The effects of stock market movements on consumption and investment: does the shock matter?

Stephen Millard

and

John Power

Working paper no. 236

Bank of England, Threadneedle Street, London, EC2R 8AH.
E-mail: stephen.millard@bankofengland.co.uk
john.power@bankofengland.co.uk

The views expressed are those of the authors and do not necessarily reflect those of the Bank of England. The authors would like to thank seminar participants at the Bank of England and London Metropolitan University as well as participants at the Society for Economic Dynamics Annual Conference 2003 in Paris for helpful comments. Any remaining errors are, of course, the fault of the authors.

Copies of working papers may be obtained from Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone 020 7601 4030, fax 020 7601 3298, e-mail mapublications@bankofengland.co.uk.

Working papers are also available at www.bankofengland.co.uk/wp/index.html.

The Bank of England's working paper series is externally refereed.

ISSN 1368-5562

Contents

Abstract

This paper uses a simple model to examine the links between equity price movements and consumption and investment. Generally, the effect of a given movement in equity prices on consumption depends on the underlying source of the shock to equity prices, and some empirical evidence is presented that supports this. Furthermore, in the model the effect of a given movement in equity prices on investment does not depend on the source of the shock. However, some theoretical arguments and empirical evidence are provided to suggest that it might in the real world.

Key words: Consumption, investment, equity prices, VARs.

JEL classification: E44.

Summary

This paper examines the impact of equity prices on consumption and investment. In particular, it considers whether the impact of any given movement in equity prices on consumption and investment depends on the source of the shock that caused the equity price movement. Clearly this is an important topical issue given the sharp falls witnessed in the stock markets a few years ago. In the United Kingdom, for example, the FTSE All-Share index fell by about 40% between September 2000 and the end of 2002. However, much of the literature on consumption and investment equations largely ignores the source of the underlying shock in determining the consumption and investment response; typically, relationships are estimated using equity prices as right-hand side variables without any effort to distinguish the source of the shock to equity prices. In the context of the consumption function, this leads researchers to present estimates of the 'marginal propensity to consume out of wealth'. It would be better to ask the question 'How large are the changes in consumption and wealth following a shock of a particular type?' If the shock does matter and it is not adequately captured in macroeconometric models, this could detract from their forecast performance or give a misleading view about the outlook for consumption.

To answer the above question, it is important first to know what sorts of shocks move equity prices. Theory suggests that asset prices represent the present value of future income to be derived from the underlying asset. So equity prices might move when either expected future income or dividend growth changes, or when the discount factor applied to them changes. The discount factor, in turn, will be the sum of the equity risk premium and the risk-free interest rate. Alternatively, of course, equity prices might not be accounted for by any of those 'fundamental' reasons, but rather reflect irrational responses to market sentiment, or noise.

The paper then approaches the problem from two angles: theoretical and empirical. It develops a simple general equilibrium model that links equity prices to consumption and investment for a small open economy. It then analyses the link between consumption and equity prices and explains why the consumption response is likely to be different depending on the source of the shock. Importantly, the response to a risk premium shock is likely to differ from that to an interest rate shock. Indeed, there are cases where risk premia movements will have no effect on consumption. Next the focus is on the links between investment and equity prices. Unlike consumption, the model suggests that investment will always respond to movements in equity prices irrespective of the source of the shock. But this might not be true in the real world. In particular, movements in equity prices that are unrelated to fundamentals are likely to have a much smaller effect on investment (if any) than those related to fundamentals. Finally, the paper presents some empirical evidence from a vector autoregression model to identify whether the source of the shock matters in the data. That analysis suggests that it does matter for consumption and, contrary to the simple predictions from the model, for investment too.

This paper has still left some questions unresolved. In particular, it would be good to model more explicitly the sorts of shocks that drive the economy and their stochastic processes. Then it would be interesting to investigate whether or not it makes sense, in a more complicated model for shocks to equity price volatility to have no effect on consumption volatility. To do this, one would have to examine the links between the equity risk premium and the volatilities of

fundamental shocks. In turn, that would mean considering models that are able to explain the magnitude of observed equity risk premia, something that is not the case in the simple model presented here. It would also be interesting to investigate whether there are any shocks (or specifications of the model) under which investment would respond more or less to equity price movements than in the baseline case.

1 Introduction and overview

In this paper, we examine the impact of equity prices on consumption and investment. In particular, we consider whether the impact of any given movement in equity prices on consumption and investment depends on the source of the shock that caused the equity price movement. Clearly this is an important topical issue given the sharp falls witnessed in the stock markets recently. In the United Kingdom, for example, the FTSE All-Share index fell by about 40% between September 2000 and the end of 2002. However, on the whole the literature in which consumption and investment 'functions' are estimated largely ignores the source of the underlying shock in determining the consumption and investment response; typically relationships are estimated using equity prices as right-hand side variables without any effort to distinguish the source of the shock. In the context of the consumption function, this leads to researchers to present estimates of the 'marginal propensity to consume out of wealth'. In our view, it would be better to ask the question 'How large are the changes in consumption and wealth following a shock of a particular type?' This is important as if the shock does matter and this is not adequately captured in macroeconometric models, then this could detract from their forecast performance or give a misleading view about the outlook for consumption.

1.1 Equities and consumption

When analysing the effect of stock market wealth on consumption, we normally look at empirical estimates of the marginal propensity to consume out of wealth. The marginal propensity to consume out of wealth measures by how much the value of consumption is expected to rise given a unit increase in wealth. The permanent income hypothesis (Friedman (1957)) implies that any shocks to income, transitory or permanent, (which will lead to a measured rise in wealth) will be consumed over the lifetime of the consumer, since consumers prefer a smoother consumption profile than an erratic one. This result implies a low marginal propensity to consume out of transitory income, roughly related to the annuity value of the change in wealth (shock to transitory income). As such, in the long run, you would expect the marginal propensity to consume to be very similar to the real interest rate.(1) Therefore, a value of around 3% to 5% would be a reasonable prior for the marginal propensity to consume out of wealth.

