Bubbles as payoffs at infinity*

Christian Gilles[1] and Stephen F. LeRoy[2]

[1] Board of Governors, Federal Reserve System, Washington, DC 20551, USA

[2] Carlson School of Management, University of Minnesota, Minneapolis, MN 55455, USA

June 1993

Revised November 1995

Revised February 1996

* This paper was circulated under the title "Stochastic bubbles in Markov economies." We acknowledge with gratitude numerous conversations with Mark Fisher, the editorial advice of David Levine and the useful comments of anonymous referees. This paper should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or its staff.

Bubbles as payoffs at infinity

Summary. We define rational bubbles to be securities with payoffs occurring in the infinitely distant future and investigate the behavior of bubble values. We extend our analysis to a setting of uncertainty. In an infinite horizon arbitrage-free model of asset prices, we interpret the money market account as the value of a particular bubble; a similar interpretation holds for other assets related to the state-price deflator and to payoffs on bonds maturing in the distant future. We present three applications of this characterization of bubbles.

1. Introduction

Formal analysis of rational speculative bubbles calls for models in which agents consume at an infinite number of dates. An immediate problem arises: if agents can trade without restriction in sequential markets, they may enjoy unbounded utility by borrowing to finance current consumption, borrowing again to repay the initial loan, borrowing yet again to repay the second loan, and so on. Such Ponzi schemes must be ruled out in some way if we are to construct equilibrium models of bubbles. Two classes of bubble models occur, depending on the strategy used to rule out Ponzi schemes.

One strategy is to impose a priori bounds on indebtedness (Santos and Woodford [1992], Magill and Quinzii [1993]).[1] For example, Magill and Quinzii assumed that agents cannot choose consumption plans which imply a positive present value of asymptotic net indebtedness. In such settings agents' choice sets are not linear spaces, implying that security prices are not necessarily representable as the values of their payoffs under linear functionals. Bubbles may then be defined as the excess of security values over the values of their payoffs. In this strand of the literature it is shown that bubbles can occur on any security that is retraded into the infinite future, as long as that security is in zero net supply. The requirement that the security be in zero net supply is needed because otherwise existence of a bubble would increase aggregate wealth without providing any

[1] Tirole (1982) is in the same spirit although, as Kocherlakota (1992) pointed out, Tirole nowhere explicitly ruled out Ponzi schemes.

new consumption opportunities. The resulting violation of Walras' law is inconsistent with equilibrium. The arbitrage that, in a linear setting, would destroy the bubble—buy the payoff generated by the security, short the security itself and never cover the position—is ruled out because it results in a positive net present value of asymptotic indebtedness.

Another modeling strategy, adopted by Gilles (1989) and Gilles and LeRoy (1992a, 1992b), maintains the convention of Arrow and Debreu that agents trade state-contingent consumption claims—securities—at one date, which is interpreted as occurring before time begins. The sequence of budget constraints of Santos–Woodford and Magill–Quinzii is replaced by a single integrated budget constraint, thereby ruling out Ponzi schemes at the outset. In the setting of Gilles and Gilles–LeRoy, choice sets are linear spaces, thereby preserving linear pricing. An attractive feature of our approach is that with linear pricing the connection to the central line of finance theory, which has asset values definitionally equal to the values of payoffs, is maintained.

In one specification of this linear approach, the choice set consists of those payoffs generated by a finite number of securities trades, together with the limit points of nets of such payoffs. This specification is adopted with varying degrees of explicitness in formal treatments of arbitrage (Clark [1993]) and discussions of the Arbitrage Pricing Theory (Ross [1976]; Gilles and LeRoy [1991]). With security values definitionally equal to the values of payoffs, the value of bubbles cannot usefully be defined as the difference between the value of securities and the value of their payoffs. Instead, the fundamental value of a security is defined as the limiting value of the initial segment of its payoff, and the bubble value is whatever is left over (Gilles [1989], Gilles and LeRoy [1992a, 1992b]).

Defining bubbles in this way creates the possibility of two types of bubbles, which we call "price bubbles" and "payoff bubbles". We analyzed price bubbles in the papers just cited; here we turn to the symmetric notion of payoff bubbles. The notion of payoff bubbles must initially appear puzzling: do not bubbles necessarily involve security prices? In our usage, the term "price" (or "price system") refers to a rule to assign values to payoffs; linear pricing then means that this rule is a linear functional. Security "prices", in contrast, are really values: they are the values taken on under the pricing rule by the particular payoffs that define securities. In speaking of bubble prices in a linear setting, economists ordinarily are referring to values, to payoffs evaluated by a price system. Therefore bubble values can in principle originate with either the price system or security payoffs.

The distinction between price bubbles and payoff bubbles can best be set out by analyzing a special case. To understand price bubbles, assume that there is no uncertainty, that time is continuous, and that all payoff streams are representable as continuous functions on $[0, \infty)$ that converge at infinity. Under this restriction the value of payoff x under an arbitrary linear functional p, $p(x)$, is always representable as

$$p(x) = p \cdot x := \int_{t=0}^{\infty} p(t)x(t)dt + p_\infty x_\infty \,,$$

where $p(t)$ is integrable, p_∞ is a scalar and x_∞ is defined as $\lim_{t\to\infty} x(t)$. The fundamental value of x under the price system p, $f_p(x)$, is defined as the limiting value of its initial segments:

$$f_p(x) := \int_{t=0}^{\infty} p(t)x(t)dt,$$

so that the bubble value of x under p, $b_p(x)$ equals the value of x less the value of its fundamental:

$$b_p(x) := p(x) - f_p(x) = p_\infty x_\infty.$$

A price bubble is identified with strictly positive p_∞, while if $p_\infty = 0$ all payoffs have values equal to their fundamental values.

To understand payoff bubbles, reverse the role of prices and payoffs here: the value of a security x under a price system p, $p(x)$, can be regarded as the value taken on by a price system p—a continuous function $p(t)$ converging at infinity, with $p_\infty := \lim_{t\to\infty} p(t)$—under a linear functional x, rather than vice-versa as above. We are led to define the payoff space as the set of linear functionals x on the space of price systems p. By symmetry, we have

$$p(x) = x \cdot p = \int_{t=0}^{\infty} p(t)x(t)dt + x_\infty p_\infty.$$

Here the interpretation is that $x(t)$, an integrable function, is the fundamental payoff and $\int_{t=0}^{\infty} p(t)x(t)dt$ is the value of the fundamental payoff. The “payoff at infinity”, x_∞, represents the bubble component of x; note that its value is determined by the long-run average price, rather than the prices at any finite set of dates.

Essentially, a payoff bubble is the part of a payoff stream that occurs "at infinity"—*i.e.*, after every finite future date. Of course, in a finite sample, looking at the realized payoffs cannot directly reveal the presence of a bubble. But postulating that security payoffs have bubble components may help explain fluctuations in security values; for example, some analysts have conjectured that the excess volatility of asset prices is indirect evidence of bubbles. Modeling bubbles as payoffs at infinity provides a rigorous device for evaluating the frequently-encountered argument that existence of valued bubbles violates a transversality condition for investors' optima—or, alternatively, allows a utility-increasing arbitrage. An explicit mathematical structure for bubbles provides the foundation for their empirical analysis by explaining clearly when theoretical arguments can and cannot be used to demonstrate their nonexistence.

From a formal point of view, then, bubbles can be modeled either as payoff bubbles or price bubbles.[2] From a substantive point of view, some economic problems are clarified by one specification, some by the other. Our concern in this paper is with payoff bubbles. In order to motivate our contention that the notion of payoff bubbles can clarify substantive economic problems, we turn to an example. We keep the treatment informal here, pending more rigorous discussion in Section 3 below.

