

Risk & Sustainable Management Group

Risk & Uncertainty Program Working Paper: R05#3

Research supported by an Australian Research Council Federation Fellowship
http://www.arc.gov.au/grant_programs/discovery_federation.htm

Cost Minimization and Asset Pricing

Robert G. Chambers

Dept of Agricultural and Resource Economics, University of Maryland,
College Park

and

John Quiggin

Australian Research Council Federation Fellow, University of Queensland

Schools of Economics and Political Science
University of Queensland
Brisbane, 4072
rsmg@uq.edu.au
<http://www.uq.edu.au/economics/rsmg>



THE UNIVERSITY
OF QUEENSLAND
AUSTRALIA

Cost Minimization and Asset Pricing

Abstract

A cost-based approach to asset-pricing equilibrium relationships is developed. A cost function induces a stochastic discount factor (pricing kernel) that is a function of random output, prices, and capital stock. By eliminating opportunities for arbitrage between financial markets and the production technology, firms minimize the current cost of future consumption. The first-order conditions for this cost minimization problem generate the stochastic discount factor. The cost-based approach is dual in nature and determines state-claim prices as the current-period marginal cost of increasing future stochastic output. A cost-based pricing kernel is estimated using annual time-series data on macroeconomic variables and returns data.

Asset pricing theory requires that an asset's price equals the inner product of a *stochastic discount factor* (or pricing kernel) and the asset's stochastic payout (Ross, 1978; Harrison and Kreps, 1979; Hansen and Singleton, 1982; Clark, 1993; Cochrane, 2001; Campbell, 2003; Duffie, 2003; and many, many others). The stochastic discount factor can be rationalized as subjective Arrow "state-claim prices". If the stochastic return on an asset is denoted by \tilde{R} and the stochastic discount factor is represented by \tilde{m} , the equilibrium implication is that

$$E \left[\tilde{m} \tilde{R} \right] = 1, \quad (1)$$

where E represents the expectation operator.

The consumption-based approach identifies \tilde{m} with the consumer's intertemporal marginal rate of substitution between nonstochastic current period consumption and stochastic future consumption. An important, and apparently unresolved, empirical challenge to the consumption-based approach is that the resulting models do not appear to fit market data particularly well (Hansen and Singleton, 1982; Hansen and Jagannathan, 1991, 1997; Campbell, Lo, and MacKinlay, 1997; Cochrane, 2001; Campbell, 2003). Perhaps the most famous manifestation of this lack of fit is the *equity-premium puzzle* introduced by Shiller (1982) and Mehra and Prescott (1985). As is well-known, for the most popular specification of the consumption-based discount factor, the consumption-based model can be reconciled with observed low volatility of aggregate consumption growth only if risk aversion is much larger than is commonly believed.

Our intent is not to resolve the equity-premium puzzle or to explain the perceived poor performance of the consumption-based model. Instead, our focus is on enlarging the range of economic models available for asset pricing analysis. We consider a representation of the stochastic discount factor that arises not from consumers optimally smoothing stochastic consumption across time but from the intertemporal optimization behavior of producing firms that have access to financial markets. The associated asset pricing rule emerges from the rational need to exploit any opportunities for risklessly raising intertemporal returns or lowering current period costs.

Even though many presentations of financial market equilibrium quite consciously ignore producers, there are a number of reasons to take their perspective in looking at asset-market

equilibrium. Most importantly, as business cycle theory suggests, the macroeconomic fluctuations that drive fluctuations in asset markets are most closely associated with production-side shocks. Financial markets react to real fluctuations. Ultimately any theory that explains asset-price behavior must be capable of portraying and measuring that linkage. The close causal nexus between production-side fluctuations and financial markets is underlined by the observation that many financial markets originally arose to manage risks associated with uncertain production. Empirically the production side of the economy also seems more variable than the consumption side. It is precisely the smoothness of the consumption side of the aggregate economy relative to the production and financial parts of the economy that makes the equity premium puzzle so compelling to theorists and empiricists.

Our analysis is most closely related to Cochrane's (1991) production-based and Cochrane's (1996) investment-based asset pricing models. Cochrane (1991, 1996) recognized that asset returns in a properly functioning market should be priced by accurate models of stochastic intertemporal marginal rates of transformation equally as well as by accurate models of stochastic intertemporal marginal rates of substitution. The differences between the consumption-based and production-based approaches reflect the differences that naturally arise from looking at an equilibrium relationship from two different sides of the market, that of the intertemporal consumer and that of the intertemporal producer.

Our approach differs from Cochrane (1991, 1996) and much of the closely related empirical literature on real-business cycles (e.g. Jermann (1998), Tallarini (2000)) in that we do not ground our analysis in a stochastic production function representation of the technology. Instead, we rely on a cost function that is dual to a more general primal representation of the technology than the stochastic production function—the state-contingent input correspondence. This approach allows us to identify the stochastic discount factor with the firm's current period marginal cost of future stochastic production. This has several advantages. Most importantly, unlike the stochastic production function approach, it does not arbitrarily impose a zero marginal rate of transformation between production in different states of Nature upon the production technology, and it allows us to estimate a stochastic discount factor without making any restrictive assumptions on the subspace spanned by financial markets.

The testable component of what we refer to as the *cost-based approach* is a moment

restriction on the joint stochastic process of asset returns, output, investment, existing capital stock, and measures of input and output prices of the general form

$$E \left[m(\tilde{z}, i, k, w, p) \tilde{R} \right] = 1,$$

where \tilde{z} denotes stochastic output, i denotes investment, k denotes the existing capital stock, w denotes input prices, and p denotes output price.

1 State-Contingent Technologies and the Asset Structure

To preserve mathematical simplicity, we develop the basic equilibrium pricing relationship in a two-period setting with a single stochastic output. Generalizing results to the case of multiple outputs and multiple time periods is straightforward.

Firms face a stochastic environment in a two-period setting. The current period, 0, is certain, but the future period, 1, is uncertain. Uncertainty is resolved by ‘Nature’ making a choice from a state space Ω . Each element of Ω is referred to as a state of nature. The random variable space is \mathfrak{R}^Ω , which we endow with the usual expectations inner product and norm (Luenberger, 1969).

The only assumption on firms’ preferences is that they are strictly increasing in period 0 consumption and nondecreasing in period 1 stochastic income.

The firm’s stochastic production technology is represented by a single-product, state-contingent continuous input correspondence that exhibits internal costs of adjustment associated with current period investment. Let $\mathbf{x} \in \mathfrak{R}_+^N$ be a vector of variable inputs (e.g. labor and nonlabor services) committed prior to the resolution of uncertainty (period 0), $i \in \mathfrak{R}$ be the level of current period investment in the capital good, k the existing (period 0) stock of capital, and $\tilde{z} \in \mathfrak{R}_+^\Omega$ the stochastic output chosen in period 0 but realized in period 1. The period 1 price of the state-contingent output is taken as nonstochastic and denoted by p .¹ The current period price of the investment good is normalized to one.