There is an extensive literature on empirically deriving the marginal propensity to consume out of stock market wealth on household consumption. While this has been largely concentrated in the United States (in part driven by the desire to explain consumption behaviour over the mid to late 1990s), recent estimates for the United Kingdom and other countries have been compiled by, among others, the OECD and IMF. In the United Kingdom, Bertaut (2002) found that the long-run marginal propensities to consume for total, financial and housing wealth are all around 0.04, or a four pence increase in consumption for an increase of one pound in wealth. Ludwig and Sløk (2001) estimate the long-run marginal propensity to consume out of equity wealth for the United Kingdom to be around 0.055, or a 5.5 pence increase in consumption for an increase of one pound in wealth. Boone, Girouard and Wanner (2001) find that the long-run marginal

(1) With the possibility that it may be a little higher if you relax the assumption of infinitely lived agents and a little lower to the extent that credit market imperfections limit the extent to which agents can smooth consumption.

propensity to consume for equity and housing wealth in the United Kingdom to be the same as Bertaut, at around 0.04. All these estimates are close to estimates of the real interest rate.

1.2 Equities and investment

Estimates of the effects of equity prices movements on investment have been based on the neoclassical model, which can be used to specify investment as a function of either the user cost or *q*.

Firms' desired capital stocks are determined not only by their planned production levels, but also by the real user cost of capital, that is, the cost faced by the firm in acquiring and holding an additional unit of capital. One important component of the user cost is the cost of finance, a weighted average of the cost of equity finance and other forms of finance. Rises in equity valuations unrelated to current or future dividend payouts are associated with falls in the cost of equity finance and should be a positive influence on firms' investment decisions.

The literature on investment and the cost of capital is vast and has already been reviewed extensively (eg Chirinko (1993) and Caballero (1999)). Although the theoretical foundations of the cost of capital and the interaction with investment are well established, empirical results have been less encouraging. Blanchard (1986) summarises that 'it is well known that to get the user cost of capital to appear at all in the investment equation, one has to display more than the usual amount of econometric ingenuity', although Ellis and Price (2003) find a well defined and sensible effect on the most recent vintage of aggregate UK business data. Caballero (1999) agrees that the short-run response of investment to changes in the cost of capital is complex and depends on the shock to the cost of capital. But he concludes that the cost of capital is nevertheless an important variable, determining the long-run relationship between capital, output and the cost of capital.(2)

Q is derived from the same theory. At its simplest, when the market value of an additional unit of capital exceeds it replacement cost, a firm can increase profits by investing. Given that this ratio, *marginal q*, is difficult to measure, economists have focused on Tobin's (1969) *average q*: the ratio of the stock market value of a firm to the replacement cost of its assets. Under strict assumptions, average *q* is a sufficient statistic for investment rates.

But *q* has also generally fared poorly in empirical studies of its predictive power for investment. First, *q* models usually imply that it would take a very long time – often decades – for firms' capital stock to fully adjust to a change in *q*. Second, even though *q* should theoretically be a sufficient statistic for investment, a large literature has found cash-flow variables to be significant in *q* regressions.(3) This has been widely interpreted as suggestive of the existence of financing constraints. The poor performance of *q* may also partly reflect measurement issues. Studies have typically focused on the role of average *q*, which is only equal to marginal *q* under restrictive

---

(2) Caballero (1999) argues that because of adjustment costs the long-run coefficient in neoclassical investment equations is biased downward in finite samples. The basic idea is that adjustment costs result in the observed capital-output ratio to be less volatile than the desired capital-output ratio and the cost of capital.
(3) See Fazzari, Hubbard and Petersen (1988) for early evidence in this area.

assumptions. And average *q* itself may be poorly measured if estimates of the capital stock are unreliable.

To get around this, Bond and Cummins (2001) followed an alternative firm-level approach. They calculated a 'fundamental *q*' based on analysts' earnings forecasts from the Institutional Brokers Estimate System (IBES), and found that when this measure was included, a simple stock market-based *q* does not have any explanatory power for investment.(4) In a similar vein, Bond *et al* (2004) found that when they controlled for expected profitability using IBES forecasts, conventional *q* measures provided little additional information.(5)

1.3 Our approach

But most of this empirical work ignores the key question: can we be confident in using such equations to predict consumption and investment when we know that equity prices will be driven by different shocks in different time periods? The dividend discount model (DDM) offers a framework which can isolate the factors or shocks that affect equity prices. Theory suggests that asset prices represent the present value of future income to be derived from the underlying asset. So we might expect equity prices to move when either expected future income or dividend growth changes (which we might expect to be closely related to firms' productivity), or when the discount factor applied to them changes. The discount factor, in turn, will be the sum of the equity risk premium and the risk-free interest rate. Alternatively, in this framework, equity prices might not be accounted for by any of those 'fundamental' reasons, rather reflecting irrational responses to market sentiment, or noise. We note however that this is just one way of characterising the shocks that can affect equity prices. In the real world, it is likely that there is a wider range of shocks driving equity prices. And the impact of those underlying shocks on the variables in the DDM could be complex. For example an oil price shock could affect both the equity risk premium (as a result of the uncertainty it generates), and future dividend growth (by affecting firms' future costs of production).

There has been little research along these lines. Lantz and Sarte (2001) have a related paper which addresses the question of how varying the shock to wealth results in varying responses in consumption, but with the emphasis on whether shocks are anticipated, a rather different emphasis from that taken in our paper. The authors use a closed-economy model in which there is only one shock: productivity. They show that consumption and wealth can actually move in different directions depending on whether the shock is anticipated or not. An unanticipated fall in the level of productivity causes consumption and wealth to fall when the shock is realised. However, with an anticipated rise in future productivity, consumption increases immediately while wealth falls. This happens since, for consumption to rise immediately, consumers in a closed economy must dispose of some of their capital. And in response, the value of equity falls, at least initially. An important insight from Lantz and Sarte's paper is that it is inappropriate to

---

(4) It is likely that IBES forecasts are biased upwards, as suggested by Chan, Karceski and Lakonishok (2001), although it is not obvious that the size of this bias will vary markedly over time.
(5) Cash flow is insignificant under this specification, suggesting that cash-flow variables widely used in investment equations capture information about firms' future profitability rather than financing constraints.

consider a general consumption response to changes in wealth (or indeed for changes in wealth to 'cause' changes in consumption). Rather wealth and consumption are jointly determined: both react to various disturbances in the economy.