Consider a firm which, at a given date, has a given amount of capital which earns a constant rate of return. If this firm pays most of its earnings out in dividends, dividends are initially high, but grow slowly over time. If, on the other hand, the firm retains and reinvests most of its earnings, the initial dividends are low, but they grow fast. In the absence of taxes and frictions and assuming that dividends are capitalized at the same rate as that which the firm earns on capital, the Modigliani–Miller proposition implies that the present values of all the dividend streams generated by different payout rates are equal.

But suppose that the dividend payout rate converges toward zero. The limiting dividend at any date is zero, and the present value of a sequence of zeros is, of course, zero. Apparently the Modigliani–Miller theorem is violated as dividends approach zero. By selling a zero-dividend firm for its market value at some future date, of course, a trader

[2] Although a model in which price bubbles and payoff bubbles coexist might be useful, we have not found a manageable mathematical structure in which this would be the case.

can generate a revenue that equals the present value of the dividend streams implied by positive dividend payout rates. But this creates a new problem: it is a commonplace in finance that the future sale of an asset for its market value should not affect the asset's present value, which is precisely what appears to happen here.

This puzzle may not be important, but it poses the kind of problem for which economic theory should have a convincing answer. The theory of rational speculative bubbles as developed in this paper provides such an answer. The dividend streams generated by successively lower dividend payoff rates converge, in the relevant topology, not to zero but to a payoff at infinity—in other words, to a bubble. The present value of this bubble equals the present value of any of the dividend streams generated by a positive dividend payout rate. The dividend streams do converge to zero in some topologies, to be sure; the example proves only the inconsequential fact that in such topologies the present-value operator is discontinuous.

2. Securities as charges

Bubble payoffs, then, are payoffs at infinity, while fundamental payoffs are payoffs at finite dates. The example of the zero-dividend firm shows that it would be desirable for bubbles to be the limits of sequences (or nets) of fundamental securities. This is not obviously the case with the duality that we used to introduce the notion of bubble payoffs, namely that between the space of converging functions and the space of summable functions, so it is desirable to turn to a richer setting.

In earlier work focusing on price bubbles (Gilles [1989], Gilles–LeRoy [1992a, 1992b]), we took as securities the elements of $L^\infty(M, \mathcal{M}, \mu)$, where $(M, \mathcal{M}, \mu)$ is a measure space. This linear space is the set of functionals on the index set M that are measurable with respect to $\mathcal{M}$ and bounded (except perhaps on a set of μ-measure zero). Price systems were the norm-continuous linear functionals on the security space. This space, which is denoted $\mathbf{ba}(M, \mathcal{M}, \mu)$, is known to be identifiable with the space of bounded charges on $\mathcal{M}$ that vanish on sets of μ-measure zero, where a *charge* is a finitely additive set function. If the charge is not only finitely but also countably additive, it is a measure.

With price system $p \in \mathbf{ba}$, the value of security $x \in L^\infty$ is

$$p(x) = \int_M x\, dp\,. \tag{1}$$

In order to focus on payoff bubbles, we reverse the role played by L^∞ and now interpret it as the price space. $L^\infty(M, \mathcal{M}, \mu)$ is the dual of $L^1(M, \mathcal{M}, \mu)$, the space of μ-integrable real-valued functions on M, which we now specify as the space of fundamental payoffs. This reversal of the payoff and price spaces was illustrated in the simple setting discussed in the introduction above. With price $p \in L^\infty$, the value of the fundamental payoff $x \in L^1$ is

$$p(x) = \int_M x\, p\, d\mu\,. \tag{2}$$

A symmetry between (1) and (2) becomes evident when we rewrite the latter as

$$p(x) = p(\lambda) = \int_M p\, d\lambda\,, \tag{3}$$

where λ is a measure on $\mathcal{M}$ defined by

$$\lambda(A) := \int_A x\, d\mu, \qquad \text{for all } A \in \mathcal{M}. \tag{4}$$

In view of (4), each fundamental payoff $x \in L^1$ determines a unique measure $\lambda \in \mathbf{ba}$; conversely, each measure $\lambda \in \mathbf{ba}$ determines a unique function $x \in L^1$ satisfying (4) (x is the Radon–Nikodym derivative of λ with respect to μ). Thus, L^1 is embedded into $\mathbf{ba}$ as its subset of measures; x and λ being indistinguishable, we write indifferently $p(x)$ or $p(\lambda)$, as in (3). Note that for any $p \in L^\infty$, $p(\lambda)$ is well defined by the right side of (3) for any λ in $\mathbf{ba}$ and not just for its subset of measures; but if λ is not a measure, it has no Radon–Nikodym derivative x, and therefore the representation of the value in (2) is not available.

Charges are partially ordered by setwise comparisons; *i.e.*, $\pi \geq \tau$ means $\pi(A) \geq \tau(A)$ for all $A \in \mathcal{M}$. A positive charge π (*i.e.*, $\pi \geq 0$) is called pure if $\lambda = 0$ is the only positive measure not exceeding π (*i.e.*, satisfying $\pi \geq \lambda$). By the Jordan decomposition theorem, any charge in $\mathbf{ba}$ can be expressed as the difference between two positive charges. An arbitrary charge is called pure if it equals the difference between two positive pure charges. By the Yosida–Hewitt theorem, any charge τ can be uniquely expressed as the sum of a measure λ and a pure charge π; hence, the elements of **ba** that are not measures (not fundamental securities) contain pure charges.

The payoff space $\mathbf{ba}(M, \mathcal{M}, \mu)$ contains pure charges if and only if the subset of M that supports μ is infinite, but pure charges do not necessarily admit an interpretation as payoffs in the infinite future. Indeed, when the set M does not index time, the notion of a payoff in the infinite future is not even available. Suppose for example that M is an infinite index set for the states of the world in a static model with uncertainty; then a pure charge corresponds to a payoff in a state occurring with arbitrarily small probability which has nonzero expected value. Even if M indexes time, a bubble payoff may not necessarily be paid in the infinitely distant future. Suppose for example that $M = \{0, 1/2, 3/4, \ldots, n/(n+1), \ldots\}$; then the horizon is finite and all payoffs, including "payoffs at infinity" are paid within one year. In continuous time models, in particular, there exist bubbles that are paid within a fixed time interval.

One way to guarantee that all pure charges are payoffs at infinity is to let M be a countable set indexing time, the type of commodity, and the state of the world, with a finite number of commodity types and a finite number of states at each of a countably infinite number of dates. To see this, note that a pure charge on M associates zero weight with any finite subset, which means that a bubble pays no dividend at any finite time. In other words, all payoffs occur "at infinity" (but note that, contrary to this loose description, there are many distinct securities that pay at infinity).

We recognize that our characterization of securities as elements of a charge space is unconventional. There is precedent for specification of securities as measures, for example in the commodity differentiation literature (Mas-Colell [1975]), and in any case measures are equivalent to integrable functions (precisely the result that yields the embedding of L^1 in **ba**), which is a conventional specification. The novel element here is our specification that securities may be modeled as set functions that are not countably additive. But nothing about the nature of securities implies countable additivity—for example, nothing about a share of stock requires us to assume that its value equals the limiting value of the initial segments of the dividend stream it generates. Assuming that, on the contrary, the two are always equal would imply that the value of the stock fundamental necessarily equals the value of the stock itself under any price system. Rather than rule out bubbles without benefit of economic reasoning in this way, we prefer to leave room for the existence of bubbles by allowing securities which are pure charges, and then let economic considerations determine whether these are valued.

The assumption that the fundamental component of a security payoff is integrable (or summable, to take its ℓ^1 representation)[3] implies that securities like χ_M are inadmissible.[4] This is so because the dividend stream χ_M would have infinite value under the price system χ_M which, as an element of L^∞, is admissible. However, we can allow constant and even increasing dividend streams as payoffs by the simple device of redefining the norm, as we now demonstrate.