¹In the empirical analysis, nothing of substance changes by taking the period 1 price of the output to be stochastic.

The continuous input correspondence, $X : \mathfrak{R}_+^\Omega \times \mathfrak{R}^2 \rightarrow \mathfrak{R}_+^N$, maps stochastic output, capital and investment into variable input sets:

$$X(\tilde{z}, k, i) = \{\mathbf{x} \in \mathfrak{R}_+^N : \mathbf{x} \text{ can produce } \tilde{z} \text{ given investment level } i \text{ and capital stock } k\}.$$

Intuitively, $X(\tilde{z}, k, i)$ is associated with all of the variable-input combinations on or above the firm's production isoquant for \tilde{z} for a given level of current-period investment and capital stock. In addition to continuity of X , the only technical restriction that we require is that $X(\tilde{z}, k, i)$ satisfies free disposability of state-contingent output, or, more precisely $\tilde{z}' \leq \tilde{z} \Rightarrow X(\tilde{z}, k, i) \subseteq X(\tilde{z}', k, i)$.

Period 0 input prices are denoted by $\mathbf{w} \in \mathfrak{R}_+^N$ and are non-stochastic. The (period 0) production cost function, $c : \mathfrak{R}_{++}^N \times \mathfrak{R}_+^\Omega \times \mathfrak{R}^2 \rightarrow \mathfrak{R}_+$, is defined

$$c(\mathbf{w}, \tilde{z}, k, i) = \min_{\mathbf{x}} \{\mathbf{w}\mathbf{x} : \mathbf{x} \in X(\tilde{z}, k, i)\} \quad \mathbf{w} \in \mathfrak{R}_+^N$$

if $X(\tilde{z}, k, i) \neq \emptyset$ and ∞ otherwise. $c(\mathbf{w}, \tilde{z}, k, i)$ is continuous and nondecreasing in \tilde{z} over the region in which it is finite, and it is nondecreasing and superlinear in \mathbf{w} .

$c(\mathbf{w}, \tilde{z}, k, i)$ is also subject to internal costs of adjustment. The internal adjustment cost model usually holds that current period investment, by diverting resources away from productive activities, raises current period cost. It is also typically presumed that cost is convex in stochastic output and that higher levels of current period capital stock lower current period costs of output. We do not need any of these restrictions to develop our representation of the pricing kernel, and thus in the interest of generality, we do not impose them. However, as we show below, our empirical results support these properties for our data set.

Financial markets are frictionless, and the *ex ante* financial security payoffs are given by the $\Omega \times J$ matrix \mathbf{A} . The stochastic payout on the j th financial asset is denoted $\tilde{A}_j \in \mathfrak{R}^\Omega$, and its period 0 price is denoted v_j . The firm's portfolio vector, corresponding to the period 0 purchases of the financial assets, is denoted $\mathbf{h} \in \mathfrak{R}^J$. Denote the subspace spanned by \mathbf{A} as $M \subset \mathfrak{R}^\Omega$, where

$$M = \{\tilde{y} : \tilde{y} = \mathbf{A}\mathbf{h}, \mathbf{h} \in \mathfrak{R}^J\}.$$

Denote the j th vector of state-contingent returns by $\tilde{R}_j = \frac{\tilde{A}_j}{v_j}$, and the state-contingent returns matrix by $\mathbf{R} = [\tilde{R}_1, \dots, \tilde{R}_J]$.

2 An Equilibrium Relationship

Given any equilibrium level of current period investment and period 1 stochastic income, denoted by $\tilde{y} \in \mathfrak{R}^\Omega$, from production and financial investments, firms must solve:

$$C(\mathbf{w}, \mathbf{v}, p, \tilde{y}, i) = \min_{\mathbf{h}, \mathbf{z}} \{c(\mathbf{w}, \tilde{z}, k, i) + \mathbf{v}'\mathbf{h} : \mathbf{A}\mathbf{h} + p\tilde{z} \geq \tilde{y}\}. \quad (2)$$

If firms did not behave in this manner, there would exist unexploited arbitrage opportunities for them to raise period 0 consumption expenditures (by lowering costs) while maintaining stochastic period 1 income (and thus future stochastic consumption). This cannot be consistent with any reasonable notion of equilibrium.

In what follows, it eases exposition if we assume that $c(\mathbf{w}, \tilde{z}, k, i)$ is Gateaux differentiable in \tilde{z} .² Define the *directional derivative* of $c(\mathbf{w}, \tilde{z}, k, i)$ in the direction $\tilde{n} \in \mathfrak{R}^\Omega$ by

$$c'(\mathbf{w}, \tilde{z}, k, i; \tilde{n}) = \lim_{t \downarrow 0} \left\{ \frac{c(\mathbf{w}, \tilde{z} + t\tilde{n}, k, i) - c(\mathbf{w}, \tilde{z}, k, i)}{t} \right\}.$$

Notice that the directional derivative, $c'(\mathbf{w}, \tilde{z}, k, i; n)$, is positively linearly homogeneous in n . $c(\mathbf{w}, \tilde{z}, k, i)$ is said to admit a *Gateaux derivative*, denoted $\partial c(\mathbf{w}, \tilde{z}, k, i)$, if this limit exists for all $\tilde{n} \in \mathfrak{R}^\Omega$ and

$$c'(\mathbf{w}, \tilde{z}, k, i; \tilde{n}) = E[\partial c(\mathbf{w}, \tilde{z}, k, i)\tilde{n}],$$

for all \tilde{n} where $E[\tilde{x}\tilde{y}]$ denotes the expectations inner product. Intuitively, the Gateaux derivative is the marginal current period cost of changing stochastic period 1 production.

If $c(\mathbf{w}, \tilde{z}, k, i)$ is Gateaux differentiable, necessary conditions for an interior solution to (2) include (Clarke, 1983):

$$c'(\mathbf{w}, \tilde{z}, k, i; \frac{\tilde{A}_j}{p}) = E\left[\frac{\partial c(\mathbf{w}, \tilde{z}, k, i)}{p}\tilde{A}_j\right] = v_j \quad (3)$$

for all j , or in returns notation

$$c'(\mathbf{w}, \tilde{z}, k, i; \frac{\tilde{R}_j}{p}) = E\left[\frac{\partial c(\mathbf{w}, \tilde{z}, k, i)}{p}\tilde{R}_j\right] = 1, \quad j = 1, \dots, J. \quad (4)$$

²Gateaux differentiability assures uniqueness of a pricing kernel. It is easy to verify following the arguments in Chambers and Quiggin (2000) that $c(\mathbf{w}, \tilde{z}, k, i)$ need not be Gateaux differentiable for arbitrary technologies. In fact, as we illustrate below, it is never differentiable at efficient production points for the popular stochastic production function technology. When the cost-structure is not differentiable, the requirement of Gateaux differentiability can be relaxed to the notion of differentiability introduced by Clarke (1983) with little true change in the argument.

