Table 4 are in general invalid. Since this generated variable is procured as the fitted-value from a regression of Δy_t on a set of regressors that does not include all of the right-hand-side variables in (5), however, obtaining the correct standard errors for their two-step method is likely to be quite complex. Fortunately, it is easy to reestimate the Davis-Palumbo specification and obtain a consistent estimate of the standard errors by following

[5] If the three exogenous conditioning variables $\widetilde{z}_{2t}$, $\widetilde{z}_{3t}$, $\widetilde{z}_{4t}$ are included in an single equation error-correction specifications for Δy_t and Δa_t, the estimate of the adjustment parameter for income growth, γ_y, is reduced further and is zero out to the third decimal place, while the adjustment parameter for wealth, γ_a remains about the same as that reported in Table 1, close to 0.38.

the recommendation of Pagan (1984) and applying instrumental variables (IV) to

$$\Delta c_t = \upsilon_c + \overline{\gamma}_c \widehat{\boldsymbol{\alpha}}' \mathbf{x}_{t-1} + \overline{\boldsymbol{\Gamma}}(L)\boldsymbol{\Delta}\mathbf{x}_{t-1} + \mathbf{C}(L)\mathbf{z}_t^* + \mathbf{u}_t, \tag{7}$$

where $\mathbf{z}_t^* \equiv (\Delta y_t, DUNEMP_{t-1}, RFF_{t-1}, UNEXP_t)'$ and the instrument set includes all of the regressors in (7) except Δy_t, in addition to those instruments Davis and Palumbo used to form a fitted value for Δy_t, namely four lags of Δc_t, Δy_t, Δa_t, the federal-funds-rate, and the unemployment rate. Note that the specification in (7) is the precisely the same as that in (5) except that a few additional instruments are used to form a fitted value for Δy_t.

The results of this IV regression (7) are reported in Panel B Table 5, again using the data and sample period employed by Davis and Palumbo. For ease of comparison, Panel A of Table 5 reproduces Panel A of Table 4 showing the estimates from a standard error-correction specification for Δc_t. Finally, in analogy to the regression carried out in Panel C of Table 4, Panel C of Table 5 reports the results of an OLS estimation of (7) where orthogonal regressors, $\widetilde{\mathbf{z}}_t^*$, replace $\mathbf{z}_t^*$. As before, this orthogonalization procedure is necessary to obtain valid estimates of the adjustment parameter, γ_c, and is accomplished by taking the residuals from a regression of each of the variables in $\mathbf{z}_t^*$ on all of the right-hand-side variables in a second-order specification of the error-correction representation (1).[6]

Notice that the estimate of $\overline{\gamma}_c$ given in Panel B is precisely the same as that obtained in Davis and Palumbo's Table 11, row 3. Thus, these findings again reproduce their result. Also, the t-statistic (which is now valid in the presence of the generated regressor, $\Delta\widehat{y}_t$) shows that this parameter is strongly statistically significant, and the standard error of the estimate is smaller than that in Panel A by the same order of magnitude as that reported in Table 4.

As before, however, the results from the alternative specification using orthogonal regressors, given in Panel C of Table 5, show that the estimate of γ_c is about zero (point estimate

[6]Recall that the first variable in $\mathbf{z}_t^* = (z_{1t}, z_{2t}, z_{3t}, z_{4t})' \equiv (\Delta y_t, DUNEMP_{t-1}, RFF_{t-1}, UNEXP_t)'$ is the current value of labor income growth. Because IV estimation is used in (7) and Δy_t is *not* included in the instrument set, the regressor z_{1t} in (7) is actually the fitted value from a regression of income growth on the all of the instrumental variables used in the IV regression (7). Thus, the orthogonal variable $\widetilde{z}_{1t}^*$ is created by first regressing Δy_t on each of these instruments, and saving the fitted values. Next, those fitted values are regressed on all of the right-hand-side variables in (1) and the residuals are stored as $\widetilde{z}_{1t}^*$.

equal to -0.004, t-statistic equal to -0.096) and about one-tenth of the size of that obtained in Panel A. In addition, efficiency gains created by including the conditioning variables $\mathbf{z}_t^*$ in the VECM specification Panel B are again preserved in Panel C where the orthogonalized variables, $\widetilde{\mathbf{z}}_t^*$ are used in place of $\mathbf{z}_t^*$. (The standard error of the estimate in Panel B of Table 5 is the same as that in Panel C.)