By defining the norm of a function x to be

$$\|x\|_\rho^1 := \int_{t=0}^{\infty} |x(t)| e^{-(\rho-1)t} dt$$

for some constant ρ rather than, as in L^1, $\int_{t=0}^{\infty} |x(t)| dt$, we can admit securities with growing payoffs, as long as the growth rate is below ρ. Define L_ρ^1 to be the space of functions x such that $\|x\|_\rho^1$ is finite; in other words, $x \in L_\rho^1$ if $\widehat{x} \in L^1$, where $\widehat{x}(t) := x(t)e^{-(\rho-1)t}$. Choosing L_ρ^1 with $\rho > 1$ as the space of fundamental securities allows us to include χ_M as a security, but at a cost: when the security space expands, the price space must contract. The space of admissible price systems, denoted L_ρ^∞, becomes the set of functions p such that

$$\|p\|_\rho^\infty := \sup_t |p(t)| e^{(\rho-1)t}$$

is finite; in other words, $p \in L_\rho^\infty$ if $\widehat{p} \in L^\infty$, where $\widehat{p}(t) := p(t)e^{(\rho-1)t}$. In economic terms, accommodating securities with payoffs that rise over time requires adopting price systems that imply sufficient discounting, so that all securities remain finite-valued. The relations among L_ρ^1, L_ρ^∞ and $\mathbf{ba}_\rho$ are exactly the same as those among L^1, L^∞ and $\mathbf{ba}$.[5]

---

[3] With M countable, we tacitly assume that $\mathcal{M}$ is the set of all subsets of M and μ the counting measure. In this case, L^1, the space of absolutely integrable functions, becomes ℓ^1, the space of absolutely summable sequences, and L^∞ becomes ℓ^∞, the space of bounded sequences.

[4] Here χ_A is the indicator function of $A \subset M$, which takes on value 1 on A and 0 elsewhere; thus, χ_M takes on unit value everywhere.

[5] The relation between $\mathbf{ba}$ and $\mathbf{ba}_\rho$ is easily characterized: for each $\tau \in \mathbf{ba}_\rho$, there corresponds a

Specification of the security space as the space of charges on an infinite index set implies that there always exist securities which are bubbles or which have bubble components. Whether or not these securities exist in positive net supply depends on the specification of production sets and endowments, and whether or not these bubbles are valued depends on the equilibrium price system. For example, with $\mathbf{ba}_1$ as the security space (*i.e.*, $\rho = 1$), any pure charge is a free good (and therefore economically inconsequential) under the price system $e^{-(\rho-1)t}$, $\rho > 1$; but under the price system χ_M any nonzero positive pure charge has positive value. Correspondingly, whether valued bubbles exist under any fixed price system depends on whether the security space is defined to be as large as possible consistent with that price system. For example, under the price system e^{-2t}, if the security space is $\mathbf{ba}_\rho$ with $\rho < 3$, then there are no valued bubbles. On the other hand, with the same price system there exist valued bubbles in $\mathbf{ba}_3$.

It follows that whether such models as Lucas (1978) or Mehra and Prescott (1985) contain valued bubbles is a purely formal question, depending solely on how the model builder chooses to specify the security space. On the other hand, whether a particular security such as the market portfolio contains a bubble component is very much a substantive question.

Because bubbles have zero payoffs at any single date, they are difficult to characterize formally. We get around this problem by characterizing bubbles as the limit points of nets of fundamental securities, since fundamentals are easy to characterize. Aside from mathematical convenience, this treatment stems from the root of our interest in bubbles (apparent in the motivating example of the zero-dividend firm): we want to know what happens to fundamental securities and their values as the payoff becomes increasingly distant or spread out over time. To illustrate the convergence of fundamental securities to bubbles, consider the sequence $\{\chi_i, i = 1 \ldots \infty\}$, which is a sequence of discount bonds. As shown in the appendix, this sequence admits a bubble as a limit point, not in the norm topology, but in the weak* topology of $\mathbf{ba}$.

---

unique $\widehat{\tau} \in \mathbf{ba}$, with the correspondence determined by

$$\int p\, d\tau = \int \widehat{p}\, d\widehat{\tau}, \quad \text{for all } p \in L^\infty_\rho,\ \widehat{p} \in L^\infty \text{ such that } \widehat{p}(t) = p(t)e^{(\rho-1)t}.$$

3. An application: "Zero"–dividend firms.

We now return to the problem of the zero-dividend firm introduced in Section 1 above. Consider the valuation of a firm initially consisting of one unit of capital which generates earnings at constant rate $\rho - 1$. The gross interest rate is also constant and equal to ρ, implying that the price system is $e^{-\rho t}$. This firm pays dividends at constant rate γ (expressed as a proportion of current earnings), and with its retained earnings it acquires new capital which has constant internal rate of return ρ.

This firm's dividend is $\rho\gamma$ at $t = 0$, and dividends grow exponentially at rate $\rho(1-\gamma)$, implying that the dividend stream is:

$$x_\gamma(t) = \rho\gamma e^{\rho(1-\gamma)t},$$

where $x_\gamma(t)$ is the dividend at date t generated by a firm with payout rate γ. To accommodate such a dividend stream for any γ in $(0, 1)$, let the space of fundamental payoffs be L^1_ρ. For $\gamma > 0$ the Miller–Modigliani proposition that the initial value of the firm is independent of γ is easily verified:

$$p(x_\gamma) = \int_{t=0}^{\infty} e^{-\rho t} x_\gamma(t) dt = 1.$$

Now let γ approach zero. For any positive γ the above proposition implies that the value of the firm is 1, yet x_γ converges (pointwise) to a zero dividend stream, which has zero value.

This apparent discontinuity in the value of the firm as γ approaches zero stems from the assumption that the dividend stream converges to zero. Clearly it does so in the topology of pointwise convergence, but the discontinuity of the valuation functional that the example implies should lead us to question the appropriateness of that choice of topology, not to conclude that some kind of deep paradox has been unearthed.

Suppose that we choose instead the topology induced by the norm

$$||x||^1_\rho = \int_{t=0}^{\infty} |x(t)| e^{-\rho t} dt. \tag{5}$$

This norm, already defined in Section 2, has the convenient property that the norm of a (positive) dividend stream coincides with its value, immediately guaranteeing the continuity of the price functional (this is saying no more than $p \in L^\infty_\rho$, which is obvious). Under the norm (5), the net $\{x_\gamma\}$ does not converge to the zero vector, since $\|x_\gamma - 0\| = 1$ for all γ. The situation is similar to that of the sequence of discount bonds discussed in the appendix: that sequence does not have any limit points in the norm topology of L^1, and neither does the net $\{x_\gamma\}$ in the norm topology of L^1_ρ.

The norm dual of L^1_ρ is L^∞_ρ. The value of $y \in L^1_\rho$ under the price system $q \in L^\infty_\rho$ is $q \cdot y := \int p(t)y(t)\,dt$. The space L^1_ρ can therefore be endowed with the weak topology, under which a net $\{y_\alpha\}$ converges to y if $\{p \cdot y_\alpha\}$ converges to $p \cdot y$ for all $p \in L^\infty_\rho$. As in the case of the discount bonds, again, the net $\{x_\gamma\}$ has no limit point in the weak topology of L^1_ρ.

It comes as no surprise that our solution is to enlarge the space of dividend streams from the L^1_ρ space of sequences normed by (5) to the associated charge space $\mathbf{ba}_\rho$, and to consider convergence in the weak* topology (if this is done the fundamentals x_γ appear as the measures μ_γ, with the correspondence defined in (4)).[6] Because the elements of the net $\{\mu_\gamma\}$ are bounded in norm, the net must have a weak* limit point μ_0, by Alaoglu's theorem (which is stated in the appendix).

It is easy to see that μ_0 must be a pure charge. For any charge τ, $\int \chi_t d\tau$ is the t-th component of the fundamental of τ. Since $x_\gamma(t) = \int \chi_t\, d\mu_\gamma$ converges to zero, a limit that must equal $\int \chi_t\, d\mu_0$ (by definition of weak* convergence, tested with the price system χ_t), it must be the case that μ_0 has a zero fundamental.