This equilibrium condition is easily explained. Suppose that starting from production position \tilde{z} , the firm replicates, via its stochastic production technology, $t\tilde{A}_j$ of the j th asset. The resulting marginal cost is

$$c(\mathbf{w}, \tilde{z} + t\frac{\tilde{A}_j}{p}, k, i) - c(\mathbf{w}, \tilde{z}, k, i).$$

In the limit, this marginal cost equals $c'(\mathbf{w}, \tilde{z}, k, i; \frac{\tilde{A}_j}{p})$. Now suppose further that $v_j > c'(\mathbf{w}, \tilde{z}, k, i; \frac{\tilde{A}_j}{p})$. The firm then could profitably sell off some of its holding of the j th asset while replacing it with a like amount produced using its stochastic technology. (If it has no holding, it could profitably sell what it has produced.) Either way, current period cost is lowered, and the firm can achieve firm an unambiguously higher current period consumption. There is thus an arbitrage available between the technology and financial markets that cannot be consistent with equilibrium behavior. On the other hand, if $v_j < c'(\mathbf{w}, \tilde{z}, k, i; \frac{\tilde{A}_j}{p})$, the firm could alter its period 1 stochastic production by $\frac{\tilde{A}_j}{p}$ and profitably replace it by purchasing the j th asset in the current period.³ Thus, for smooth technologies, unless $v_j = c'(\mathbf{w}, \tilde{z}, k, i; \frac{\tilde{A}_j}{p})$ holds for all j , there exist arbitrage opportunities between the physical technology and financial markets that the firm can risklessly exploit to raise its current period consumption. This cannot be consistent with rational behavior by the firm.

Thus,

$$\tilde{m} = \frac{\partial c(\mathbf{w}, \tilde{z}, k, i)}{p}, \tag{5}$$

represents an appropriate stochastic discount factor for the firm of the same basic form as (1). Before discussing this version of the stochastic discount factor further, it is important to make several points about its derivation. This asset-price relationship is an entirely production-based asset pricing model. Apart from monotonicity, it does not rely on any restrictions on the firm's attitudes towards risk. Therefore, it is valid for any firm with monotonic preferences. The development does not require any restrictions on the subspace spanned by the market assets. In particular, it does not require any of the following: complete markets,

³Here we rely on the fact that Gateaux differentiability ensures that

$$c' \left(\mathbf{w}, \tilde{z}, k, i; \frac{A_j}{p} \right) = -c' \left(\mathbf{w}, \tilde{z}, k, i; -\frac{A_j}{p} \right).$$

More generally this need not be true.

investment returns lie within the market span, M , or that the mean-variance frontier for M be contained within the span of factors defined by production or investment opportunities. The asset pricing relationship can be inferred solely from the firm's first-order conditions for the removal of arbitrages between financial and production opportunities. And because it arises from the firm's optimization behavior, it yields clear theoretical predictions.

3 Interpretation of the Stochastic Discount Factor

One way to examine the properties of the stochastic discount factor (5) is to compare it to a stochastic discount factor derived from the consumption-based model. Figure 1 illustrates the market span, M , as a ray through the origin of state-contingent income space, \mathfrak{R}^Ω . The pricing kernels, \tilde{m} , that satisfy

$$E \left[\tilde{m} \tilde{A} \right] = v$$

for M are given by the hyperplane perpendicular to M labelled m in that figure. The no-arbitrage prices are the strictly positive elements of the hyperplane m . The consumption-based approach to asset pricing isolates a particular element of m , the consumption-based stochastic discount factor, by finding a point of tangency (not drawn) between a hyperplane parallel to m and the representative-consumer's indifference curve in state-contingent income space.

The approach that we are advocating looks at the other side of the market. It isolates a particular element of m by finding a point of tangency between it and the representative firm's isocost curve in state-contingent income space (as normalized by p). This is illustrated by point A in the figure. For clarity's sake, we have drawn this isocost curve as though $c(\mathbf{w}, \tilde{z}, i, k)$ were quasi-convex in stochastic output. The assumption of quasi-convexity, however, is not essential to the derivation of the equilibrium pricing relationship just as the assumption of risk aversion is not essential to the derivation of the equilibrium pricing relationship in the consumption-based approach.

Intuitively, therefore, our derivation of the pricing kernel says little more than that the firm should equate its marginal rate of substitution between state-contingent incomes to its marginal rate of transformation of these incomes as a producer. *Formally, however,*

the validity of (5) as a stochastic discount factor only requires that the firm's preferences are decreasing in current period cost and increasing in future income. Its existence and properties do not hinge on any assumption about the firm's risk attitudes. This latter point is essential because (5) suggests a fundamentally different way of looking at pricing relationships than the consumption-based approach.

As usual, by rewriting the equilibrium pricing relationship in terms of covariances and means, we obtain for any payoff that

$$v_j = E \left[\tilde{A}_j \right] E [\tilde{m}] + Cov \left(\tilde{m}, \tilde{A}_j \right).$$

The second term, $Cov \left(\tilde{m}, \tilde{A}_j \right)$, is usually thought of as a risk adjustment. Intuitively, an asset that covaries positively with the stochastic discount factor has its price raised by the risk adjustment while an asset that covaries negatively with the stochastic discount factor has its price lowered by the risk adjustment.

The usual intuition for the risk adjustment comes from the consumption-based model where \tilde{m} is defined in terms of the marginal utility of stochastic period 1 consumption. Assuming expected-utility preferences and risk aversion, this marginal utility covaries negatively with period 1 consumption. Thus, the risk adjustment lowers an asset's price if that asset varies positively with period 1 consumption and raises its price if the asset varies negatively with period 1 consumption. Risk-averse individuals are willing to pay a premium for an assets that balance their consumption risk.

Different forces are at work for (5). The stochastic discount factor now depends upon stochastic production levels, input prices, current period investment, the capital stock, and implicitly the current state of technical knowledge. Assets that covary positively with the stochastic discount factor still have their price raised by the risk adjustment. But now the intuition and the theoretical predictions are different. Suppose that marginal cost, as one typically expects, is increasing in stochastic output. Then assets that covary positively with stochastic output have their prices raised by the risk adjustment, and assets that covary negatively with stochastic output have their prices lowered. The explanation is that (5), instead of evaluating the marginal utility of consumption, measures the marginal cost of replicating assets via the production technology. Thus, assets that covary positively with

this marginal replication cost should have their prices raised precisely because they are more costly to replicate physically when output is high than assets that covary negatively with the marginal replication cost. The risk adjustment now manifests a *replication effect* rather than a risk-averse consumer's response to consumption risk.

We have argued in passing that the specification of the technology in terms of general input correspondences and cost structures has the threefold advantage of generality, tractability and empirical flexibility over the more familiar stochastic production function approach used in both the cost-of-adjustment investment and production-based asset pricing literatures. This seems to be a good point at which to explain why.