On the investment side, Lettau and Ludvigson (2002) have noted that changes in current stock returns (and equity prices) which also affect future discount rates, would have implications for both current and future investment. Under that condition (which they believe is valid for US data), a fall in the equity risk premium could stimulate investment growth in the near term but reduce it over the longer run.

In this paper, we approach the problem from two angles – theoretical and empirical. In Section 2 we develop a simple general equilibrium model of equity prices of consumption and investment for a small open economy (in contrast to the closed-economy model used by Lantz and Sarte (2001)). In Section 3 we analyse the link between consumption and equity prices and conclude that the consumption response is likely to be different depending on the source of the shock. Importantly, the model implies that the response from a risk premium shock is likely to differ than that from an interest rate shock. Indeed, in the model there are cases where risk premia movements will have no effect on consumption. Section 4 outlines the links between investment and equity prices. Unlike consumption, the model suggests that investment will always respond to movement in equity prices irrespective of the source of the shock. In Section 5 we present some empirical evidence from a VAR which tries to identify whether the source of the shock matters in the data. That analysis implies that it does matter for consumption and, contrary to the simple predictions from the model, for investment too.

2 The model

We use a fairly standard, small open-economy model such as can be found in, say, Obstfeld and Rogoff (1996). The model is simple and highly stylised: we do not suggest that our model is at all representative of all the channels linking equity prices to consumption and investment in the United Kingdom. However it is rich enough to distinguish between shocks to interest rates, productivity and risk premia and how they might affect consumption/investment.

2.1 Consumers

The representative consumer either consumes or invests in financial assets in order to maximise his utility subject to an asset accumulation constraint. The consumer can invest in either a risky asset or an internationally traded, risk-free asset. We assume that he can borrow or lend as much as he wants on the world capital markets at the prevailing world real interest rate (subject to a 'No Ponzi Game' condition). This assumption amounts to assuming that the consumer lives in a small open economy. It is access to the world capital markets that allows consumers to insure themselves against movements in domestically generated income and, in turn, this drives the key result that shocks to domestic productivity will have different effects on consumption than shocks to the world real interest rate.

Mathematically, we can write his problem as:

Maximise $$E_0 \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\tau} - 1}{1-\tau} \quad \textbf{(1)}$$

Subject to: $$V_t x_t + B_t = (V_t + D_t)x_{t-1} + (1 + r_{t-1})B_{t-1} + w_t - c_t \quad \textbf{(2)}$$

where c is consumption, x is the consumer's end-of-period holdings of equity shares, V is the real value of a share in the domestic firm, D are the real dividends paid on each share, B is the consumer's holdings of real bonds, r is the world real interest rate and w is real labour income. We assume that he takes dividends, real labour income, share prices and the world real interest rate as given. Further, we have assumed that domestic consumers are unable to hold foreign equity and *vice versa*.

2.2 Firms

We assume that the representative firm operates in a perfectly competitive world market and that the domestic economy forms a negligible fraction of the world (ie it is 'small'). That implies that the domestic firm can always sell its entire output at the world market price. As the representative consumer owns the representative firm, the firm's problem will be to maximise the present discounted utility value of its current and expected future dividends. It does this by choosing the optimal paths for its inputs of capital and labour.

If we let λ denote the marginal utility of a unit of consumption, then its problem will be:

Maximise $$E_0 \sum_{t=0}^{\infty} \beta^t \lambda_t D_t \quad \textbf{(3)}$$

Subject to $$D_t = y_t - \left(I_t + \frac{\chi}{2} \frac{I_t^2}{k_{t-1}} \right) - w_t h_t \quad \textbf{(4)}$$

$$y_t = A_t k_{t-1}^{\alpha} h_t^{1-\alpha} \quad \textbf{(5)}$$

$$k_t = I_t + (1-\delta)k_{t-1} \quad \textbf{(6)}$$

where A is a productivity shock, h is labour input, which will equal unity in equilibrium since we have assumed that labour is inelastically supplied, I is investment and k is the end-of-period capital stock. Note that the quadratic term in **(4)** implies that the firm faces costs of adjusting its capital stock.

2.3 *Exogenous variables*

Finally, we assume that the exogenous variables in the model – world real interest rates and productivity – follow known stochastic processes that are independent of each other.[6]

3 Equities and consumption

In this section, we examine the links between the determinants of equity prices and consumption and the effects of a temporary interest rate shock, temporary and permanent productivity shocks and shocks to the determinants of the equity risk premium. We first consider the consumer's decision to invest in the risk-free bond. This simply implies that:

$$c_t^{-\tau} = \beta(1+r_t)E_t c_{t+1}^{-\tau} \quad \textbf{(7)}$$

and assuming that consumption growth is log-normally distributed, we can rearrange equation **(7)** to obtain:

$$E_t\left(\ln\left(\frac{c_{t+1}}{c_t}\right)\right) = \frac{r_t + \ln(\beta) + \frac{1}{2}\tau^2\sigma_c^2}{\tau} \quad \textbf{(8)}$$

Here σ_c^2 is the variance of consumption growth.

Despite their relative simplicity, in this model these equations yield important results. Most notably they determine consumption growth. In the absence of uncertainty, consumption growth will be determined by real interest rates **(7)**.

Under uncertainty, consumption growth will also be affected by the volatility of consumption outturns; as consumption growth becomes more volatile, individuals might wish to save more for any given real interest rates.

Equation **(8)** shows that an increase in the risk-free rate (which would be expected to lead to a fall in equity prices) would cause consumption to fall today relative to tomorrow. That reflects the substitution effect: consumption today becomes more costly relative to tomorrow.