The net $\{x_\gamma\}$, then, admits as limit points, not zero, but rational bubbles. That a net of fundamental securities can converge to a bubble proves only the inconsequential fact that the projection operator mapping each security onto its fundamental is discontinuous in the weak* topology; equivalently, it reflects the fact that the linear subspace of measures is not weak* closed in $\mathbf{ba}$.

In a precise sense, by setting $\gamma = 0$ the firm pays infinite dividends in the infinite future. Because the current value of this bubble is unity rather than zero, the apparent anomalies attending "zero"-dividend firms disappear.

---

[6] In the weak* topology of $\mathbf{ba}_\rho$, a net of charges $\{\tau_\alpha\}$ converges to τ if $p \cdot \tau_\alpha$ converges to $p \cdot \tau$ for all $p \in L^\infty_\rho$. As a result, when restricted to measures, this topology is identical to the weak topology of L^1_ρ.

4. Uncertainty

The foregoing characterization of bubbles extends directly to a setting of uncertainty. We adopt the uncertainty model described in Duffie (1992), to which we refer the reader for the terms we leave undefined. We fix a probability space $(\Omega, F, \mathcal{M})$ and the standard filtration $\{F_t : t \geq 0\}$ generated by a standard Brownian motion $W(t)$ in $\Re^d$. An *adapted process* is a function $f : \Omega \times [0, \infty)$ such that $f(t) := f(\cdot, t)$ is measurable with respect to F_t for all t. Suppose that the payoff stream of security x is characterized by its gain $g_x(t)$, where $g_x(t)$ is an adapted process that equals the cumulative payoff from date 0 to date t. Doing this allows us to handle simultaneously the case of a continuous and atomless payoff stream, in which case $g_x(t)$ is differentiable (on each sample path) and also the case of a single lump-sum payoff (or, in an obvious extension, a finite number of lump-sum payoffs), in which case $g_x(t)$ has discontinuities. We will restrict our attention to payoff streams that satisfy $\int_0^\infty dg_x^2(t) < \infty$, almost surely.

In the absence of arbitrage, there exists a positive-valued adapted process m, with $m(0) = 1$, such that the value of x at t, $V(t, x)$, is given by

$$\frac{V(t,x)}{m(t)} = E_t^{\mathcal{M}} \left[\int_{\tau=t}^{\infty} \frac{1}{m(\tau)} dg_x(\tau) \right]. \tag{6}$$

Here $E_t^{\mathcal{M}}$ denotes expectation taken with respect to the natural or physical probability measure $\mathcal{M}$ restricted to the σ-field F_t. If markets are complete m is unique. If markets are incomplete, there exist many processes m satisfying (6); for example, the stochastic process for the marginal utility of any expected-utility-maximizing agent (whose optimum is interior) can serve as $1/m$ in (6).

The process m is called the *state-price deflator*.[7] Eq. (6) implies that for any payoff stream x satisfying $dg_x(t) = 0$ for $t < T$, the discounted value $V(t, x)/m(t)$ is a martingale under $\mathcal{M}$ until time T.

[7] The literature does not adopt a standard terminology here: Duffie (1992) also used the term "state-price deflator", but entered the term corresponding to our $m(\tau)$ in the numerator rather than the denominator, suggesting that m should be called a factor rather than a deflator. Many authors use the term "pricing kernel", with various related meanings. For example, Hansen and Jagannathan (1991) characterized the pricing kernel as the representative agent's marginal rate of substitution for consumption at successive dates, a usage that has no analogue in continuous-time settings.

We saw in Sec. (2) that in the deterministic case payoff bubbles were defined by using the equilibrium price system to define a norm on the payoff space, and then characterizing bubbles as limit points of fundamental payoffs. Exactly the same strategy is employed here. Define the norm of payoff x by

$$\|x\|_m^1 \ := \ \int_{\tau=0}^{\infty} \frac{\|dg_x(\tau)\|_2}{\|m(\tau)\|_2}. \tag{7}$$

Here $\|z\|_2$ denotes the Hilbert space norm $\sqrt{E(z^2)}$.

Define the space of *fundamental payoffs* as the set of payoffs with finite norm, as defined by (7). A continuous linear functional P on the L^1 space of fundamental payoffs is representable by a member p of the L^∞ space of bounded adapted L^2-valued functions on $[0, \infty)$, where $P(x) = \int_0^\infty E_0^{\mathcal{M}}[p(t)\, dg_x(t)]$. On this dual space, the norm is defined by

$$\|p\|_m^\infty \ := \ \text{ess sup}_\tau \|p(\tau)\|_2 \ \|m(\tau)\|_2.$$

Because $1/m$ is a convex function of m, this norm and Jensen's inequality imply that

$$\|1/m\|_m^\infty \geq 1,$$

with equality if and only if $m(t)$ is a deterministic process. There is no guarantee that $\|1/m\|_m^\infty$ is finite, however. If it is infinite, there exist fundamental payoffs to which $1/m$ assigns infinite value. With such a payoff x, the variance of the payoff at t exists almost everywhere (since $\|dg_x(t)\|_2$ is well-defined), but it rises too fast (or does not fall fast enough), a problem which does not arise in the absence of uncertainty. We call an adapted process $p(t)$ an admissible price system if it is positive and $\|p\|_m^\infty < \infty$ and we *assume* that $1/m$ is an admissible price system so as to guarantee that all finite-norm payoffs have finite value.

The state-price deflator $m(t)$ can be interpreted as the value of a fundamental payoff. To see this, consider the value at t of μ_τ, a claim to a single payment equal to $m(\tau)$ at the fixed future date τ. In this case the gain process g_{μ_τ} is a single jump equal to $m(\tau)$ at time τ, so (6) reduces to $V(t, \mu_\tau)/m(t) = E_t^{\mathcal{M}}(1) = 1$ (trivially, a martingale), and therefore $V(t, \mu_\tau) = m(t)$, for any $t < \tau$ and any $\tau > 0$. Note that, since $\|\mu_\tau\|_m^1 = \|m(\tau)\|_2/\|m(\tau)\|_2 = 1$ for any τ, μ_τ is an admissible fundamental payoff. In fact the payoff norm was designed to ensure that all of these payoffs have unit norm (which implies, as we will see below, that they admit a bubble as a limit point).

As it happens, there is nothing special about portfolio $m(t)$ that allows it alone to be used as a deflator: any positive-valued portfolio satisfying the numeraire choice $m(0) = 1$ would do, although the probability measure with respect to which the expectation is taken depends on which security or portfolio serves as a deflator. To illustrate this statement, we will consider pure discount default-free bonds. Let χ_τ denote a deterministic payoff of one unit of consumption at date τ, so that the gain process $g_{\chi_\tau}(t)$ is a unit step at $t = \tau$. According to equation (6), the value $V(t, \chi_\tau)$, for $t < \tau$, is given by

$$\frac{V(t, \chi_\tau)}{m(t)} = E_t^{\mathcal{M}} \left[\frac{1}{m(\tau)} \right]. \tag{8}$$

The yield to maturity on this bond is given by

$$r(t, \tau) := -\frac{\log(V(t, \chi_\tau))}{(\tau - t)},$$

and the short rate $r(t)$ is defined by

$$r(t) := \lim_{\tau \to t} r(t, \tau).$$

Now define the *money market account* $b(t)$, which keeps track of the value at t of investing at the short rate, by

$$b(t) := \exp\left(\int_0^t r(s) ds \right). \tag{9}$$

Since $b(0) = 1$ and $b(t) > 0$, we can use $b(t)$ as a deflator. Absence of arbitrage implies and is implied by the existence of a probability measure $\mathcal{B}$, equivalent to $\mathcal{M}$ in the sense that they both assign measure zero to the same sets, such that the value of any fundamental payoff x, $V(t, x)$, satisfies

$$\frac{V(t, x)}{b(t)} = E_t^{\mathcal{B}} \left[\int_{\tau=t}^{\infty} \frac{1}{b(\tau)} dg_x(\tau) \right]. \tag{10}$$

As before, $b(t)$ can be interpreted as the value at t of a fundamental payoff, because $b(t) = V(t, \beta_\tau)$, where β_τ is a claim to a single payment of $b(\tau)$ at the fixed date τ, for any $\tau > t$. The validity of this statement follows easily from (10). The probability measure $\mathcal{B}$ is known as the *risk-neutral probability measure*, since it gives the value of any security as the discounted value of its payoff deflated by the interest rate, as would be the case under the natural probabilities $\mathcal{M}$ if agents were risk neutral (in a sense to be discussed below).