Suppose that the production technology is modeled by a stochastic production function, subject to adjustment costs, of the form

$$\tilde{z} = F(k, i, \tilde{\varepsilon}),$$

where $\tilde{\varepsilon} \in \mathfrak{R}^\Omega$ may be interpreted as a random production shock or a random input. Here the interpretation is that i and k are determined in period 0, and then Nature intervenes by choosing a realization of $\tilde{\varepsilon} \in \mathfrak{R}^\Omega$ to determine *ex post* output in period 1. One approach to generating a stochastic discount factor is to recognize that marginal changes in current period investment define a random variable of the form

$$\tilde{\delta} = F_i(k, i, \tilde{\varepsilon}).$$

Then following Cochrane (1991) so long as $\tilde{\delta} \in M$, a mimicking portfolio for $\tilde{\delta}$ can be constructed, and its price forms the basis for a production-based model of asset pricing.

Unfortunately, as Chambers and Quiggin (1998, 2000) have shown, this specification of the technology suffers from an obvious technical shortcoming that appears to have been largely overlooked in both the theoretical and empirical literatures. This shortcoming is most clearly illustrated by assuming for the moment that $\Omega = \{1, 2\}$ and $\tilde{\varepsilon} = (\varepsilon_1, \varepsilon_2)$. It then follows immediately that

$$\tilde{z} = (F(k, i, \varepsilon_1), F(k, i, \varepsilon_2)),$$

and

$$\tilde{\delta} = (F_i(k, i, \varepsilon_1), F_i(k, i, \varepsilon_2)).$$

Notice, in particular, that the first of these equalities says that for any given level of investment and the capital good, only one possible pair of state-contingent outputs can emerge. If one were to illustrate this production technology in terms of a state-contingent product transformation curve (with free disposability of outputs), one would obtain a right-angled product transformation curve that is the mirror image of a Leontief indifference set (Chambers and Quiggin (1998, 2000)). *There is, by assumption, no substitutability between state-contingent outputs.*⁴ Regardless of the dimension of Ω , this remains true. Thus, the stochastic production function cannot lead to a cost function representation that is Gateaux differentiable at technically efficient points (more on this below).⁵

The second equality leads to a similar conclusion in terms of the production perturbation. For a given level of investment and capital, firms can only arrange their investment activities to incur a single pattern of marginal changes in future returns. When it is realized that the most common specification for the stochastic production function in the literature is of the multiplicative form:

$$F(k, i, \tilde{\varepsilon}) = \tilde{\varepsilon} f(k, i),$$

then the starkness of this assumption becomes even more apparent. Here the firm faces production shocks to which it cannot react in making its production choices. All it can do is choose the magnitude of the production risk that it faces, much in the same fashion that an individual producer, when faced with a single asset chooses the magnitude of the risk he or she faces by choosing his holding of the asset.⁶

In particular, if it is assumed that $\tilde{\delta} \in M$, it follows that $\tilde{\varepsilon} \in M$. Thus, the physical production technology is effectively redundant in the presence of financial markets. In other words, the uncertainty of physical production plays no truly independent role in determining the ultimate level of uncertainty that the economy faces. Or put another way, all production risk can be modelled as though it arises in financial markets and not from real phenomena.

⁴Cochrane (1996, p.574) recognizes that the production function specification that he chooses ensures that "...there is nothing a producer can do to transform goods *across states*". (Italics in original.)

⁵Similarly, it does not lead to a distance or transformation function representation that is Gateaux differentiable at efficient points (Chambers and Quiggin, 2000).

⁶It is on this basis that it is routinely argued that the problem of the firm facing a stochastic technology is isomorphic to the simple portfolio selection problem (Gollier, 2001).

It is as though all of the uncertainty that consumers and producers face arise not from the real side of the economy, but from the financial sector.

The fact that the stochastic production function representation leads to a situation where the state-contingent transformation curve for the production technology exhibits zero substitutability between state-contingent outputs also has important implications for firm behavior. As pointed out above, in geometric terms, this is manifested by a right-angled kinked product transformation curve, whose supporting hyperplanes span \mathfrak{R}_+^Ω . Thus, any explanatory power that can emerge from a production-based asset pricing model or the closely related q -theory of investment under this specification must emerge solely from the elimination of intertemporal arbitrage opportunities as opposed to removing (period 1) intratemporal arbitrage opportunities made available to the firm by the simultaneous existence of financial markets and stochastic physical technologies.

The consequences of this difficulty may perhaps be best grasped by considering its mirror reflection in consumption-based asset pricing models. There the stochastic discount factor is given by the stochastic marginal rate of substitution between consumption in different states of Nature divided by the consumer's subjective discount rate. The assumption of a zero state-contingent marginal rate of transformation would be mirrored in a zero marginal rate of substitution between period 1 consumption in different states of Nature. *The parallel assumption is that investors are perfectly risk averse.* Asset pricing would have to be explained entirely in terms of discounting sure returns back to the current period because all investors would rationally strive for portfolios that yielded a sure return.

4 An Empirical Model

To illustrate (5), we attempt to estimate it using annual U.S. macroeconomic data on aggregate production (Gross Domestic Product) and its price, aggregate investment (Gross Private Domestic Investment), unit labor cost, unit nonlabor cost, stock price returns (returns on the Standard & Poor's 500), and returns on commercial paper for the period 1929-1995.⁷

⁷The Standard and Poor's returns data and the return on commercial paper were drawn from <http://kuznets.fas.harvard.edu/~campbell/data.htm>. They correspond to the data that underlie some of

We assume that aggregate production, as measured by Gross Domestic Product (which we take to be stochastic), can be modelled as though there exists a representative producer who rationally removes any arbitrage opportunities between the physical technology and financial markets.

To implement the theoretical model empirically, we must first specify an econometrically estimable form for the pricing kernel. This requires specification of a cost function. Specifying a cost function for a stochastic production technology presents a number of difficulties not present in specifying estimable versions of nonstochastic technologies. The cost function, $c(\mathbf{w}, \tilde{z}, i, k)$, and its dual input set, $X(\tilde{z}, i, k)$, both depend on the random variable $\tilde{z} \in \mathbb{R}^\Omega$. As with all random variables, \tilde{z} is only incompletely observed because one typically only has observations on one *ex post* realization of any random variable for any observation point.

The usual tactic pursued in such situations is to make an identification assumption that permits estimation of a nonstochastic portion of the technology, and then use that estimated knowledge to construct an approximation of the underlying distribution. The familiar stochastic production function with variable inputs and a multiplicative error structure illustrates (Cochrane, 1991; Jermann, 1998; Tallarini, 2000). There it is typically assumed that stochastic output is related to inputs by a relationship of the form

$$\tilde{z} = f(\mathbf{x}, i, k) \tilde{\varepsilon},$$

where $\tilde{\varepsilon}$ is now a positive random variable with $E[\tilde{\varepsilon}] = 1$, and $f(\mathbf{x}, i, k)$ is a suitably nonstochastic parametric representation of expected output. Once a stochastic structure for $\tilde{\varepsilon}$ is specified, it can be treated as an error term in the estimation of $f(\mathbf{x}, i, k)$. This permits the estimation of the stochastic technology using only observations on the *ex post* output realization.