In summary, the results in this section provide no support for the conclusion in Davis and Palumbo (2001) that aggregate consumer spending adjusts only gradually to movements in income or wealth, or that a "sudden increase in wealth leads toward a period of faster than normal consumption growth" (page 39). Instead, the adjustment parameter for consumption growth–when estimated correctly–appears to be about zero, indicating that spending typically adapts within the span of roughly one quarter to fluctuations in income and wealth. Since consumption does not participate in the error-correction, and since lagged values of wealth growth have little impact on consumption growth, movements in wealth can have important implications for consumption contemporaneously, but they bear little relation to *future* consumer spending.[7]

We emphasize that none of the results we have presented here, nor those presented in Ludvigson and Steindel (1999) or Lettau and Ludvigson (2001d), imply that wealth has no impact on aggregate consumption. Indeed, because consumption is tied to wealth and labor earnings in the long-run, our findings imply that permanent movements in wealth *must* influence spending, and that they typically do so within the span of about one quarter. As we stress in Lettau and Ludvigson (2001d), however, not all movements in wealth appear to be permanent, and the aggregate data suggest that unsustainable, or transitory, changes in wealth have little influence on consumer spending. A more detailed discussion of the evidence supporting these conclusions can be found in (Lettau and Ludvigson (2001d)).

4 Other Remarks

Although the results in Davis and Palumbo (2001) do not say much about the speed and extent to which consumption adjusts to fluctuations in income and wealth, they do bear on

[7]This point was emphasized in Ludvigson and Steindel (1999).

an important forecasting question, namely "what observable indicators are related to future consumption growth?" Their specification (C) in Table 11 shows that forecastable movements in income growth are correlated with consumption growth, a finding that has been reported elsewhere (e.g., Campbell and Mankiw (1989)). But they also show that the Michigan Unemployment Expectations index and the real federal-funds-rate are strongly related to future consumption growth. (Because the unemployment index is a contemporaneous regressor, the equation they estimate is not quite a forecasting equation using historical data, but it is a forecasting regression in real time since the unemployment index is available much sooner than are consumption data.) These forecasting results are potentially interesting to both academic researchers and practitioners.

Nevertheless, some caveats are in order even about these forecasting results. The variables that Davis and Palumbo selected to include in the consumption growth regression (5) were obviously chosen with the benefit of hindsight. Although this observation does not by itself eradicate the possibility that there is real information for consumption in these forecasting variables, there is a well known problem with choosing predictive variables after their predictive power has been verified that nevertheless seems warranted in this instance.

Other caveats should be given about Davis and Palumbo's conclusions. First, in the discussion on the top of page 26, the authors seem to rationalize ignoring empirical evidence, documented in Lettau and Ludvigson (2001a) and Lettau and Ludvigson (2001d), that wealth participates in the error correction on the grounds that such evidence would imply predictability of equity returns, "leaving open a relatively easy way to make money in the stock market." There is now a large and growing body of empirical evidence documented in the field of asset pricing that suggests equity returns *are* forecastable, and that the error-correction in asset wealth discussed here reflects precisely this forecastability (Lettau and Ludvigson (2001a)). Nevertheless, it is reasonable to ask whether such predictability in asset values makes economic sense. As it turns out, research in field of theoretical finance over the last 15 years has shown that predictability in the stock market is not necessarily inconsistent with market efficiency. Accordingly, it is not necessarily the case that the average investor can make money from such predictability, a phenomenon that can be generated by time-variation in the rate at which rational, utility maximizing investors discount expected future

income from risky assets. Nice overview treatments of these issues along with a review of the empirical evidence can be found in Campbell, Lo, and MacKinlay (1997), chapters 7 and 8, and Cochrane (2001), chapter 20.

Second, the authors argue on page 16 that the unusually low ratio of consumption to asset wealth in the last half of the 1990s tells us little about how responsive consumption has been during this period to movements in wealth. While this argument may be valid as a general statement, it begs an important question about this recent period, namely why the cointegrating residual, $c_t - \widehat{\alpha}_a a_t - \widehat{\alpha}_y y_t$–driven by the surge in stock values–reached such unusually low levels in the late 1990s? One possible explanation for this phenomenon, given in Lettau and Ludvigson (2001a), is that this cointegrating residual is low when excess returns on equity are expected to be low in the future.

Third, in most of their empirical analysis, Davis and Palumbo analyze the wealth effect by asking whether end-of-period wealth is associated with consumption, the latter measured as a flow over the period (consumption data are time-averaged). A timing convention for wealth is needed since the level of consumption is a flow during the quarter rather than a point-in-time estimate. Their discussion is less clear, however, about the assumptions required to make this a reasonable modeling strategy for measuring the structural effects of wealth on consumption. In order to address these structural questions (as opposed to questions of pure forecasting), a_t must be in the information set when consumption, c_t is chosen. If we think of consumption for a given quarter as measuring spending at the beginning of the quarter, then the appropriate measure of wealth for these investigations is not end-of-period wealth, but beginning-of-period wealth. We view this is a reasonable timing convention since in this scenario households can "stock their refrigerator" at the beginning of the period and consume by running down that stock during the period. In order for the end-of-period measure used by Davis and Palumbo to be valid, one must implicitly assume that households consume in one instant on the last day of the period after the markets close, and starve the rest of the period. We find the alternative beginning-of-period assumption more sensible because is permits consumption to take place before the last instant at the end of the period.