Another deflator is provided by the value at t of the fundamental payoff π_τ, defined as a default-free discount bond that matures at τ, and has payoff then of $1/V(0,\chi_\tau)$. From (8) this price is given by

$$p(t,\tau) := V(t,\pi_\tau) = \frac{V(t,\chi_\tau)}{V(0,\chi_\tau)} = m(t)\frac{E_t^{\mathcal{M}}[1/m(\tau)]}{V(0,\chi_\tau)}. \tag{11}$$

The scale of the payment at maturity was chosen to ensure that, for any τ, $p(0,\tau)=1$. Because $p(t,\tau)>0$ for all $t<\tau$, it can be used as a deflator (until time τ). Define the probability measure associated with this *τ-year bond deflator* $p(t,\tau)$ as $\mathcal{P}(\tau)$, as is appropriate since each τ results in a different probability measure. Then the value of any payoff x which does not extend beyond τ obeys

$$\frac{V(t,x)}{p(t,\tau)} = E_t^{\mathcal{P}(\tau)}\left[\int_{s=t}^{\tau}\frac{1}{p(s,\tau)}dg_x(s)\right], \quad \text{for } t\le\tau.$$

The various probability measures are equal whenever the deflators are equal, which occurs under several simplifying assumptions. If interest rates are deterministic, then $\mathcal{B}=\mathcal{P}(T)$ for all T, since then the deflators are equal as well. Similarly, $\mathcal{M}=\mathcal{P}(T)$— and $m(t)=p(t,T)$—if the representative agent is risk neutral at horizon T. This will occur, for example, if agents maximize the expected value of consumption at date T, which will result in risk-neutral pricing for all payoffs at date T, but generally not for payoffs at other dates, except under strong restrictions on preferences or production technologies.

To investigate the interpretation of $\mathcal{M}=\mathcal{B}$, assume that $1/m(t)$ follows the diffusion process

$$\frac{d(1/m(t))}{1/m(t)} = \gamma_m(t)dt + \sigma_m(t)\cdot dW(t), \tag{12}$$

where $W(t)$ is the standard Wiener process in $\Re^d$ that describes the evolution of fundamental sources of uncertainty (and that generates the filtration $\{F_t,\ t\ge 0\}$), and

$\sigma_m(t) \cdot dW(t)$ equals $\sum_i \sigma_{mi}(t)\ dW_i(t)$. Similarly, let $V(t,x)$, the value of some payoff stream x, follow the diffusion:

$$\frac{dV(t,x)}{V(t,x)} = \gamma_x(t)dt + \sigma_x(t) \cdot dW(t).$$

By Ito's lemma, the drift of $V(t,x)/m(t)$ is given by $\gamma_x(t) + \gamma_m(t) + \sigma_x(t) \cdot \sigma_m(t)$. If $g_x(t) = 0$ for $t < T$, then $V(t,x)/m(t)$ is a martingale until time T under the natural probability measure $\mathcal{M}$, so that this drift must equal zero, implying that

$$\gamma_x(t) = -\gamma_m(t) - \sigma_x(t) \cdot \sigma_m(t). \tag{13}$$

Now let $x = \beta_T$, the money-market account interpreted as a claim to a single payoff equal to $b(T)$ payable at time T. As noted earlier, $V(t, \beta_T) = b(t)$, which from equation (9) and Ito's lemma follows the process

$$\frac{db(t)}{b(t)} = r(t)dt.$$

In other words, the money market account has drift $\gamma_b(t) = r(t)$ and is predictable ($\sigma_b(t) = 0$). Substituting these values for γ_x and σ_x in (13), we conclude:

$$\gamma_m(t) = -r(t). \tag{14}$$

Substituting (14) in (13), the latter becomes

$$\gamma_x(t) - r(t) = \sigma_x(t) \cdot (-\sigma_m(t)), \tag{15}$$

which is recognized as the security market line of the consumption CAPM. Because of this CAPM relation, we refer to $-\sigma_m(t)$ as the market price (strictly, prices) of risk.

Now consider the stochastic process $b(t)/m(t) = V(t, \beta_T)/m(t)$, which, as noted in the discussion of eq. (6), is a martingale under $\mathcal{M}$ until time T, for arbitrary T. Since $b(t)$ is predictable (has zero instantaneous variance), Ito's lemma implies that $b(t)/m(t)$ has volatility $(b(t)/m(t))\sigma_m(t)$ under $\mathcal{M}$, so $b(t) = m(t)$ if and only if $\sigma_m(t) = 0$ for all t. The condition that the market price of risk equals zero—equivalently, that the state-price deflator is predictable—is often identified with "risk-neutral pricing".[8] Despite

[8] Equation (15) proves the well-known fact that the local expectations hypothesis, as defined by Cox, Ingersoll and Ross (1981), is valid if and only if the market price of risk is zero, or, as we now see, if and only if the state-price deflator is predictable.

this characterization, it should be noted that under the restriction $\sigma_m(t) = 0$ actuarial pricing occurs only for payoffs with an infinitely short horizon. As just seen, we can equally well identify risk-neutral pricing with a T-year horizon instead of an instantaneous horizon, implying $\mathcal{M} = \mathcal{P}(T)$ instead of $\mathcal{M} = \mathcal{B}$. Risk neutral pricing at all horizons simultaneously (*i.e.*, $\mathcal{M} = \mathcal{B} = \mathcal{P}(T)$ for all T) obtains if and only if interest rates are deterministic (so $\mathcal{B} = \mathcal{P}(T)$ for all T) and the state-price deflator $m(t)$ is predictable (so $\mathcal{M} = \mathcal{B}$).

5. Stochastic bubbles

Now we are in a position to characterize bubbles in stochastic environments. We have defined an L^1 space of payoffs: those stochastic processes with finite norms, as defined by (7). Now we embed this space in the corresponding charge space $\mathbf{ba_m}$. Then, just as in the deterministic case, the fundamental payoff associated with any positive payoff x is the largest fundamental which is less than or equal to x, the bubble component being whatever is left over. A pure bubble is a payoff with a zero fundamental.

To define our first stochastic bubble, consider the payoff μ_τ defined earlier as a claim on a single payment equal to $m(\tau)$ at date τ. Recall that the value of this payoff satisfies $V(t, \mu_\tau) = m(t)$, for all $t \leq \tau$. The set $\{\mu(\tau),\ \tau \geq 0\}$ constitutes a net, each element of which has unit norm. Therefore Alaoglu's theorem implies the existence of a charge μ such that $\{\mu(\tau)\}$ converges to μ along some subnet. Call this charge the *state price deflator bubble*. Now, μ_τ admits μ as a limit point, $V(t, x)$ is continuous in x, and $V(t, \mu_\tau) = m(t)$ for all τ; therefore $V(t, \mu) = m(t)$ as well.

We thus have a generalization of the concept of a state-price deflator that works for payoffs with any finite horizon, or with an infinite horizon: the state-price deflator at any date is the value at that date of a self-financing portfolio the payoff of which is the state-price deflator bubble μ in the infinite future.