A negative productivity shock, if it persists into the next period, lowers equity prices in this model. If the shock is permanent, the level of output will be permanently lower; this will be passed on to consumers as permanently lower income and, hence, they will have to permanently consume less. If the shock is temporary, then initially, consumption will fall by less than output resulting in a trade deficit; this deficit can only be financed by borrowing from abroad (selling real bonds to foreigners). As consumers must pay future interest payments on the bonds, that

---

(6) We might expect shocks to the world real interest rate to be related to shocks to world productivity and, hence, domestic productivity. But we assume that the diffusion of productivity shocks between countries takes a substantial amount of time so that movements in domestic productivity follow that of world productivity (and world interest rates) at a long lag and *vice versa*.

implies that consumption has to be permanently lower than income (and the initial level of consumption) in the future.[7] This story is identical to that put forward by Friedman (1957) in his original exposition of the permanent income hypothesis, via a permanent negative productivity shock lowers the consumer's permanent income and, hence, permanent consumption; whereas a temporary negative productivity shock will lower the consumer's transitory income, and only lower permanent consumption by its amortised value.

In the case of a permanent shock to productivity, equity prices will be permanently lower; in the case of a temporary shock to productivity equity prices will recover to their initial trend as output recovers. So, clearly, the relationship between the initial falls in consumption and equity prices will depend on whether the shock is temporary or permanent. In addition, the fall in consumption will not necessarily be the same as the fall brought about by higher interest rates and we can recall that for an interest rate shock the fall is temporary with consumption higher than its initial trend in the future.

We next consider equities more specifically. The first-order condition for the consumer's desired equity holdings will be:

$$V_t c_t^{-\tau} = \beta E_t \left(c_{t+1}^{-\tau} \left(V_{t+1} + D_{t+1} \right) \right) \qquad \textbf{(9)}$$

If we assume that consumption growth and equity returns are jointly log-normally distributed, we can combine equations **(9)** and **(7)** to get:

$$E_t \left(\frac{V_{t+1} + D_{t+1}}{V_t} \right) = 1 + r_t + \tau Cov_t \left(\frac{V_{t+1} + D_{t+1}}{V_t}, \frac{c_{t+1}}{c_t} \right) \qquad \textbf{(10)}$$

In words, the return on equities depends on the risk-free rate and an equity risk premium. In equilibrium, equities are priced (the required return is set) as a function of the covariance of the yield with marginal utility. When equity returns are covarying positively with consumption growth (negatively with marginal utility) then consumers are requiring a return higher than the risk-free rate. That reflects risk-averse consumers' desire to maintain a smooth consumption stream: equities which give investors a high payoff when consumption is high (and marginal utility low) are less desirable than those which are expected to be a good hedge against poor consumption outturns. Hence they require a higher rate of return.

We can further break down the equity risk premium as follows:

---

[7] We can note, again, that the small open-economy assumption implies that the fall in consumption will be less than the fall in output. If the consumer did not have access to the world capital markets or could only access them at a premium, then – for a given impact on equity prices – the effect of a temporary productivity shock on consumption would be much larger.

$$ERP_t = \tau Cov_t\left(\frac{V_{t+1}+D_{t+1}}{V_t}, \frac{c_{t+1}}{c_t}\right)$$
$$= \tau StDev_t\left(\frac{V_{t+1}+D_{t+1}}{V_t}\right) StDev_t\left(\frac{c_{t+1}}{c_t}\right) Corr_t\left(\frac{V_{t+1}+D_{t+1}}{V_t}, \frac{c_{t+1}}{c_t}\right) \quad \textbf{(11)}$$

This equation enables us to make a simple point: an increase in the degree of risk aversion, the variance of consumption growth, the variance of equity returns and/or the correlation between equity returns and consumption growth would lead to an increase in the equity risk premium. But equation **(8)** shows that in a two-period model, current consumption will only fall – as a result of precautionary saving – if the degree of risk aversion or the variance of consumption rises. A rise in equity return volatility that was not accompanied by a rise in consumption volatility would have no effect on consumption.

What sort of shock could produce this? In the model, as laid out in Section 2, the variances of equity prices and consumption will depend on the variances of the underlying shocks to world interest rates and productivity. An increase in the variance of either shock will increase the variance of both equity prices and consumption. But, thinking outside the model, suppose equity returns are partly driven by 'noise' and this is common knowledge. Examples of this could be thin markets at certain times of the day or certain days of the year or particular trades that were large enough to drive prices for a short period of time. Suppose that the variance of this 'noise' increases. In this case, we might again expect the equity risk premium to increase but without any effect on the volatility of consumption growth and, hence, on the level of consumption, although in the absence of a formal model this is speculation.

To conclude this section, we have shown that the effects of a given movement in equity prices on consumption depend critically upon the source of the underlying shock. In our simple model, a temporary positive productivity shock will raise equity prices but have no effect on consumption. A negative shock to the real interest rate would raise equity prices and lead to a rise in the level of consumption from the wealth effect and a fall in the growth rate of consumption from the substitution effect. A fall in the equity risk premium that was not purely due to less volatile share prices would raise consumption today by decreasing precautionary savings. And, finally, any remaining movement in equity prices unrelated to 'fundamentals' might not have any effect on consumption.(8)

4 Equities and investment

In this section, we analyse the links between equity price movements and investment, again paying particular attention to the extent to which a given movement in equity prices may have different implications for investment depending on the source of the shock. We start by noting that the firm's first-order conditions imply:

---

(8) Although, arguably, if equity prices are driven purely by sentiment (which subsequently proved to be unfounded), consumption might be affected through its effects on consumer confidence. That channel is not modelled in this paper.

$$I_t = \frac{q_t - 1}{\chi} k_{t-1} \qquad \textbf{(12)}$$

and

$$\lambda_t q_t = \beta E_t \left(\lambda_{t+1} \left(\frac{\chi}{2} \left(\frac{I_{t+1}}{k_t} \right)^2 + \frac{\alpha y_{t+1}}{k_t} + q_{t+1}(1-\delta) \right) \right) \qquad \textbf{(13)}$$

where q is the Lagrange multiplier on the capital accumulation equation and, in our model, will be equivalent to Tobin's q.

If we multiply both sides of equation **(13)** by k and combine the result with equation **(4)**, we get:

$$\lambda_t q_t k_t = \beta E_t \lambda_{t+1} (q_{t+1} k_{t+1} + D_{t+1}) \qquad \textbf{(14)}$$

and comparing this to the consumer's first-order condition shows that:

$$V_t = q_t k_t \qquad \textbf{(15)}$$

In other words the value of the stock market, V, is equal to the value of the capital stock multiplied by Tobin's q.