Finally, on pages 22, 23 and 25 the authors argue that the estimated cointegrating coefficients, $\widehat{\alpha}_a$ and $\widehat{\alpha}_y$, (the regression coefficients presented in their Table 5) will reflect the

influence of c_t on a_t and y_t (rather than the other way around), unless it is the case that consumption participates in the error-correction subsequent to an equilibrium-distorting shock. For example, on page 25: "This highlights the fact that one needs to find evidence that spending exhibits error-correction behavior in the short-run to assert that changes in income or wealth eventually generate changes in consumption in the long-run." There are two elements to this argument: the first concerns endogeneity bias itself (the possibility that the estimated cointegrating coefficients $\widehat{\alpha}_a$ and $\widehat{\alpha}_y$ will reflect the influence of c_t on a_t and y_t rather than the marginal influence of a_t and y_t on c_t), and the second suggests that there is a connection between such endogeneity bias and absence of error-correction behavior in consumption.

From a statistical perspective, the first element (endogeneity bias) is unlikely to be a problem. In sufficiently large samples, dynamic least squares estimates of the cointegrating coefficients are robust to endogeneity of the regressors because they are "superconsistent," converging to the true parameter values at rate proportional to the sample size T rather than proportional to $\sqrt{T}$ as in ordinary applications (Stock (1987)). Moreover, Monte Carlo simulations, discussed in Lettau and Ludvigson (2001d) and calibrated to match the data generating processes of consumption, asset wealth, and labor income, suggest that samples of the size currently encountered are indeed sufficiently large: the cointegrating coefficient estimates can be very accurately recovered in simulated samples of the size now available.

The second element in this argument is invalid, for two reasons. First, the superconsistency results just discussed are a purely statistical property of cointegrated systems that is not related to which variables participate in the error-correction subsequent to an equilibrium-distorting shock. Second, there is no logical inconsistency between the presence of a wealth or income effect that will influence consumption in the long-run on the one hand, and the absence of error-correction behavior in consumption on the other. A notable counterexample is a frictionless permanent income model in which only permanent changes in wealth and income influence consumption and spending adjusts fully to such changes within the period. Examples can be found in the classic permanent income models investigated by Campbell (1987) and Galí (1990), in which only labor income displays error-correction, or in more general permanent income models that allow for time-varying expected asset returns,

in which case both asset wealth and labor income may display error-correction (Lettau and Ludvigson (2001d)). In each of these examples, none of the error-correction is done by consumption, yet the cointegrating coefficients α_a and α_y give the effect on consumption of permanent changes (changes that are sustainable in the long-run) in wealth and income. There is no error-correction in consumption in these models because spending adjusts *within the period* to permanent movements in wealth and income. It follows that the presence of a "long-run" wealth effect need not be inextricably linked to error-correction behavior in consumption, a phenomenon that tells us something about the length of time over which consumption adjusts to permanent changes in wealth, but nothing about the magnitude of those effects.

5 Concluding Remarks

The surge in asset values in recent years has presented a legion of new challenges to macroeconomists faced with analyzing the real economy. To meet these challenges, it would seem necessary to combine insights from finance with those from macroeconomics, and we think that doing so constitutes an important step in understanding the relation between aggregate consumption and household wealth. Yet research in financial economics has often proceeded independently of that in macroeconomics. For example, empirical work in the former literature suggests that expected returns on aggregate stock market indexes in excess of a short-term interest rate vary significantly over time. Perhaps because many macroeconomists are accustomed to thinking in terms of constant discount rates, these findings are often overlooked in the latter literature.

The remarks in this comment draw largely from what we have learned in our recent work that attempts to integrate these two literatures empirically (Lettau and Ludvigson (2001a); Lettau and Ludvigson (2001b); Lettau and Ludvigson (2001d); Lettau and Ludvigson (2001c)). We believe that such an integrated focus will continue to prove fruitful, as economists fumble their way forward in the quest to understand how asset markets and the real economy are interrelated.

References

Campbell, J. Y. (1987): "Does Saving Anticipate Declining Future Labor Income? An Alternate Test of the Permanent Income Hypothesis," *Econometrica*, 55, 1249–73.

Campbell, J. Y., A. W. Lo, and C. MacKinlay (1997): *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ.