The money-market account bubble can be defined in exactly the same way: consider the security β_τ which is a claim on $b(\tau)$ at τ and nothing else. By Alaoglu's theorem, to show that the net $\{\beta_\tau,\ \tau \geq 0\}$ admits a nontrivial bubble β as a limit point, it suffices that the net be bounded in norm. To prove the claim that $\|\beta_\tau\|_m^1 \leq \|\mu_\tau\|_m^1 = 1$, consider the stochastic process for $m(t)$ which is implied by that for $1/m(t)$ given in (12) (recall that $\gamma_m(t) = -r(t)$):

$$\frac{dm(t)}{m(t)} = (r(t) + \sigma_m(t) \cdot \sigma_m(t))dt - \sigma_m(t) \cdot dW(t).$$

The process for $b(t)$ has thus lower drift, namely $r(t)$, and lower variance, namely zero, than the process for $m(t)$. Given that $m(0) = b(0) = 1$, we conclude that $\|b(t)\|_2 \leq \|m(t)\|_2$, for any t. Since $\|\beta_\tau\|^1_m = \|b(\tau)\|_2/\|m(\tau)\|_2$, the claim is proved. We thus have shown the existence of the *money market account bubble* β, and any security can be valued as the expected value of its payoff discounted by the value of the money-market account, where the expectation is taken with respect to the risk-neutral probability measure $\mathcal{B}$.

Finally, we can define $\pi(\tau)$ as the payoff consisting of $p(\tau,\tau) = 1/V(0,\chi_\tau)$ at date τ, and nothing at any other date. Again, the net $\{\pi(\tau)\}$ is bounded in norm, since

$$\|\pi_\tau\|^1_m = \frac{1}{V(0,\chi_\tau)\|m(\tau)\|_2} \leq \frac{1}{E_0^{\mathcal{M}}[1/m(\tau)]E_0^{\mathcal{M}}[m(\tau)]} \leq 1\,,$$

where the first inequality follows from $\|m(\tau)\|_2 = \sqrt{E_0^{\mathcal{M}}[m(\tau)^2]} \geq E_0^{\mathcal{M}}[m(\tau)]$ and $V(0,\chi_\tau) = E_0^{\mathcal{M}}[1/m(\tau)]$, while the second is Jensen's inequality (the terms are all equal if and only if $m(\tau)$ is deterministic). Therefore, the net $\{\pi(\tau)\}$ admits a weak* limit π, which limit we call the *very long discount (VLD) bubble* (the term comes from Kazemi [1992]). Now, by definition $p(t,\tau) = V(t,\pi_\tau)$; since π is a limit point of the net $\{\pi(\tau)\}$ and $V(t,\cdot)$ is continuous, $p(t,\tau)$ has a corresponding limit point, say $p(t) = p(t,\infty)$, which is the value at time t of the bubble π. Since $p(t)$ is positive, it is a valid deflator, and there exists an equivalent martingale measure, say $\mathcal{P}(\infty)$, such that the value of any payoff x satisfies

$$\frac{V(t,x)}{p(t)} = E_t^{\mathcal{P}(\infty)}\left[\int_{\tau=t}^{\infty} \frac{1}{p(\tau)} dg_x(\tau)\right].$$

Now, as promised, we derive some conditions for equality among the various bubbles we have defined. First, we have that $\beta = \mu$ if and only if the state price deflator is predictable. To see that the condition is necessary, recall that if the bubbles β and μ are equal, so must be their values at date t, $b(t)$ and $m(t)$, for any t. We showed above that this equality occurs if and only if $m(t)$ is predictable. Sufficiency is obvious, since

$\sigma_m(t) = 0$ implies that $b(t) = m(t)$, but because β and μ are just the respective limit points of $\beta(\tau)$ and $\mu(\tau)$ (which are claims on $b(\tau)$ and $m(\tau)$ payable at time τ), these must be equal also (assuming the same subnet is used in each case).

Turn now to conditions for $\beta = \pi$. This equality holds if and only if predictions about $1/b(\tau)$, for large τ, are not updated in light of current information under the equivalent martingale measure $\mathcal{B}$. Deterministic interest rates imply this condition, but the nonrandomness assumption is stronger than needed. To understand the absence of updating condition, note that if we specialize (10) twice (first to $x = \chi_\tau$ and then to $t = 0$), there results

$$\frac{V(t,\chi_\tau)}{b(t)} = E_t^{\mathcal{B}}\left[1/b(\tau)\right] \tag{16}$$

and

$$V(0,\chi_\tau) = E_0^{\mathcal{B}}\left[1/b(\tau)\right]. \tag{17}$$

Dividing (16) by (17), there results

$$V(t,\pi(\tau)) = \frac{V(t,\chi_\tau)}{V(0,\chi_\tau)} = b(t)\frac{E_t^{\mathcal{B}}[1/b(\tau)]}{E_0^{\mathcal{B}}[1/b(\tau)]}. \tag{18}$$

Now let τ become large. $V(t,\pi(\tau))$ converges to the value at t of π, and $b(t)$ is the value at t of β, so if $\pi = \beta$ these values approach equality, implying that the ratio of expectations in (18) approaches unity. Conversely, if the ratio of expectations approaches unity, then as τ gets large, $V(t,\pi_\tau)$ approaches $b(t)$, which is the payoff at time t on β_t. Therefore, for any subnet of $\{\pi_\tau\}$ converging to π, there exists a subnet of $\{\beta_\tau\}$ converging to π as well. In other words, by considering appropriate subnets, we have $\beta = \pi$.

We have $\pi = \mu$ under analogous conditions, except that the absence of updating applies to the pricing kernel $1/m(t)$ under $\mathcal{M}$ rather than to $1/b(t)$ under $\mathcal{B}$, since (11) replaces (18). Specifically, we have

$$V(t,\pi(\tau)) = m(t)\frac{E_t^{\mathcal{M}}[1/m(\tau)]}{E_0^{\mathcal{M}}[1/m(\tau)]}.$$

The condition for absence of updating for $\pi = \mu$ is similar to that for $\pi = \beta$ but not identical. In the latter case the condition is satisfied if $b(t) = f(t)\alpha(t)$, where $f(t)$ is deterministic and $\alpha(t)$ is predictable, stationary and ergodic under $\mathcal{B}$. In the former case the condition is satisfied if $m(t) = f(t)\alpha(t)$, where again $f(t)$ is deterministic and $\alpha(t)$ is stationary and ergodic under $\mathcal{M}$, but now $\alpha(t)$ need not be predictable since μ is not generally predictable.

6. Applications

In earlier sections we argued that bubbles play a role whenever infinite-horizon models are used (and, we will see in subsection (6.3), not only then). Making the role of bubbles explicit can clarify otherwise puzzling results. For example, we suggested that by regarding firms which retain all their earnings as paying out an infinite dividend in the infinite future, rather than as not paying a dividend, the apparent discontinuity of firm value as a function of the dividend payout rate is avoided, and the Miller-Modigliani theorem is extended to firms that retain all their earnings.

In this section we indicate some other areas of application of the theory of bubbles. In the interest of brevity the arguments are sketched rather than stated formally.

6.1. Measuring the state-price deflator Kazemi (1992) proposed using the value of a VLD bond as a proxy for the state-price deflator, on the grounds that fluctuations in the value of bonds with payouts in the very distant future are dominated by fluctuations in the current marginal utility of consumption. In our terms Kazemi assumed that the VLD bubble π equals the state-price deflator bubble μ, so that the value of the VLD bubble equals the reciprocal of the current marginal utility of consumption of any agent who consumes at an interior point. Necessary and sufficient conditions for the validity of this use of π as a proxy for μ were stated in Section 4 above.

In a footnote Kazemi suggested that current values of even stochastic payouts can be used as proxies for μ, as long as the payout occurs in the distant future. The results of this paper show that Kazemi is incorrect here: it is precisely the argument of this paper that different bubbles have different current values, and most bubbles will have values that differ from the reciprocal of the marginal utility of consumption. To see this, consider a firm which retains all its earnings and is subject to firm-specific productivity shocks and shocks to prices of substitute goods, for example. This firm can be considered to have a dividend payout consisting of a bubble, but its current value will be determined by the current realization of the productivity shocks and the shocks to prices of substitutes as well as consumers' marginal utility of consumption.