But, as discussed above, it brings a an economic cost in loss of generality because it assumes that the underlying technology admits zero substitutability between outputs for different realizations of $\tilde{\varepsilon}$. It is precisely this assumption (zero substitutability) that allows the empirical analysis in Campbell (2003). The data on the macroeconomic variables (real gross domestic product (gdp), unit labor cost, unit nonlabor cost, and gross private domestic investment) are from the US National Income Product Account website.

one to infer the entire output distribution from a single *ex post* observed output given knowledge of $f(\mathbf{x}, i, k)$. The case where Ω is finite and given by

$$\Omega = \{1, 2, \dots, S\},$$

with

$$\tilde{\varepsilon} = \{\varepsilon_1, \varepsilon_2, \dots, \varepsilon_S\},$$

and

$$\begin{aligned} \tilde{z} &= \{f(\mathbf{x}, i, k) \varepsilon_1, f(\mathbf{x}, i, k) \varepsilon_2, \dots, f(\mathbf{x}, i, k) \varepsilon_S\} \\ &= \{z_1, z_2, \dots, z_S\} \end{aligned}$$

illustrates. Once a single z_i is known, then the remaining elements of \tilde{z} can be determined via

$$z_s = \frac{z_i}{\varepsilon_i} \varepsilon_s.$$

But as Chambers and Quiggin (1998, 2000) show:

$$c(\mathbf{w}, \tilde{z}, i, k) = \max_{1, 2, \dots, S} \left\{ \min_{\mathbf{x}} \left\{ \mathbf{w}' \mathbf{x} : f(\mathbf{x}, i, k) \geq \frac{z_s}{\varepsilon_s} \right\} \right\}.$$

for this technology. This cost structure is nondifferentiable at economically efficient points.

The approach taken in this paper is to assume at time t a quasi-homothetic cost function of the form:

$$c(\mathbf{w}_t, \tilde{z}_{t+1}, k_t, i_t, t) = \beta(\mathbf{w}_t, i_t, k_t) + \gamma(\mathbf{w}_t) \left[a + c \frac{i_t}{k_t} \right] E[\tilde{z}_{t+1}] + \gamma(\mathbf{w}_t) \frac{b}{2} E[\tilde{z}_{t+1}^2],$$

where $\beta(\mathbf{w}_t, i_t, k_t)$ is nondecreasing and superlinear in \mathbf{w}_t , and

$$\gamma(\mathbf{w}_t) = 100 \left(w_{lt}^{\frac{1}{2}} w_{nt}^{\frac{1}{2}} \right),$$

with w_{lt} denoting unit labor cost and w_{nt} unit nonlabor cost.⁸

Empirically, k_t is constructed recursively using data on i_t by assuming an annual depreciation of .04 as

$$k_{t+1} - k_t = i_t - .04k_t,$$

⁸By convention, w_{lt} and w_{nt} are taken to be the unit costs reported for the same year as z_{t+1} and p_{t+1} .

with initial period (1929) capital stock arbitrarily normalized to one.

This specification yields a linear stochastic discount factor at time t for period $t + 1$ that is given by:

$$\tilde{m}_t = \frac{\partial c(\mathbf{w}_t, \tilde{z}_{t+1}, k_t, i_t, t)}{p_{t+1}} = \frac{\gamma(\mathbf{w}_t)}{p_{t+1}} \left[a + c \frac{i_t}{k_t} + b \tilde{z}_{t+1} \right].$$

Taking expectations conditional on the information available at time t yields for random asset return, \tilde{R}_{t+1} ,

$$\gamma(\mathbf{w}_t) E_t \left[a \frac{\tilde{R}_{t+1}}{p_{t+1}} + c \frac{i_t}{k_t} \frac{\tilde{R}_{t+1}}{p_{t+1}} + b \tilde{z}_{t+1} \frac{\tilde{R}_{t+1}}{p_{t+1}} \right] = 1. \quad (6)$$

The law of iterated expectations then implies the following unconditional expectation

$$h_t \equiv E \left[\gamma(\mathbf{w}_t) \left[a \frac{\tilde{R}_{t+1}}{p_{t+1}} + c \frac{i_t}{k_t} \frac{\tilde{R}_{t+1}}{p_{t+1}} + b \tilde{z}_{t+1} \frac{\tilde{R}_{t+1}}{p_{t+1}} \right] - 1 \right] = 0, \quad (7)$$

for any stochastic return.

Our estimation procedure (see below) is based upon the generalized method of moments (GMM). Thus, (7) for a single asset does not contain enough sample information to permit identification of all three parameters of the pricing kernel. To permit identification we pursue the strategy of introducing instrumental variables into (6) using variables that can be plausibly taken as known at time t (and thus statistically predetermined), which we denote by v_t , to generate additional unconditional expectations of the form

$$g_t \equiv E \left[v_t \gamma(\mathbf{w}_t) \left[a \frac{\tilde{R}_{t+1}}{p_{t+1}} + c \frac{i_t}{k_t} \frac{\tilde{R}_{t+1}}{p_{t+1}} + b \tilde{z}_{t+1} \frac{\tilde{R}_{t+1}}{p_{t+1}} \right] - v_t \right] = 0. \quad (8)$$

In estimation, we used three sets of instruments. The first set of instruments corresponds to (w_{lt}, w_{nt}) and yields an exactly identified system in the single return estimation results presented below. The second set of instruments corresponds to $(w_{lt}, w_{nt}, \mathbf{R}_t)$ where \mathbf{R}_t is the two-vector containing the observed return on the S&P 500 and commercial paper, and the third set corresponds to $(w_{lt}, w_{nt}, \mathbf{R}_t, \mathbf{R}_{t-1})$.

A few further comments are in order. Most importantly, this cost structure (like the expected-utility functional) is additively separable across states of Nature. This facilitates estimation because it allows us to replace the random variable \tilde{z}_t by its *ex post* realization, z_t , in the construction of the sample analogues of (7) and (8) that form the basis of the GMM estimation procedure.

But it also places restrictions on the underlying technology. It implies, for example, that the marginal cost associated with a small change in the s th realization of the random variable \tilde{z} , $z(s)$, $s \in \Omega$ is

$$\gamma(\mathbf{w}_t) \left[a + c \frac{i_t}{k_t} + bz_{t+1}(s) \right].$$

This marginal cost is independent of any other potential realization $z(k)$, $k \neq s$. Literally, this implies that the marginal cost of preparing output for one state of Nature is independent of the output levels chosen for the other states of Nature.⁹ Thus, where the stochastic production function specification assumes that firms cannot adjust to production risks in different states at all, this specification assumes that firms have an almost perfect ability (at some cost) to adjust to these risks.