Campbell, J. Y., and G. Mankiw (1989): "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," in *NBER Macroeconomics Annual: 1989*, ed. by O. Blanchard, and S. Fischer, pp. 185–216. MIT Press, Cambridge, MA.

Cochrane, J. H. (2001): *Asset Pricing*. Princeton University Press, Princeton, NJ.

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

Engle, R. F., and C. W. J. Granger (1987): "Co-integration and Error Correction Representation, Estimation and Testing," *Econometrica*, 55, 251–276.

Engle, R. F., D. F. Hendry, and J.-F. Richard (1983): "Exogeneity," *Econometrica*, 51(2), 277–304.

Galí, J. (1990): "Finite Horizons, Life-Cycle Savings, and Time-Series Evidence on Consumption," *Journal of Monetary Economics*, 26, 433–452.

Johansen, S. (1992): "Cointegration in Partial Systems and the Efficiency of Single-Equation Analysis," *Journal of Econometrics*, 52, 389–402.

Lettau, M., and S. Ludvigson (2001a): "Consumption, Aggregate Wealth and Expected Stock Returns," *Journal of Finance,* forthcoming.

——— (2001b): "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia are Time-Varying," *Journal of Political Economy,* forthcoming.

——— (2001c): "Time-Varying Risk Premia and the Cost of Capital: An Alternative Implication of the Q Theory of Investment," Unpublished Paper, Federal Reserve Bank of New York.

——— (2001d): "Understanding Trend and Cycle in Asset Values," Unpublished Paper, Federal Reserve Bank of New York.

Ludvigson, S., and C. Steindel (1999): "How Important Is the Stock Market Effect on Consumption?," *Economic Policy Review, Federal Reserve Bank of New York*, 5, 29–51.

Pagan, A. (1984): "Econometric Issues in the Analysis of Regressions with Generated Regressors," *International Economic Review*, 25(1), 221–247.

Seo, B. (1998): "Statistical Inference on Cointegrating Rank in Error-Correction Models with Stationary Covariates," *Journal of Econometrics*, 85, 339–385.

Stock, J. H. (1987): "Asymptotic Properties of Least Squares Estimators of Cointegrating Vectors," *Econometrica*, 55, 113–144.

Table 1: VECM Estimates using U.S. Data

Variable	Estimate	t-Stat
	Δc_{t+1}	
k	0.015	0.857
$\widehat{\alpha}'\mathbf{x}_{t-1}$	-0.021	-0.672
Δc_{t-1}	**0.222**	2.459
Δc_{t-2}	0.076	0.859
Δa_{t-1}	0.000	0.039
Δa_{t-2}	0.016	0.789
Δy_{t-1}	0.085	1.799
Δy_{t-2}	-0.022	-0.474
$\overline{R}^2$	0.11	
	Δa_{t+1}	
k	**-0.021**	-3.239
$\widehat{\alpha}'\mathbf{x}_{t-1}$	**0.387**	3.287
Δc_{t-1}	**0.676**	1.974
Δc_{t-2}	-0.140	-0.420
Δa_{t-1}	0.092	1.220
Δa_{t-2}	0.004	0.059
Δy_{t-1}	0.348	1.960
Δy_{t-2}	-0.179	-1.024
$\overline{R}^2$	0.11	

Notes for Table 1: The table reports results taken from Lettau and Ludvigson (2001b) of the single equation estimates for Δc_{t+1} and Δa_{t+1} from the VECM $\Delta\mathbf{x}_t = k + \boldsymbol{\gamma}\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1} + \boldsymbol{\Gamma}(\mathbf{L})\boldsymbol{\Delta}\mathbf{x}_{t-1} + \mathbf{v}_t$,for $\mathbf{x}_t = (c, y, a)'$ using data from 1952:Q4 to 1999:Q4. Significant coefficients at the 5% level are highlighted in bold face.