Now, combining equations **(6)**, **(12)** and **(15)** gives us:

$$\frac{\chi}{V_{t-1}} I_t^2 + \left(1 + (1-\delta) \frac{\chi}{\chi I_{t-1} + V_{t-2}} \right) I_t + (1-\delta) \frac{V_{t-1} V_{t-2}}{\chi I_{t-1} + V_{t-2}} = V_t \qquad \textbf{(16)}$$

So, we can see that unlike the result we obtained for consumption, a given movement in equity prices will be associated with a given movement in investment irrespective of the source of the shock:

$$\frac{dI_t}{dV_t} = \frac{1}{\frac{2\chi I_t}{V_{t-1}} + \left(1 + (1-\delta) \frac{\chi}{\chi I_{t-1} + V_{t-2}} \right)} \qquad \textbf{(17)}$$

Q is the shadow price of capital, embodying the relevant marginal conditions in a single number, and consequently is a sufficient statistic to tell us whether marginal investment is profitable. A change in q caused by a change in equity prices causes that marginal profitability to change: clearly, the source is irrelevant. So once we know what has happened to stock prices, and hence q, we will know what has happened to investment.(9) But, there are, of course, other shocks that could affect equity prices that are not captured by this simple model. A shift in firms' mix of

(9) Typically, the neoclassical model is also extended to consider the role of output in determining investment. In accelerator models, higher output leads to higher investment (and often there is a role for output in empirical estimates of the investment equation). Although there is no output term in equation **(16)**, the simple q theory is consistent with the accelerator model: an increase in output should raise the market value of the firm's equity (capital cannot adjust instantaneously to higher demand for the firm's product) thereby increasing q and stimulating investment.

debt and equity finance might change the cost of equity without any change in the cost of capital or investment, for example, under the Modigliani-Miller conditions when variations in the mix of debt to equity leave the value of the firm and all real decisions, including investment, unaffected. The model ignores 'financial accelerator' type effects resulting from the change in corporate balance sheets that movements in equity prices bring about. For example, by reducing firm's collateral, or net worth, a large fall in equity prices might lead to a reduction in lending by financial institutions to firms. That might act to restrain investment for credit-constrained firms. Finally, the source of the shock to equity prices may matter if company managers respond differently to changes they perceive as permanent rather than transitory, or if they believe equity price changes reflect 'fundamentals' or noise.

MacGorain and Thompson (2002) examined this last point. Consistent with simple investment theory, they hypothesised that movements in expected future profits (productivity in our model) would be likely to affect both share prices and investment. But they argued that, at times, the expectations of a firm's future earnings implied by the firm's share price may differ from managers' own opinions about their firm's future profitability. They calculated a measure of q based on analysts' forecasts of earnings per share growth (as provided by IBES) arguing that this was a good indicator of managers' opinions. They then estimated a simple aggregate econometric equation for investment including both Tobin's q and this 'analysts' q'. They found that to the extent that the equation could explain variations in the rate of investment, it did so almost entirely by the q measure based on analysts' earnings forecasts rather than the stock market. This suggests that movements in equity prices unrelated to fundamentals have little effect on investment. Their results corroborate the firm-level results of Bond *et al* (2004) reported in the introduction.

To conclude this section, we have shown that the effects of a given movement in equity prices on investment do not, in a standard model, depend upon the source of the underlying shock. However, that ignores some theoretical arguments and empirical evidence to the contrary.

5 Some empirical evidence

Given the theoretical predictions of the effect of the source of the equity price shock on consumption and investment, it is useful to examine those conclusions within an identified vector autoregression (VAR) model for the United Kingdom. This also allows us to work through the dynamics of the change in consumption and investment. To do this we used a VAR in which we identified a transitory shock to world interest rates, a permanent shock to domestic productivity and a transitory shock to the equity risk premium. More specifically, we estimated a VAR containing the following variables: world three-month nominal interest rates (constructed as a GDP-weighted average of non-UK G7 rates), i_f, consumption or investment (depending on which variable we are considering), C or I, retail prices (as measured by RPIX), P, the UK bank base rate, i, FTSE All-Share dividends, D, and the FTSE All-Share index, V. We used quarterly data from 1978 Q1 to 2002 Q1.[10]

In our VAR, we assumed that world interest rates were exogenous; we used the VAR to calculate the response of consumption (investment) and equity prices to a shock to world interest rates.[11] By dividing the response of consumption (investment) by the initial response of equity prices, we obtained the consumption (investment) response to a unit 'shock' to equity prices where the actual source of the shock was world interest rates. In doing this we were asking ourselves the question, 'What would happen to consumption/investment over the next few quarters if we see equity prices move by 1% today as a result of a world interest rate shock?' Over time, the response of equity prices themselves will differ depending on the source of the shock – meaning that the ratio of the consumption/investment response to the equity price response will not be equal to that shown in our charts – but this does not detract from the point we wish to make in this paper that a given observed movement in equity prices will be associated with different responses of consumption and investment depending on the source of the shock.

To identify the domestic productivity shock we used the approach of Christiano *et al* (2003). We first note that, if we are prepared to assume that shocks to the world real interest rate and the equity risk premium are temporary, our model implies that the only shock that can permanently affect the level of output, investment and the capital stock is a permanent productivity shock. So, consider the relationship:

---

(10) There has been much debate among researchers working on estimating consumption functions as to what are the best variables to use in the empirical work – in particular, whether or not to use total labour income or disposable labour income, total consumption or consumption of non-durables – and how to define wealth. (See Blake, Fernandez-Corugedo and Price (2003) for a much fuller description of these issues.) However, this is not an issue in this paper since we are not concerned about explaining consumption behaviour *per se*, we are simply trying to identify different shocks.