6.2. "Long interest rates can never fall" Dybvig, Ingersoll and Ross (1995) proved that long-term interest rates can never fall. To understand the intuitive content of their result, consider a discrete-time uncertainty tree with one state at date 0 and two states at dates 1, 2, ..., so that all uncertainty is resolved at date 1. Suppose that state-claim prices equal 1 at date 0 and at each node on the upper branch of the uncertainty tree. State-claim prices on the lower branch equal 2^{-t} at date t, for $t = 1, 2, \ldots$. Thus in date 1 the gross one-period interest rate is revealed to be 1 on the upper branch and 2 on the lower branch; after date 1 these interest rates never change .

Clearly, the yield to maturity on a t-period pure discount bond issued at date 0 equals $(1 + 2^{-t})^{1/t}$; this value converges to 1 when t converges to infinity, so call this limit the yield to maturity on a VLD bond issued at date 0. Then on date 1 the interest rate on a new VLD bond is either 1, as on date 0, or 2. Thus interest rates may rise; they cannot fall. Essentially, this occurs because the date–0 cost of buying one unit of consumption at date t is, for high t, dominated by the lowest of the possible interest rate realizations.

When the rate of interest is 1, the original VLD bond has value 1 and yield to maturity 1 (so there is no need to issue a new VLD bond to get this result), but when the rate of interest rises to 2, the original VLD bond issued at date 0 becomes worthless (the payoffs on the sequence of discount bonds used to define the bubble do not increase fast enough) and therefore its own yield to maturity is not well defined. Observe that 1 is the correct price at date 0 for an asset that will be worth 1 in the low interest rate state and 0 in the other state.

To recast Dybvig, Ingersoll and Ross's result in the language of this paper, observe that the duration of the VLD bubble is infinite. Thus an increase in prevailing long-term interest rates corresponds to a decrease in the value of a VLD bubble to zero. Similarly, a decrease in prevailing long-term interest rates would correspond to an infinite increase in value; the fact that an infinite increase in value cannot occur implies the impossibility of a decrease in yields.

As the example indicates, the yield to maturity on a VLD bubble can never change: the fact that the duration of a VLD bubble is infinite implies that its yield to maturity stays the same over time, or the security becomes worthless, in which case its yield to maturity is undefined.

6.3. Doubling Strategies Up to now we have interpreted the index set as representing time. However, the same mathematical structure applies when the index set refers to alternative possible events, and the application of the theory of bubbles in this context produces useful insights.

Consider a gambling strategy consisting of betting 1 at even odds, and then doubling the bet after each loss. The game ends with the first win. A naive interpretation of this game is that it yields a payoff of 1 with probability 1, and therefore is an arbitrage. A more sophisticated argument would reject the implication that a potentially infinite sequence of fair-game bets can possibly produce an arbitrage. This argument is made formal by observing that the sequence of payoffs to the truncated games does not converge, implying that the payoff to the game as a whole is not well defined. To see this, observe that the cumulated payoff for the game truncated after t stages consists of 1 with probability $1 - 2^{-t}$ and $-(2^t - 1)$ with probability 2^{-t}. This payoff converges pointwise to 1 except on a set of probability measure zero. However, just as with the zero-dividend firm, pointwise convergence is not the relevant criterion. In the L_2 norm the payoff to the doubling strategy is not well defined since the payoff variance becomes infinite: thus the doubling strategy is inadmissible.

A different analysis emerges if one imposes the L_1 norm on payoffs rather than the L_2 norm. It remains true that the sequence of payoffs to the truncated games diverges in the L_1 norm. However, under the L_1 norm, the norm of the payoff of the truncated doubling strategy is bounded above by 2. Therefore, reinterpreting the space of payoffs as **ba** rather than L_1, Alaoglu's theorem guarantees that a subnet of the sequence of payoffs to the truncated games converges. The limiting payoff is 1 on a set of measure 1, and a (negative) bubble on a set of measure zero. Despite occurring on a set of measure zero, the bubble contributes exactly -1 to the expectation of the game, which is therefore zero even in infinite time. Thus if bubbles are admitted the payoff on the infinitely repeated game is well defined.

As the foregoing examples suggest, admitting bubbles expands economists' ability to analyze settings in which an infinite number of security trades take place.[9]

7. Conclusion

Economics and finance are replete with examples of securities which have values greater than those of their payoffs in the finite future. These securities are sometimes labeled bubbles, but sometimes they are called very long discount bonds, intrinsically useless money, zero-dividend firms, or whatever. Such securities are usually analyzed in a setting which is not completely explicit theoretically, and, as a result, does not provide satisfactory answers to such basic questions as how, under a linear pricing rule, a security which has zero payoff at each date can still have positive value or how, under a nonlinear pricing rule, arbitrage opportunities are avoided. Our approach provides a theoretically explicit answer to these questions, and therefore allows a unified treatment within the general equilibrium paradigm of phenomena that might seem to be unrelated, and to lie outside that paradigm.

[9] Of course, the rules of the game change when the payoff space is **ba** rather than L_1 or L_2. For example, one can no longer routinely ignore sets with probability measure zero.

Appendix

We wish to characterize bubbles as the limit points of nets of fundamental securities. The desired characterization raises the immediate question of whether—more precisely, in what topology—the indicated limits exist. Rather than provide a full answer, consisting of a formal exposition of integration with respect to a charge, we address these issues less formally in the context of an example. The index set is assumed countable.

Let the security space initially be $\ell^1_1 = \ell^1$ (recall that the norm of x in ℓ^1 is $\|x\|^1_1 := \sum_{i=1}^{\infty} |x_i|$). The payoff stream on a unit-norm pure discount bond χ_i that matures at date i is

$$\chi_i(t) = \begin{cases} 1 & \text{if } t = i; \\ 0 & \text{otherwise.} \end{cases} \tag{19}$$

The price space is L^∞, where the norm is $\|p\|^\infty_1 := \sup_i |p_i|$.

The sequence $\{\chi_i\}$ does not converge in the norm topology, nor does it have any limit points, since $\|\chi_i - \chi_j\|^1_1 = 2$ for all $i \neq j$. Given that our intention is to characterize bubbles as limits of fundamentals, we wish to specify a topology in which, unlike in the norm topology, χ_i converges. The example of the zero-dividend firm indicates how to proceed. In that example the value of the dividend streams is constant and positive over different dividend payout rates; this suggests that the dividend streams do not converge to zero, and also that the appropriate test of convergence involves values. Thus it seems natural to turn to the notions of weak and weak* convergence. A net of fundamental securities $\{x_\alpha \mid \alpha \in A\}$ in ℓ^1 converges weakly to $y \in \ell^1$ if the net of real numbers $\{p(x_\alpha)\}$ converges to $p(y)$ for any price system $p \in L^\infty$. The same notion of convergence applied to **ba** instead of ℓ^1 yields weak* convergence in **ba**.[10]

The sequence of discount bonds $\{\chi_i\}$ does not have a weak limit point in ℓ^1. To show this, suppose that $y \in \ell^1$ were such a limit point. By testing against the price system $\chi_t \in L^\infty$, we obtain $y_t = 0$, since $y_t = \chi_t(y)$ must be equal to $\lim_i \chi_t(\chi_i)$; hence $y = 0$. But by testing against $\chi_M \in L^\infty$, we obtain $\chi_M(y) = 1$ since $\chi_M(\chi_i) = 1$ for all i, a contradiction. On the other hand, the sequence is contained in the unit ball of ℓ^1,

[10] More generally, let X be a topological linear space and Y a set of continuous linear functionals on X. If convergence in X is tested by convergence of values for all elements of Y, then the convergence is called weak if Y is the dual of X and weak* if X is the dual of Y.

which is contained in the unit ball of **ba** (the set of charges with total variation norm not exceeding unity), and Alaoglu's theorem states that the unit ball of **ba** is compact in the weak* topology. This theorem therefore implies that the sequence $\{\chi_i\}$ has weak* convergent subnets, although it does not say that the sequence itself converges (false), or that it has weak* converging subsequences (also false).[11] We now construct a convergent subnet; not surprisingly, the limit point is not in ℓ^1: it is a bubble.