Table 2: Long-Horizon Regressions Using Simulated Data

H	Panel A: $\gamma_c = 0, \gamma_a = 0.38$				Panel B: $\gamma_c = -0.12, \gamma_a = 0$				Panel C: $\gamma_c = -0.12, \gamma_a = 0.38$			
	$\sum_{h=1}^{H} \Delta c_{t+h}$		$\sum_{h=1}^{H} \Delta a_{t+h}$		$\sum_{h=1}^{H} \Delta c_{t+h}$		$\sum_{h=1}^{H} \Delta a_{t+h}$		$\sum_{h=1}^{H} \Delta c_{t+h}$		$\sum_{h=1}^{H} \Delta a_{t+h}$	
	Estimate (t-Stat)	$\overline{R}^2$	Estimate (t-Stat)	$\overline{R}^2$	Estimate (t-Stat)	$\overline{R}^2$	Estimate (t-Stat)	$\overline{R}^2$	Estimate (t-Stat)	$\overline{R}^2$	Estimate (t-Stat)	$\overline{R}^2$
1	0.00 (0.38)	0.00	**0.38** (68.40)	0.48	**-0.12** (-21.04)	0.09	-0.01 (-1.41)	0.00	**-0.12** (-15.85)	0.05	**0.38** (50.61)	0.33
2	0.00 (0.39)	0.00	**0.36** (69.52)	0.63	**-0.12** (-22.09)	0.16	-0.01 (-1.26)	0.00	**-0.11** (-16.66)	0.08	**0.33** (51.16)	0.44
4	0.00 (0.45)	0.00	**0.32** (60.62)	0.67	**-0.10** (-23.29)	0.26	-0.01 (-1.21)	0.00	**-0.08** (-16.11)	0.12	**0.26** (42.36)	0.45
8	0.00 (0.76)	0.00	**0.26** (41.15)	0.58	**-0.08** (-25.23)	0.39	-0.01 (-1.09)	0.00	**-0.05** (-14.78)	0.13	**0.17** (27.62)	0.32
12	0.00 (0.93)	0.00	**0.21** (29.48)	0.47	**-0.07** (-25.15)	0.42	-0.00 (-0.40)	0.00	**-0.04** (-13.39)	0.11	**0.12** (19.40)	0.23
16	0.00 (1.06)	0.00	**0.18** (23.23)	0.39	**-0.05** (-23.06)	0.40	0.00 (0.13)	0.00	**-0.03** (-11.85)	0.09	**0.10** (14.94)	0.17
20	0.01 (1.45)	0.00	**0.15** (18.89)	0.32	**-0.04** (-21.09)	0.37	0.00 (0.47)	0.00	**-0.02** (-10.58)	0.08	**0.08** (12.48)	0.14

Notes for Table 2: The table reports output from long-horizon regressions of consumption and asset wealth on the cointegrating error for consumption, wealth and labor income, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$, using simulated data. The dependent variables in the h-period regressions are $\Delta x_{t+1} + ... + \Delta x_{t+h}$, where $x \in \{c, a\}$. For each regression, the table reports OLS estimates of the regressors, Newey-West corrected t-statistics (in parentheses) and adjusted R^2 statistics. Significant coefficients at the 5% level are highlighted in bold face.

Table 3: Long-Horizon Regressions Using U.S. Data

Panel A: $\sum_{h=1}^{H} \Delta c_{t+h}$ regressed on					
Horizon H	Δc_t	Δy_t	Δa_t	$c_t - \widehat{\alpha}_a a_t - \widehat{\alpha}_y y_t$	$\bar{R}^2$
1	**0.23**	**0.09**	0.00	-0.02	0.14
	(2.89)	(2.00)	(0.20)	(-0.68)	
2	**0.40**	0.10	0.02	-0.04	0.12
	(2.98)	(1.14)	(0.88)	(-0.44)	
4	**0.67**	0.08	0.05	-0.01	0.10
	(3.38)	(0.71)	(1.19)	(-0.07)	
8	0.67	0.07	-0.02	-0.00	0.03
	(1.58)	(0.47)	(-0.31)	(-0.01)	
12	0.68	0.16	-0.09	-0.16	0.03
	(1.36)	(0.87)	(-0.84)	(-0.50)	
16	0.65	0.20	-0.11	-0.29	0.03
	(1.13)	(0.89)	(-0.96)	(-0.69)	
20	0.47	0.18	-0.09	-0.31	0.02
	(0.68)	(0.63)	(-0.65)	(-0.62)	
Panel B: $\sum_{h=1}^{H} \Delta a_{t+h}$ regressed on					
Horizon H	Δc_t	Δy_t	Δa_t	$c_t - \widehat{\alpha}_a a_t - \widehat{\alpha}_y y_t$	$\bar{R}^2$
1	**0.52**	**0.33**	0.08	**0.41**	0.13
	(1.98)	(2.03)	(0.92)	(3.26)	
2	0.69	0.21	0.09	**0.74**	0.12
	(1.75)	(1.17)	(0.69)	(2.45)	
4	1.21	0.27	0.15	**1.37**	0.19
	(1.79)	(0.93)	(0.86)	(2.46)	
8	-0.11	0.81	0.29	**2.93**	0.31
	(-0.13)	(2.22)	(1.09)	(4.86)	
12	-0.32	0.98	0.18	**3.67**	0.37
	(-0.27)	(2.64)	(0.64)	(4.81)	
16	0.74	0.63	0.04	**3.48**	0.31
	(0.66)	(1.81)	(0.17)	(3.53)	
20	-0.71	0.69	0.04	**3.53**	0.25
	(-0.78)	(1.77)	(0.18)	(4.06)	

Notes for Table 3: The table reports output from long-horizon regressions of consumption growth and asset wealth growth on lags of these variables and the cointegrating error for consumption, wealth and labor income, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$. The dependent variables in the h-period regressions are $\Delta x_{t+1} + ... + \Delta x_{t+h}$, where $x \in \{c, a\}$. For each regression, the table reports OLS estimates of the regressors, Newey-West corrected t-statistics (in parentheses) and adjusted R^2 statistics. Significant coefficients at the 5% level are highlighted in bold face.