(11) Ideally we would like to examine the effects of shocks to world real interest rates so as to tie the empirical work in more closely with the theoretical model. However, there are several problems associated with obtaining the data required to do this calculation. One way of doing this is to construct a world price index, use this series to derive world inflation expectations and subtract this series from the world nominal interest rate used in our work. Problems with this approach include the fact that individual country's price indices are often not directly comparable, converting them into a common currency is not straightforward – do you use market exchange rates or some estimate of PPP – and then the estimates of inflation expectations will rely on the assumed process driving the world price level. Calculating a real interest rate country-by-country and then averaging gets around some of these problems but not all. In particular, the question of how inflation expectations are formed still remains. For a much fuller discussion of these issues see Jenkinson (1996).

$$\Delta I_t = g + \beta\Delta I_{t-1} + \alpha(L)\mathbf{X}_t + \varepsilon_{A,t} \quad \textbf{(18)}$$

where $\alpha(L)$ is a second-order lag-polynomial and **X** contains the world interest rate, consumption, the price level, the domestic interest rate, dividends and equity prices. Assuming that non-technology shocks have no effect on the level of investment in the long run implies $\alpha(1) = 0$. This restriction can be used to estimate a time series for the productivity shock, $\varepsilon_{A,t}$. We calculated the response of consumption (investment) and equity prices to this shock and, using the same method as before, went on to calculate the consumption (investment) response to a unit 'shock' to equity prices where the actual source of the shock was productivity.

Finally, to identify the equity return shock, we draw on the Campbell and Shiller (1988) dynamic version of the DDM described earlier. By log-linearising a simple expression for equity returns, we can write the current equity price, V, as the present value of expected future dividends, D, and returns, r:

$$\hat{V}_t = E_t \sum_{j=0}^{\infty} \beta^j [(1-\beta)\hat{D}_{t+j+1} - (r_{t+j+1} - r)] + \kappa \quad \textbf{(19)}$$

where 'hat' variables represent log-deviations from trend and r is the steady-state level of returns. It is frequently argued that dividends follow a random walk and so are not conditionally forecastable. Using this assumption, we can take the dividend term in equation **(19)** out of the expectation. That implies that equity price movements reflect shocks to future dividends and expected future returns. We assume that the FTSE All-Share index responds to shocks to all the other variables within period. And assuming that shocks to expected returns account for almost all movement in the equity price, we can interpret the errors in the equation for the FTSE All-Share as a shock to expected equity returns, that is, the equity risk premium. But, even if we believe that dividends are conditionally forecastable, it seems likely that the other information in the system – consumption/investment, the price level, world and domestic interest rates, and dividends – should capture most of the predictable component. Given these assumptions, we calculated the response of consumption (investment) and equity prices to our identified equity risk premium shock and, using the same method as before, went on to calculate the consumption (investment) response to a unit equity price 'shock'.(12)

We summarise these restrictions as follows. The structural moving average representation for our VAR is:

$$\mathbf{x_t} = \mathbf{C}(L)\varepsilon_t \quad \textbf{(20)}$$

where $\mathbf{x}_t = \left(i_{F,t} - i_F \quad \hat{C}_t \text{ or } \hat{I}_t \quad \hat{P}_t \quad i_t - i \quad \hat{D}_t \quad V_t\right)'$, ε is our vector of shocks, i_F is the average level of foreign interest rates, i is the average level of domestic interest rates and hat variables refer to log-deviations from trend. C(L) is a matrix lag polynomial:

(12) As outlined earlier equity prices can move for reasons unrelated to fundamentals. In our framework that noise would also be captured within the equity risk premium.

$$\mathbf{C}(L) = \mathbf{C0} + \mathbf{C1}L + \mathbf{C2}L^2 + \ldots \tag{21}$$

If we let our world interest rate shock be the first element of ε, the productivity shock be the second element of ε and the equity risk premium shock be the third element of ε, then our restrictions imply the following:

$$\mathbf{C0} = \begin{pmatrix} x & 0 & 0 & 0 & 0 & 0 \\ x & x & 0 & x & x & x \\ x & x & 0 & x & x & x \\ x & x & 0 & x & x & x \\ x & x & 0 & x & x & x \\ x & x & x & x & x & x \end{pmatrix} \tag{22}$$

$$\mathbf{C1}_{1,2-6} = \mathbf{C2}_{1,2-6} = \mathbf{C3}_{1,2-6} = \ldots = 0 \tag{23}$$

$$\mathbf{C}(1)_{1,2} = \mathbf{C}(1)_{3,2} = \mathbf{C}(1)_{4,2} = \mathbf{C}(1)_{5,2} = \mathbf{C}(1)_{6,2} = 0 \tag{24}$$

Charts 1 to 3 plot the responses of consumption to a 1% rise in equity prices where this rise is caused by a world interest rate shock, a permanent productivity shock and an equity risk premium shock, respectively.

Chart 1: Response of consumption to a world interest rate shock

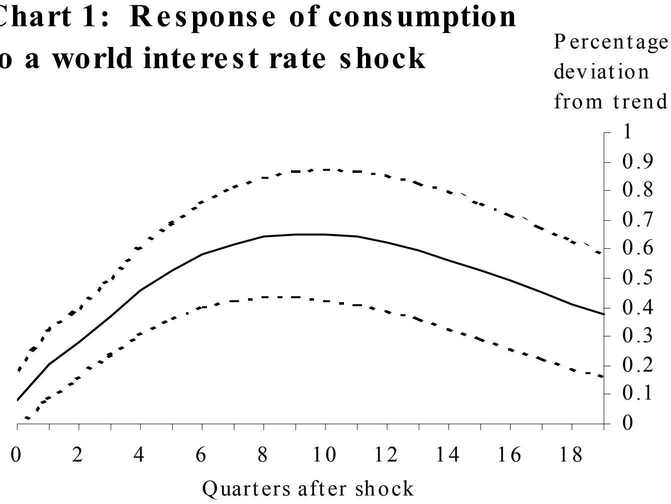

Note: The shock is such that it would result in an immediate 1% rise in equity prices.

Chart 2: Response of consumption to a productivity shock

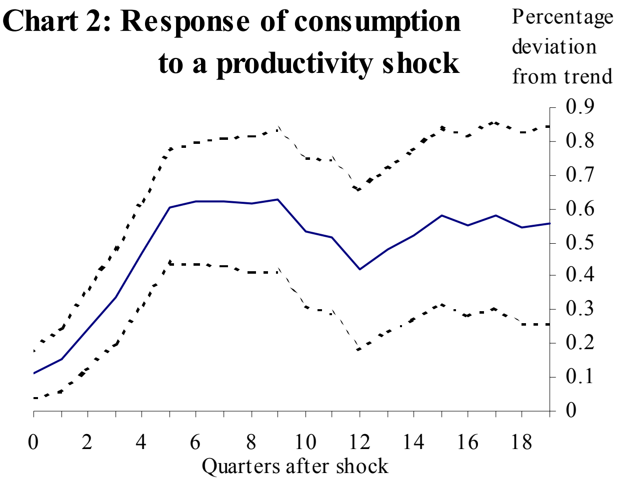

Note: The shock is such that it would result in an immediate 1% rise in equity prices.