Second, the marginal rate of transformation between output realizations $z(s)$ and $z(k)$ is

$$-\frac{\gamma(\mathbf{w}_t) \left[a + c \frac{i_t}{k_t} + bz_{t+1}(s) \right]}{\gamma(\mathbf{w}_t) \left[a + c \frac{i_t}{k_t} + bz_{t+1}(k) \right]}.$$

This marginal rate of transformation is not parametrically set to zero for efficient outcomes, as it would be for a stochastic production function specification. And because it is dependent upon the levels at which the different output realizations are chosen, it is not parametrically set to a constant. But, it is symmetric across states of Nature.

Finally, although we notationally model the cost function as dependent upon t , $c(\mathbf{w}_t, \tilde{z}_{t+1}, k_t, i_t, t)$, the actual specification is not directly dependent upon the time period, although it of course depends upon it indirectly through its dependence on investment and the capital stock. This implies, for example, that if the producer chose the same stochastic output, the same investment and faced the same capital stock and input prices, he or she would incur the same cost in 1995 as in 1929. This is implausible.

There are a number of different ways to incorporate the phenomenon of technical change into cost functions suitable for estimation using time-series data. The most common, of course, is to specify the technology as depending directly upon a time trend term. While

⁹Chambers and Quiggin (2000) show that this specification corresponds to a variable production technology that has what they refer to as state-allocable inputs. In short, this requires that inputs, such as labor and materials, can be allocated to state-specific tasks, which do not overlap across states of Nature.

this may be a tractable alternative for relatively short time periods, for longer series, it is implausible to presume that such a time trend term is stationary.

Here we tackle the problem through our treatment of the output variable. In particular, instead of measuring the random variable \tilde{z}_{t+1} in terms of levels, we measure it in terms of year-to-year changes. Thus, the actual assumption is that the cost structure at time is conditioned upon the *ex post* realization of output at time t , z_t , as follows:

$$c(\mathbf{w}_t, \tilde{z}_{t+1}, k_t, i_t, t) = \beta(\mathbf{w}_t, i_t, k_t) + \gamma(\mathbf{w}_t) \left[a + c \frac{i_t}{k_t} \right] E[\tilde{z}_{t+1} - z_t] + \gamma(\mathbf{w}_t) \left[\frac{b}{2} E[\tilde{z}_{t+1} - z_t]^2 \right].$$

Such a cost structure can be rationalized by assuming a particularly simple (and tractable) form of learning by doing over time.

5 Results and Discussion

In Table 1, we report estimation results for the parameters of the pricing kernel. There are three separate set of estimates for the parameters: those obtained by estimating the pricing kernel with different sets of instruments using only data on the S&P 500 returns; those obtained by estimating the pricing kernel with different sets of instruments using only data on commercial paper returns; and results obtained by estimating the pricing kernel using data on commercial paper and S&P 500 returns jointly.

These estimates were obtained using iterated GMM with an optimal weighting structure for the respective moment conditions. Specifically, the parameters of the pricing kernel were estimated by minimizing a weighted combination of the sample moments analogous to (7) and (8). Letting the sample moments of (7) and (8) be denoted by, respectively, g_T and h_T , the weighted combination is given by

$$J_T = [g_T, h_T] \Sigma^{-1} [g_T, h_T]',$$

where Σ is the spectral density matrix for the implied pricing errors. In estimation, Σ was estimated using the Newey-West procedure with lag length set to 4.¹⁰

¹⁰All estimation was done in a Matlab framework using the GMM program library developed and maintained by M. T. Cliff of Purdue University.

In viewing these results, several observations are apparent. First, although there are naturally differences depending on the number of instruments included (and thus included moment conditions), the estimated parameters are quite similar across all estimated versions of the model. And in most instances, the estimated parameters seem to be highly significant suggesting that deflated nonlabor and labor unit cost, output, and the investment-capital ratio all can play an important role in pricing these assets.. The estimates for a run from about .65 to approximately 1.10 and can be judged significantly different from zero at traditional confidence levels in all but two instances. The estimates for b range from .18 to .37 and are statistically different from zero in all but one instance at traditional confidence levels, and in that one instance it would be judged significantly different from zero at roughly the .07 level. The estimates for c range from 4.07 to 7.17 and can also be judged significantly different from zero in all but two instances.

Besides giving us information on the role that $\left(\frac{\mathbf{w}}{p}, z, \frac{i}{k}\right)$ play in pricing these assets, the estimated parameters provide information on the underlying cost structure.¹¹ The parameters a and b , respectively, measure the effect that changes in the first and second moments of \tilde{z} have on cost, while c measures the effect that increasing current period investment (as well as current period capital holdings) have on current period costs. The estimates of c are all positive which implies that increasing current period investment raises current period costs. This is the usual maintained hypothesis in the internal cost of adjustment literature. It can be explained intuitively by noting that deploying investment draws away variable inputs that could otherwise be used to produce period 1 output. So long as those diverted inputs exhibit free disposability (have positive marginal products), this diversion will raise the variable cost of producing output. In essence, purchasing and installing capital diverts one's attentions away from other productive activities.

The estimates of a and b suggest that independent increases in either of the first two moments of \tilde{z} raises the costs of producing aggregate output. The finding that increasing the first moment of output increases costs is not at all surprising. This just reflects the usual economic notion of positive marginal cost. The finding that the increases in the second moment tend to raise cost, however, may be less intuitively obvious. It implies, for example,

¹¹To obtain a complete picture, one would have to obtain estimates of $\beta(\mathbf{w}, i, k)$.

that increasing the variance of \tilde{z} increases cost, while lowering variance decreases cost. At an intuitive level, one might believe that reducing the variation of a stochastic output is costly because firms would be required to devote scarce resources to measures that prevent Nature's actions from having an adverse impact on production. Our empirical result suggests just the opposite.

There are a number of explanations. One is that aggregate production may not be inherently risky in the sense of Chambers and Quiggin (2000). If true, then firms tolerate riskiness in their production portfolio as a way of self insuring against the stochastic demand variations that they face in their product markets. Then, the positive (and highly significant) sign for b reflects the cost of dealing with stochastic variation in product markets. More generally, the sign of b will be positive if, in an appropriate sense, demand uncertainty is a more important source of variation than the inherent riskiness of production.

Second, and perhaps more important empirically, b also measure the presence or absence of economies of size for a given capital structure. When interpreted as a "size effect", the estimates of b imply a convex cost structure so that all state-contingent marginal costs are increasing in output. This is exactly what one expects from most production technologies. In fact, it is routinely imposed in most stochastic production function specifications. Standard intuition from the theory of the firm would suggest that size effects can lead to a 'replication effect' that leads to a positive adjustment in the price of assets that covary positively with stochastic production. Firms find it more costly to replicate those assets precisely when production is higher.