Table 4: Consumption Growth Regressions

Variable	Estimate	t-Stat
Panel A: Error-Correction Representation		
$\Delta c_t = k_A + \gamma_c \widehat{\alpha}' \mathbf{x}_{t-1} + \mathbf{\Gamma}(\mathbf{L})\mathbf{\Delta}\mathbf{x}_{t-1} + \mathbf{v}_t$		
k	**0.004**	5.272
$\widehat{\alpha}'\mathbf{x}_{t-1}$	-0.041	-0.990
Δc_{t-1}	0.125	1.553
Δc_{t-2}	**0.210**	2.214
Δa_{t-1}	**0.114**	3.289
Δa_{t-2}	0.023	0.777
Δy_{t-1}	0.018	0.205
Δy_{t-2}	-0.011	-0.138
$\overline{R}^2$	0.182	
SEE	0.006	
SSR	0.006	
Panel B: DP Consumption Growth Regression		
$\Delta c_t = k_B + \gamma_c \widehat{\alpha}' \mathbf{x}_{t-1} + \mathbf{\Gamma}(\mathbf{L})\mathbf{\Delta}\mathbf{x}_{t-1} + \mathbf{C}(\mathbf{L})\mathbf{z}_{t-1} + \mathbf{u}_t$		
k	-0.004	-1.823
$\widehat{\alpha}'\mathbf{x}_{t-1}$	**-0.128**	-2.769
Δc_{t-1}	-0.099	-1.403
Δa_{t-1}	**0.063**	2.322
$z_{1,t}$	**0.261**	2.146
$z_{2,t}$	-0.000	-0.328
$z_{3,t}$	**-0.069**	-2.595
$z_{4,t}$	**0.000**	3.904
$\overline{R}^2$	0.446	
SEE	0.005	
SSR	0.004	
Panel C: Consumption Growth Regression with Orthogonal Regressors		
$\Delta c_t = k_C + \gamma_c \widehat{\alpha}' \mathbf{x}_{t-1} + \mathbf{\Gamma}(\mathbf{L})\mathbf{\Delta}\mathbf{x}_{t-1} + \widetilde{\mathbf{C}}(\mathbf{L})\widetilde{\mathbf{z}}_{t-1} + \tilde{\mathbf{u}}_t$		
k	**0.005**	6.774
$\widehat{\alpha}'\mathbf{x}_{t-1}$	-0.005	-0.108
Δc_{t-1}	0.097	1.400
Δc_{t-2}	**0.160**	2.361
Δa_{t-1}	**0.133**	5.488
Δa_{t-2}	0.037	1.586
Δy_{t-1}	0.005	0.082
Δy_{t-2}	-0.004	-0.069
$\tilde{z}_{1,t}$	**0.335**	2.680
$\tilde{z}_{2,t}$	**0.000**	3.127
$\tilde{z}_{3,t}$	**-0.067**	-2.524
$\tilde{z}_{4,t}$	0.001	0.649
$\overline{R}^2$	0.446	
SEE	0.005	
SSR	0.004	

Notes for Table 4: See next page.

Notes for Table 4: The table reports single equation estimates from three regression specifications for Δc_{t+1}. The regression in Panel A is a second-order specification of the standard error-correction representation for Δc_{t+1} assuming the trivariate cointegrating relation given by $\widehat{\alpha}'\mathbf{x}_{t-1} = c_t - \widehat{\alpha}_a a_t + \widehat{\alpha}_y y_t$. Panel B is the specification estimated by Davis and Palumbo (DP) (2001) corresponding to the results reported in their Table 11, row three. $\mathbf{z}_t \equiv (\Delta\widehat{y}_t, DUNEMP_{t-1}, RFF_{t-1}, UNEXP_t)'$ where $z_1 = \Delta\widehat{y}_t$ is the forecasted income growth variable used in Davis and Polumbo (2001); $z_2 = DUNEMP_{t-1}$, the first difference of the unemployment rate; $z_3 = RFF_{t-1}$, the real Fed Funds rate used in Davis and Polumbo (2001); $z_4 = UNEXP_t$, the Michigan Unemployment Expectations Index. The regression in Panel C uses the orthogonalized regressors, $\widetilde{z}_t$, computed in the manner described in the text. Each $\tilde{z}_i$ is the residual from a regression of z_i on the explanatory variables in the Error Correction Representation of Panel A. SEE denotes the standard error of the estimate and SSR is the sum of squared residuals. Significant coefficients at the 5% level are highlighted in bold face.