Chart 3: Response of consumption to an equity risk premium shock

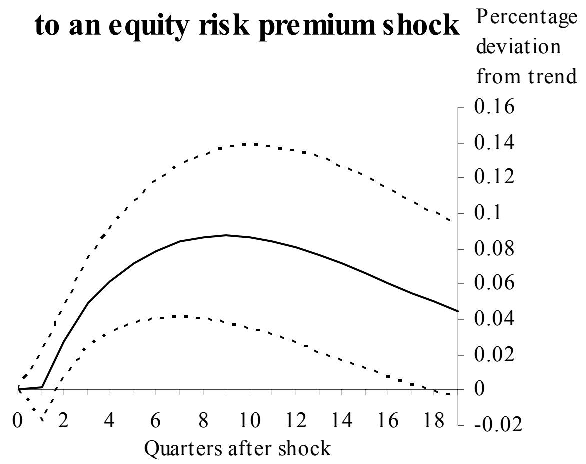

Note: The shock is such that it would result in an immediate 1% rise in equity prices.

Shocks to world interest rates or the equity risk premium have a qualitatively similar effect on consumption. Consumption rises for about nine quarters before falling back towards trend. But the interest rate shock has a much bigger impact on consumption being approximately eight times larger than that from the equity risk premium. This is in line with the discussion of Section 3, which suggested that movements in the equity risk premium could, at times, be associated with no movement in consumption. The immediate response of consumption to a productivity shock is similar in magnitude to the response to a world interest rate shock but after about a year consumption flattens out staying about 0.6% higher than its previous trend. The fact that this shock results in permanently higher output means that domestic consumers can afford to permanently raise their consumption.

Charts 4 to 6 show the effects on investment of a 1% rise in equity prices brought about by shocks to world interest rates, productivity and the equity risk premium, respectively.

Chart 4: Response of investment to a world interest rate shock

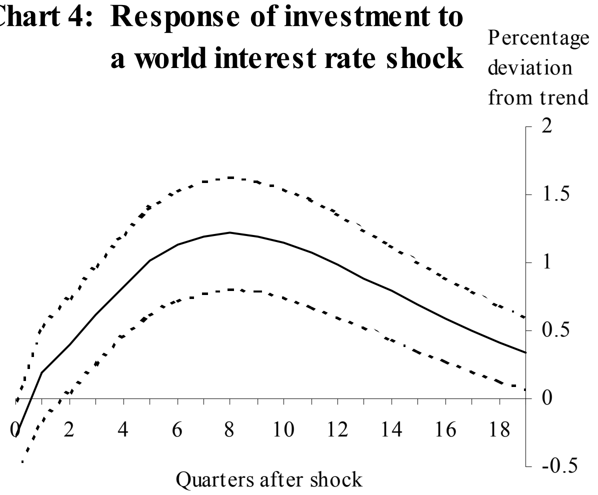

Note: The shock is such that it would result in an immediate 1% rise in equity prices.

Chart 5: Response of investment to a productivity shock

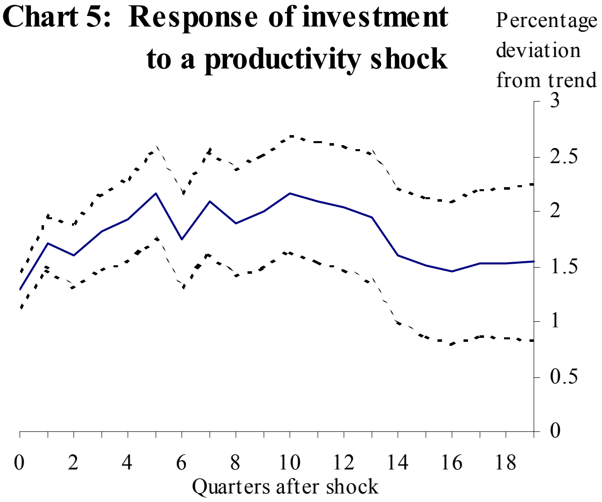

Note: The shock is such that it would result in an immediate 1% rise in equity prices.

Chart 6: Response of investment to an equity risk premium shock

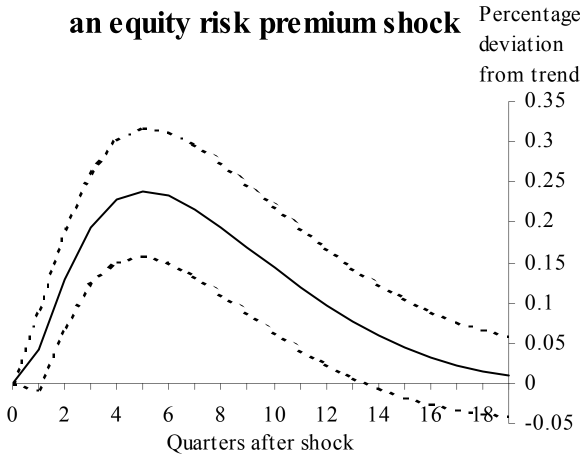

Note: The shock is such that it would result in an immediate 1% rise in equity prices.