Ultimately, it would be desirable to disentangle the "size" and "variation" effects empirically. However, the convolution of "size" and "variation" effects is a cost imposed by a cost function that is additively separable across states of Nature. That specification was chosen for several reasons. Most importantly, it leads to a pricing kernel that can be approximated in moment terms by interactions between *ex post* observations on the random variables, \tilde{z} and \tilde{R} . But using sample moments to approximate the true moments without direct observations on unrealized components of \tilde{z} and \tilde{R} potentially confounds "size" and "variation" effects. The problem is similar to that of identifying technical change and size effects using only time-series data for nonstochastic technologies. To sort the "size" effect

from the "variation" effect, either a richer data set is needed, or even further structure must be imposed. We leave both problems to future research.

The theory requires that (7) holds exactly. Theory also suggests that (8) holds exactly for each instrument so long as the instrument can be viewed as known at the time of making the production and investment decisions. In the exactly identified case, the GMM estimation procedure, by definition, ensures that parameter estimates are chosen so that (7) holds exactly for the corresponding sample moment. In the cases where the parameters of the pricing kernel are overidentified, a straightforward test of (7) is offered by comparing the mean sample forecast error of the estimated pricing kernel to its estimated standard deviation. A test of the joint exactness of (7) and the relevant versions of (8) is offered by a test of overidentifying restrictions based on computed values of J_T . As is well-known, TJ_T is distributed as χ^2 with degrees of freedom equal to the number of moments less the number of estimated parameters (Hansen, 1982; Hansen and Singleton, 1982; Hamilton, 1994).

Computed values of the sample mean forecast error, its standard error, the relevant J -statistic, its probability value, and the relevant degrees of freedom are reported in Table 2.

In all the single-asset estimation results, the sample mean forecast error is not significantly different from zero at the .05 level. Thus, there is significant statistical evidence in support of the exactness of (7). The calculated values of the J statistics that were obtained in the single-asset estimation results also provide evidence in favor of the exactness of both (7) and (8).

Turning to the results from estimating the pricing kernel parameters jointly using data on S&P 500 returns and returns on commercial paper, the mean sample forecast error for S&P returns is positive and significantly different from zero in both instances. On the other hand, the mean sample forecast error for commercial paper is not significantly different from zero in either instance. Thus, the pricing kernel jointly estimated using both returns series tends to overpredict the return on the S&P 500 portfolio, while it seems to do well in pricing the return on commercial paper. Given the tendency to overpredict the S&P 500 returns, it is not surprising that the evidence in favor of the joint exactness of both version of (7) and the associated instrument-moment conditions is much weaker when the system is estimated

jointly.

Although the empirical pricing kernel has been deduced from a cost function specification, as executed, it corresponds to a stochastic discount factor that is linear in a set of macroeconomic variables. This specification highlights the potential empirical similarity of the approach that we are advocating to existing asset pricing models based on macroeconomic factors or 'state' variables (e.g. Chen, Roll, and Ross, 1986; Cochrane, 1996). Letting

$$\beta \mathbf{X}_t = \frac{\gamma(\mathbf{w}_t)}{p_{t+1}} \left[a + c \frac{i_t}{k_t} + b \tilde{z}_{t+1} \right],$$

our empirical specification implies for any return that

$$E_t \left[\tilde{R}_t \right] = \frac{1}{\beta E_t [\mathbf{X}_t]} - Cov \left(\beta \mathbf{X}_t, \tilde{R}_t \right).$$

Thus, our model can be converted to a form that is identical to 'factor pricing' models that take as factors or 'state' variables the macroeconomic variables that are in \mathbf{X}_t (Ross, 1978). These variables measure innovations in aggregate production, inflation and wage levels, and investment.

This observation merits some further discussion to distinguish the two approaches. Our "state variables" or factors are derived from and motivated by economic theory. Moreover, although applied at an aggregate level, the model being proposed here is a firm-level theory. It does not require any formal assumptions on the space of assets. It makes clear theoretical predictions at the micro level, and it is amenable to testing at that level. Thus, it is not inherently a macroeconomic model of asset pricing. In particular, the theory applies regardless of whether our "factors" can span asset space either exactly or approximately. Instead, our admittedly one-sided explanation of asset price behavior is grounded in the marginal cost of replicating financial assets by physical technologies and in the elimination of arbitrage opportunities between the firm's production opportunities and financial markets. It makes clear predictions, grounded in theory, about the relationship between production variables and asset prices.

It is also possible to interpret our empirical model as an approximate or reduced-form consumption-based model where general-equilibrium considerations and the production side of the economy have been used to determine proxies for or to "solve out" aggregate consumption. Thus, one might view our model as arising from a consumption-based pricing

model where aggregate consumption is specified to be of the form $C(\tilde{z}, i, k, \mathbf{w})$, and the resulting pricing kernel is then properly linearized. Notice, however, that such a reduced-form approach can lead to fundamentally different predictions than the ones that arise from our model. Suppose, in fact, that one takes a consumption-based approach, but that one models consumption as $C(\tilde{z}, i, k, \mathbf{w})$. Presumably, consumption is positively related to income in the form of GDP or GNP. Consumption and production would then presumably covary positively. The implication for asset pricing would be that if consumers are risk averse, then the price of assets that covary positively with production (and thus consumption) should receive a negative risk adjustment. This is the opposite of what the replication effect for convex technologies predicts in our model. The fact that our predictions do not coincide with the predictions that emerge from a reduced-form approach neither contradicts or invalidates the reduced-form approach. Rather, it reflects the fact that (at least) two sets of market forces are at play in determining equilibrium asset prices. It also highlights the role that economic theory can play in disentangling these competing forces in both empirical and theoretical work.

6 Conclusion

A cost-based approach to asset-pricing equilibrium relationships is developed. It is shown that a cost function subject to internal costs of adjustment induces a stochastic discount factor (pricing kernel) that is a function of random output, input and output prices, existing capital stock, and investment. The only assumption on firm preferences is that they are increasing in current period consumption and future stochastic consumption. This suffices to ensure that the firm will always strive to remove any opportunities for arbitrage between existing financial markets and its production technology. This ensures that the firm will always act to minimize current period cost of providing future consumption, and it is the first-order conditions for this cost minimization problem that generate the stochastic discount factor. Neither the theory or the empirical application requires any further restrictions on firm preferences or on the asset space. Where existing production-based asset-pricing models determine state-claim prices by modelling the stochastic intertemporal marginal rate

of transformation, the cost-based approach is dual in nature and determines state-claim prices as the current-period marginal cost of increasing future stochastic output. As an illustration, a cost-based pricing kernel is estimated using annual time-series data on macroeconomic variables and returns data for the S&P 500 and commercial paper.

7 References

Campbell, J. Y. "Consumption-Based Asset Pricing." Handbook of the Economics of Finance. G. M. Constantinides, M. Harris, and R. M. Stulz Amsterdam: Elsevier, 2003.

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

Chambers, R. G., and J. Quiggin. "Cost Functions and Duality for Stochastic Technologies." American Journal of Agricultural Economics 80 (1998): 288-95.

———. Uncertainty, Production, Choice, and Agency: The State-Contingent Approach. New York: Cambridge University Press, 2000.