Table 5: Consumption Growth Regressions

Variable	Estimate	t-Stat
Panel A: Error-Correction Representation		
$\Delta c_t = k_A + \gamma_c \widehat{\alpha}' \mathbf{x}_{t-1} + \mathbf{\Gamma}(\mathbf{L})\mathbf{\Delta}\mathbf{x}_{t-1} + \mathbf{v}_t$		
k	**0.004**	5.272
$\widehat{\alpha}'\mathbf{x}_{t-1}$	-0.041	-0.990
Δc_{t-1}	0.125	1.553
Δc_{t-2}	**0.210**	2.214
Δa_{t-1}	**0.114**	3.289
Δa_{t-2}	0.023	0.777
Δy_{t-1}	0.018	0.205
Δy_{t-2}	-0.011	-0.138
$\overline{R}^2$	0.182	
SEE	0.006	
SSR	0.006	
Panel B: DP Consumption Growth Regression - IV		
$\Delta c_t = k_B + \gamma_c \widehat{\alpha}' \mathbf{x}_{t-1} + \mathbf{\Gamma}(\mathbf{L})\mathbf{\Delta}\mathbf{x}_{t-1} + \mathbf{C}(\mathbf{L})\mathbf{z}^*_{t-1} + \mathbf{u}_t$		
k	-0.001	-0.455
$\widehat{\alpha}'\mathbf{x}_{t-1}$	**-0.116**	-2.424
Δc_{t-1}	**-0.124**	-2.206
Δa_{t-1}	**0.082**	4.899
$z_{1,t}$	**0.248**	2.258
$z_{2,t}$	-0.136	-1.025
$z_{3,t}$	**-0.089**	-5.019
$z_{4,t}$	**0.000**	3.281
SEE	0.005	
SSR	0.003	
Panel C: Consumption Growth Regression with Orthogonal Regressors		
$\Delta c_t = k_C + \gamma_c \widehat{\alpha}' \mathbf{x}_{t-1} + \mathbf{\Gamma}(\mathbf{L})\mathbf{\Delta}\mathbf{x}_{t-1} + \widetilde{\mathbf{C}}(\mathbf{L})\widetilde{\mathbf{z}}^*_{t-1} + \tilde{\mathbf{u}}_t$		
k	0.005	6.741
$\widehat{\alpha}'\mathbf{x}_{t-1}$	-0.004	-0.096
Δc_{t-1}	0.106	1.483
Δc_{t-2}	0.149	2.289
Δa_{t-1}	0.132	5.061
Δa_{t-2}	0.043	2.041
Δy_{t-1}	-0.023	-0.333
Δy_{t-2}	-0.009	-0.171
$\tilde{z}_{1,t}$	0.330	1.932
$\tilde{z}_{2,t}$	0.000	2.239
$\tilde{z}_{3,t}$	-0.073	-2.802
$\tilde{z}_{4,t}$	0.081	0.452
$\overline{R}^2$	0.411	
SEE	0.005	
SSR	0.004	

Notes for Table 5: See next page.

Notes for Table 5: See notes for Table 4. Panel B gives the Davis-Palumbo type specification estimated using instrumental variables so as to produce a consistent estimate of the standard errors (see discussion in the text). $\mathbf{z}_t^* \equiv (\Delta y_t, DUNEMP_{t-1}, RFF_{t-1}, UNEXP_t)'$; the instrument set includes all of the regressors in the equation of Panel B except Δy_t, in addition to four lags of Δc_t, Δy_t, Δa_t, the federal-funds-rate, and the unemployment rate.