An *ultrafilter* F is (in our case) a family of subsets of the natural numbers N such that

(a) the empty set is not in F;

(b) A, $B \in F$ implies $A \cap B \in F$;

(c) $A \in F$ and $A \subset B$ implies $B \in F$; and

(d) for any $A \subset N$, $A \notin F \Leftrightarrow N \setminus A \in F$.

For any ultrafilter F a charge τ_F can be defined by

$$\tau_F(A) := \begin{cases} 1 & \text{if } A \in F; \\ 0 & \text{otherwise.} \end{cases} \tag{20}$$

Note that the family of sets that contain any fixed number $i \in N$ is an ultrafilter and that the bond χ_i—or, more exactly, the measure δ_i in **ba** to which χ_i corresponds as x corresponds to λ in (4)—is a 0–1 charge on this ultrafilter. Ultrafilters of this type are called *fixed ultrafilters* (they fix on the number i).

Zorn's lemma implies the existence of other ultrafilters, called *free ultrafilters*, which do not contain singletons. These free ultrafilters define 0–1 charges in just the same way as fixed ultrafilters. As a consequence of the defining properties of an ultrafilter, any free ultrafilter excludes all the finite sets; this being so, the value of the charge defined in eq. (20) is necessarily zero on any finite set when F is a free ultrafilter. Therefore 0–1 charges defined on free ultrafilters are pure charges rather than measures.

Our result is that if F is any free ultrafilter, there exists a subnet of $\{\chi_i\}$ (or $\{\delta_i\}$) that converges in the weak* topology to τ_F. In economic terms, this means that a subnet

[11] A subnet of a sequence is not necessarily a subsequence, since the directed set indexing the subnet may be richer than the natural numbers; the range of the subnet, of course, is contained in that of the sequence and is thus countable. It is possible for a sequence to have convergent subnets but no convergent subsequence.

of the sequence of discount bonds $\{\delta_i\}$ can be found which converges to a pure bubble, where "converges" means that the values of the discount bonds converge to the value of the bubble under any admissible price system.

To construct the required subnet we need to specify a directed index set and to associate a discount bond with each index. Then, we need to verify that the net so defined is a subnet of the original sequence. We take the index set to be the free ultrafilter F itself, and we direct it by inverse inclusion: $B > A$ if $B \subset A$, so that higher sets have fewer elements. Note that F plays two distinct roles: it determines the charge to which the subnet of discount bonds will be shown to converge and it serves as the directed set indexing the chosen subnet. Now, with each $A \in F$ associate the discount bond $\delta_{k(A)}$, where $k(A)$ is the smallest element in A. If $B > A$, then $B \subset A$ so that $k(B) \geq k(A)$, which means that $\{\delta_{k(A)} \mid A \in F\}$ is a subnet of $\{\delta_i \mid i \in N\}$.

In order to show that $\{\int p\, d\delta_{k(A)}\}$ converges to $\int p\, d\tau_F$ for any price system p, we need to understand integration with respect to a 0–1 charge. For any sequence $p \in L^\infty$ there exists a real number r_p equal to the supremum of the set of real numbers r satisfying: $\{i \mid p_i > r\} \in F$. Then define

$$\int p\, d\tau_F := r_p\,. \tag{21}$$

To understand why this definition makes sense, note that for any $\epsilon > 0$, the set B of indices i such that p_i is within ϵ of r_p (that is, $B := \{i \mid r_p - \epsilon \leq p_i \leq r_p + \epsilon\}$) belongs to the family F, and its complement does not.[12] The pure charge τ_F gives unit weight to elements of F and zero weight to their complement; thus for any ϵ, the set B is given unit weight, while its complement is given no weight, motivating (21).

It remains to show that $\{\int p\, d\delta_{k(A)} \mid A \in F\}$ converges to $r_p = \int p\, d\tau_F$ for any $p \in L^\infty$. This is done by showing that for any $\epsilon > 0$, there exists $B \in F$ such that $C > B$ implies that $\int p\, d\delta_{k(C)}$ is within ϵ of r_p. Fix $\epsilon > 0$ and set $B := \{i \mid r_p - \epsilon \leq p_i \leq r_p + \epsilon\}$, which we know belongs to F. Since $k(B)$ is an element of B, we must have

$$r_p - \epsilon \leq p_{k(B)} \left(= \int p\, d\delta_{k(B)}\right) \leq r_p + \epsilon\,.$$

---

[12] It is immediate that $\{i \mid p_i > r\}$ is not in F for any $r > r_p$, because r_p is defined as the supremum of all numbers with this property. Similarly, $\{i \mid p_i < r_p - \epsilon\}$ cannot be in F, since it is the complement of an element of F. It follows that the set left over must be in F.

Now, elements of F greater than B are subsets of B, so we must have

$$r_p - \epsilon \leq p_{k(C)} \leq r_p + \epsilon$$

for any $C > B$. That establishes the result.

References

Clark, S. A.: The valuation problem in arbitrage price theory, Journal of Mathematical Economics **22**, 463–478 (1993)

Cox, J. C., Ingersoll, J. E., Ross, S. A.: A reexamination of traditional hypotheses about the term structure of interest rates, Journal of Finance **36**, 321–346 (1981)

Duffie, D.: Dynamic asset pricing theory, Princeton University Press: Princeton, New Jersey (1992)

Dybvig, P. H., Ingersoll, J. E., Ross, S. A.: Long forward and zero-coupon rates can never fall, manuscript, Yale University, (forthcoming in Journal of Business, January 1996) (1995)

Feldman, M., Gilles, C.: An expository note on individual risk without aggregate uncertainty, Journal of Economic Theory **35**, 26–32 (1985)

Gilles, C.: Charges as equilibrium prices, and asset bubbles, Journal of Mathematical Economics **18**, 155–167 (1989)

Gilles, C., LeRoy, S. F.: On the arbitrage pricing theory, Economic Theory **1**, 213–229 (1991)

Gilles, C., LeRoy, S. F.: Bubbles and charges, International Economic Review **33**, 323–339 (1992a)

Gilles, C., LeRoy, S. F.: Asset price bubbles, The New Palgrave Dictionary of Money and Finance (1992b)

Hansen, L. P., and Jagannathan, R.: Implications of security market data for models of dynamic economies, Journal of Political Economy **99**, 225–262 (1991)

Kazemi, H. B.: An intertemporal model of asset prices in a markov economy with a limiting stationary distribution, Review of Financial Studies **5**, 85–104 (1992)

Kocherlakota, N.: Bubbles and constraints on debt accumulation, Journal of Economic Theory **57**, 245–256 (1992)

Lucas, R. E., Jr.: Asset prices in an exchange economy, Econometrica **46**, 1426–1446 (1978)

Magill, M., Quinzii, M.: Equilibrium with incomplete markets over an infinite horizon, manuscript, University of Southern California (February) (1993)

Mas-Colell, A.: A model of equilibrium with differentiated commodities, Journal of Mathematical Economics **2**, 263–295 (1975)

Mehra, R., Prescott, E. C.: The equity premium: a puzzle, Journal of Monetary Economics **15**, 145–161 (1985)

Ross, S. A.: The arbitrage theory of asset pricing, Journal of Economic Theory **13**, 341–360 (1976)

Santos, M., Woodford, M.: Rational speculative bubbles, manuscript, University of Chicago (1992)

Tirole, J.: On the possibility of speculation under rational expectations, Econometrica 50, 1163–1181 (1982)

Werner, J.: Diversification and equilibrium in securities markets,reproduced, University of Minnesota (1994)