Chen, N., R. Roll, and S. A. Ross. "Economic Forces and the Stock Market." Journal of Business 59 (1986): 383-403.

Clark, S. A. "The Valuation Problem in Arbitrage Price Theory." Journal of Mathematical Economics, no. 463-478 (1993).

Clarke, F. H. Optimization and Nonsmooth Analysis. New York: J. Wiley and Sons, 1983.

Cochrane, J. H. Asset Pricing. Princeton: Princeton University Press, 2001.

———. "A Cross-Sectional Test of an Investment-Based Asset Pricing Model." Journal of Political Economy 104 (1996): 572-621.

———. "Production-Based Asset Pricing and the Link Between Stock Returns and Economic Fluctuations." Journal of Finance 46 (1991): 209-37.

Duffie, D. "Intertemporal Asset-Pricing Theory." Handbook of the Economics of Finance. G. M. Constantinides, M. Harris, and R. M. Stulz Amsterdam: Elsevier, 2003.

Gollier, C. The Economics of Risk and Time. Cambridge: MIT Press, 2001.

Hamilton, J. Time-Series Analysis. Princeton: Princeton University Press, 1994.

Hansen, L. P. "Large Sample Properties of Generalized Method of Moments Estimators." Econometrica 50 (1982): 1029-54.

Hansen, L. P., and R. Jagannathan. "Assessing Specification Errors in Stochastic Discount Factor Models." Journal of Finance 52 (1997): 557-90.

———. "Implications of Security Market Data for Models of Dynamic Economies."

Journal of Political Economy 99 (1991): 225-62.

Hansen, L. P., and K. J. Singleton. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica* 50 (1982): 1269-86.

Harrison, J., and D. Kreps. "Martingales and Arbitrage in Multiperiod Securities Markets." *Journal of Economic Theory* 20 (1979): 381-408.

Jermann, U. J. "Asset Pricing in Production Economies." *Journal of Monetary Economics* 41 (1998): 257-75.

Luenberger, D. *Optimization by Vector-Space Methods*. New York: John Wiley and Sons, 1969.

Mehra, R., and E. Prescott. "The Equity-Premium Puzzle." *Journal of Monetary Economics* 15 (1985): 145-61.

Ross, S. A. "A Simple Approach to the Valuation of Risky Streams." *Journal of Business* 51 (1978): 453-75.

Shiller, R. J. "Consumption, Asset Markets, and Macroeconomic Fluctuations." *Carnegie-Rochester Series on Public Policy* 17 (1982): 203-38.

Tallarini, T. D. "Risk-Sensitive Real-Business Cycles." *Journal of Monetary Economics* 45 (2000): 507-32.

Table 1: Estimation Results

Instruments	a	t-stat.	b	t-stat	c	t.stat
			S&P			
\mathbf{w}_t	.76045	1.58	.197681	3.05	7.176833	1.84
$\mathbf{w}_t, \mathbf{R}_t$	1.016918	4.57	.266946	5.24	4.356948	2.83
$\mathbf{w}_t, \mathbf{R}_t, \mathbf{R}_{t-1}$.99578	6.47	.184143	4.70	5.253640	4.01
			Comm. Paper			
\mathbf{w}_t	.651918	.64	.370929	1.84	6.859074	1.13
$\mathbf{w}_t, \mathbf{R}_t$	1.06614	3.41	.268319	3.20	5.171290	2.25
$\mathbf{w}_t, \mathbf{R}_t, \mathbf{R}_{t-1}$	1.108862	5.29	.188931	4.94	5.847707	3.81
			System			
\mathbf{w}_t	1.069422	2.61	.353752	6.60	4.079130	1.43
$\mathbf{w}_t, \mathbf{R}_t$	1.028659	4.77	.317309	8.28	4.699451	3.19

Table 2: Pricing Tests

Asset	Instruments	Forecast Error	s.e.	TJ_T	p-value	df
S&P	$\mathbf{w}_t, \mathbf{R}_t$	-.0056	.0141	.6053	.7388	2
S&P	$\mathbf{w}_t, \mathbf{R}_t, \mathbf{R}_{t-1}$	-.018381	.01098	3.5804	.4658	4
Comm. Pap.	$\mathbf{w}_t, \mathbf{R}_t$.028166	.035378	1.7210	.4229	2
Comm. Pap	$\mathbf{w}_t, \mathbf{R}_t, \mathbf{R}_{t-1}$.014568	.013813	3.1937	.5259	4
System Results						
S&P	\mathbf{w}_t	.048451	.017015			
Comm. Pap	\mathbf{w}_t	-.006469	.024602			
				10.4974	.0148	3
S&P	$\mathbf{w}_t, \mathbf{R}_t$.053293	.023317			
Comm. Pap.	$\mathbf{w}_t, \mathbf{R}_t$.000444	.036144			
				15.2608	.0328	7

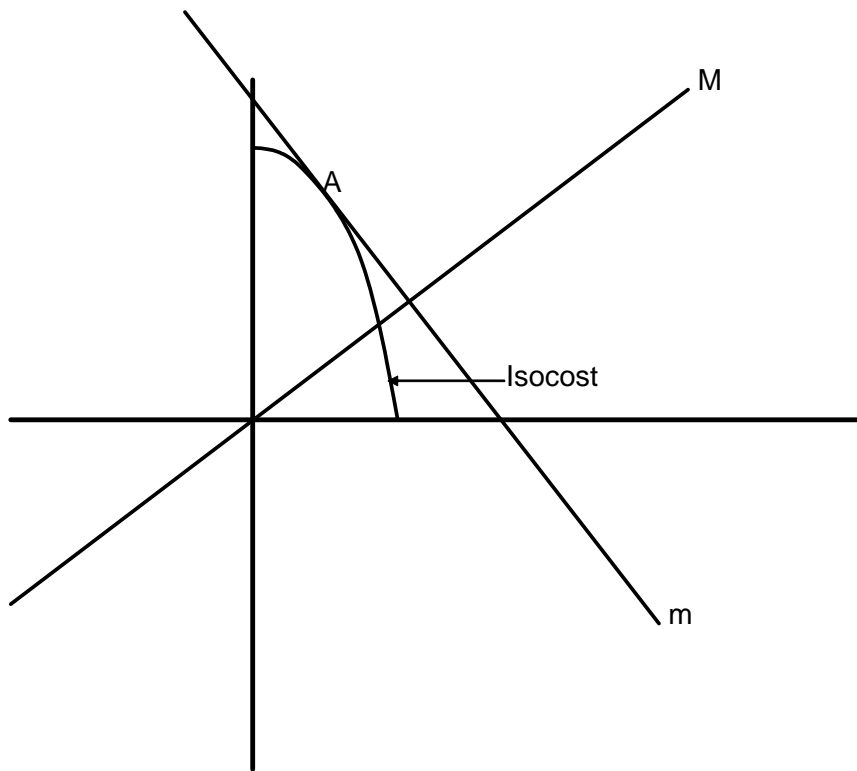


Figure 1: The Cost-based pricing kernel