Figure 1: Impulse Response in Example 1

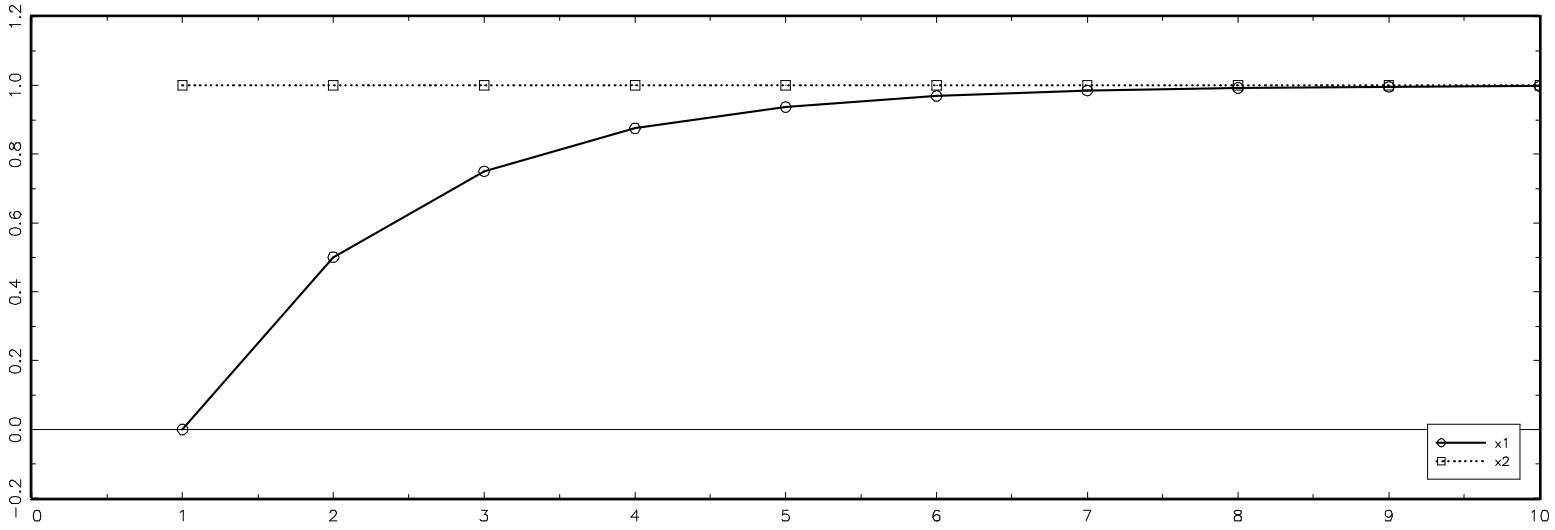
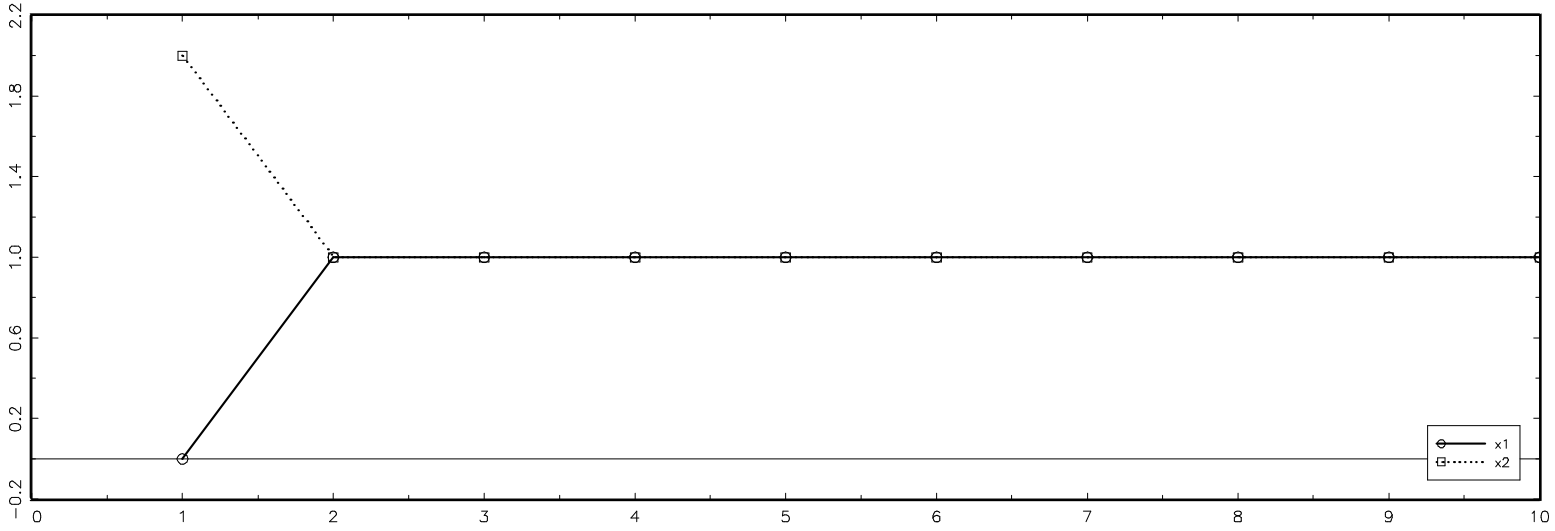
Figure 2: Impulse Response in Example 2