The response of investment to each of these shocks is qualitatively similar to that of consumption. In the case of a shock to world interest rates or the equity risk premium, investment rises initially before subsequently falling back to trend. However, there are some differences between the responses of investment to these two shocks. A shock to world interest rates leads investment to fall initially, then to rise for about two years before falling back to trend. A shock to the equity risk premium has, by assumption, no immediate effect on investment; investment then rises for only five quarters before falling back to trend. Again, the magnitude of the response to an equity risk premium shock is much smaller – about a fifth – than that to a world interest rate shock. This goes against the results of our theoretical model. However, the empirical results might also reflect the fact that the equity risk premium shock also picks up movements in equity prices unrelated to fundamentals. As we argued earlier, it is possible that investment does not respond at all to movements in equity prices that are seen as noise. And the results of MacGorain and Thompson (2002) lend further evidence to this suggestion. In the case of the productivity shock, investment jumps immediately to a level 1.3% above its initial trend and settles around 1.5% higher than its initial trend. This is a much stronger response than to either of the other two shocks, and that is shown by consumption to the same shock. Again, the productivity shock has led to a permanent rise in investment as output and the capital stock will be permanently higher.

6 Conclusions

In this paper we used a simple model to examine the links between equity price movements, on the one hand, and consumption and investment, on the other. We showed that the effect of a given movement in equity prices on consumption would be different depending on what fundamental shock had caused equity prices to move in the first place. In particular, a shock to the risk premium could have a different effect on consumption than a shock to the risk-free rate, even if the two have identical impacts on wealth. We showed some empirical evidence to support this contention, which also enabled us to explore dynamics. We found that, in our simple model, the effect of a given movement in equity prices on investment did not depend on the source of the shock. But we provided some arguments as to why this might not be true in the real world. In particular, we suggested that movements in equity prices unrelated to fundamentals may have a smaller effect on investment (if any) than those related to fundamentals. Our empirical work, together with the earlier work of MacGorain and Thompson (2002) and Bond *et al* (2004), suggests that this argument seems to hold true in UK data.

This paper has still left some questions unresolved. In particular, it would be good to model more explicitly the sorts of shocks that drive the economy and their stochastic processes. Then it would be interesting to investigate whether or not it makes sense in a more complicated model for shocks to equity price volatility to have no effect on consumption volatility. To do this, one would have to examine the links between the equity risk premium and the volatilities of fundamental shocks. In turn, that would mean considering models that are able to explain the magnitude of observed equity risk premia, something that is not the case in the simple model we presented. It would also be interesting to investigate whether there are any shocks (or specifications of the model) under which investment would respond more or less to equity price movements than in the baseline case.

References

Bertaut, C (2002), 'Equity prices, household wealth, and consumption growth in foreign industrial countries: wealth effects in the 1990s', *Federal Reserve Board International Finance Discussion Paper no. 724.*

Blake, A, Fernandez-Corugedo, E and Price, S (2003), 'The dynamics of consumers' expenditure: the UK consumption ECM redux', *Bank of England Working Paper no. 204*.

Blanchard, O J (1986), 'Comments and discussion', *Brooking Papers on Economic Activity*, Vol. 1.

Bond, S and Cummins, J (2001), 'The stock market and investment in the New Economy: some tangible facts and intangible fictions', *Brookings Papers on Economic Activity*, Vol. 1, pages 61-108.

Bond, S, Klemm, A, Newton-Smith, R, Syed, M and Vlieghe, G (2004), 'The role of expected profitability, Tobin's Q and cash flow in econometric models of company investment', *Bank of England Working Paper no. 222*.

Boone, L, Girouard, N and Wanner, I (2001), 'Financial market liberalisation, wealth and consumption', *OECD Working Paper no. 308.*

Caballero, R J (1999), 'Aggregate investment', in Taylor, J B and Woodford, M (eds), *Handbook of macroeconomics*, Vol. 1, pages 813-62.

Campbell, J Y and Shiller, R J (1988), 'The dividend-price ratio and expectations of future dividends and discount factors', *Review of Financial Studies*, Vol. 1, pages 195-228.

Chan, L, Karceski, J and Lakonishok, J (2001), 'The level and persistence of growth rates', *NBER Working Paper no. 8282.*

Chirinko, R S (1993), 'Business fixed investment spending: modeling strategies, empirical results and policy implications', *Journal of Economic Literature*, Vol. 31, Issue 4, pages 1,875-911.

Christiano, L J, Eichenbaum, M and Vigfusson, R (2003), 'What happens after a productivity shock?', Northwestern University, *mimeo*.

Davis, M A and Palumbo, M G (2001), 'A primer on the economics of time series econometrics of wealth effects', *Federal Reserve Board Finance and Discussion Series Working Paper no. 2001-09.*

Ellis, C and Price, S (2003), 'UK business investment: long-run elasticities and short-run dynamics', *Bank of England Working Paper no. 196.*

Fazzari, S M, Hubbard, R G and Petersen, B C (1988), 'Financing constraints and corporate investment', *Brookings Papers on Economic Activity*, Vol. 1, pages 141-95.

Friedman, M (1957), *A theory of the consumption function*, Princeton: Princeton University Press.

Hall, S (2001), 'Financial effects on corporate investment in UK business cycles', *Bank of England Quarterly Bulletin*, Winter, pages 449-59.

Jenkinson, N (1996), 'Saving, investment and real interest rates', *Bank of England Quarterly Bulletin*, February, pages 51-62.

Lantz, C D and Sarte, P G (2001), 'Consumption, savings and the meaning of the wealth effect in general equilibrium', *Federal Reserve Bank of Richmond Economic Quarterly*, Vol. 87/3, Summer, pages 53-71.

Lettau, M and Ludvigson, S C (2002), 'Time-varying risk premia and the cost of capital: an alternative implication of the Q theory of investment', *Journal of Monetary Economics*, Vol. 49, pages 31-66.

Lettau, M and Ludvigson, S C (2003), 'Understanding trend and cycle in asset values: reevaluating the wealth effect on consumption', *American Economic Review*, forthcoming.

Ludwig, A and Sløk, T (2001), 'Impact of stock prices and house prices on consumption in OECD countries', IMF, *mimeo*.

MacGorain, S and Thompson, J (2002), 'Profit expectations and investment', *Bank of England Quarterly Bulletin*, Winter, pages 404-09.

Obstfeld, M and Rogoff, K (1996), *Foundations of international macroeconomics*, Cambridge, Massachusetts: MIT Press.

Robertson, D and Wright, S (2002), 'What does *q* predict?', University of London, *mimeo*.

Tobin, J (1969), 'A general equilibrium approach to monetary theory', *Journal of Money, Credit and Banking*, Vol. 1, pages 15-29.