Infinite Portfolio Strategies[*]

Stephen F. LeRoy
University of California, Santa Barbara
and Federal Reserve Bank of San Francisco

March 14, 2012

Abstract

In continuous-time stochastic calculus a limit in probability is used to extend the definition of the stochastic integral to the case where the integrand is not square-integrable at the endpoint of the time interval under consideration. When the extension is applied to portfolio strategies, absence of arbitrage in finite portfolio strategies is consistent with existence of arbitrage in infinite portfolio strategies. The doubling strategy is the most common example. We argue that this extension may or may not make economic sense, depending on whether or not one thinks that valuation should be continuous. We propose an alternative extension of the definition of the stochastic integral under which valuation is continuous and absence of arbitrage is preserved. The extension involves appending a date and state called ∞ to the payoff index set and altering the definition of convergence under which gains on infinite portfolio strategies are defined as limits of gains on finite portfolio strategies.

In continuous-time finance theory, portfolio gains (prices with past dividends added in) are defined using the Ito definition of the stochastic integral. Security gains p are modeled as a martingale—most simply, as a Brownian motion—and a portfolio strategy is represented by a predictable stochastic process θ. When the integrator is a Brownian motion the definition of the stochastic integral implies that the gain z_τ at time τ of the portfolio strategy, given by

$$z_\tau = \int_{t=0}^{\tau} \theta dp, \tag{1}$$

1

is a random variable with finite mean and variance.

Existence of this integral requires that θ satisfy

$$\int\limits_{t=0}^{\tau} E(\theta_t^2)dt < \infty. \tag{2}$$

In many expositions it is pointed out that the definition of the Ito stochastic integral, and with it the characterization of gains at T of portfolio strategies, can be extended to cases in which the condition (2) is not satisfied for $\tau = T$, but is satisfied for all $\tau < T$. In that case the integral (1) can be characterized as a random variable for each $\tau < T$. In interesting cases these random variables converge in probability as τ approaches T, implying that the portfolio gain at T can be defined as the indicated limit. We will call this limit the *usual definition* of the gain at T when (2) is not satisfied for $\tau = T$.

The problem with the usual definition of the stochastic integral as applied to portfolio payoffs is that even though the expected value of the gain at $\tau < T$ necessarily equals the initial cost of the portfolio strategy, the expected value of the gain at T as just defined does not necessarily equal the initial cost of the portfolio strategy. In this case arbitrage occurs. Neither the mathematics-oriented sources for this extension (for example, Chung and Williams [3]) nor the finance-oriented sources (for example, Duffie [5], Appendix D) provide much in the way of motivation for identifying the extension with a limit in probability. The reader is left uncertain whether the extended definition of the integral involves economic assumptions beyond those involved in the standard case. Can the valuation discontinuity that occurs when condition (2) on θ is not satisfied for $\tau = \infty$ be avoided if the payoffs of infinite portfolio strategies are defined differently?

Discussion of this issue is simplified if we switch to a discrete-time setting with τ an integer and T equal to ∞. The portfolio strategy identified with θ terminated at $\tau < \infty$ is a finite portfolio strategy, the definition of which is unambiguous, so the question is how to take the limit when τ approaches ∞. We will discuss this question in the context of the doubling strategy, which fits the setup just described.

1   The Doubling Strategy

The doubling strategy is usually cast as a gambling strategy in which a gambler bets on the outcome of the toss of a fair coin, doubling the size of the bet until he wins. It can equally well be treated as a portfolio problem in which an investor invests repeatedly in a zero-price security that yields ± 1 with equal probability. The investor doubles the investment at each date until he wins, after which he bets zero. In the truncated version of the doubling strategy the investor terminates the portfolio strategy at date τ whether or not he has won.

The state ω is defined as the first date at which the outcome $+1$ occurs. At date t the observable events consist of the state ω if $\omega \leq t$, and the event $\omega > t$ otherwise. This assumption defines a filtration on $N \times N$. The doubling strategy is defined by

$$\theta_t^\tau(\omega) = \begin{cases} 0, & \omega \leq t \text{ and } t \leq \tau, \text{ or } t > \tau \\ \\ 2^t, & \omega > t \text{ and } t \leq \tau \ . \end{cases} \tag{3}$$

Its gain is

$$z_t^\tau(\omega) = \begin{cases} 1 & \omega \leq \min(t, \tau) \\ \\ -2^{\min(t,\tau)} + 1 & \omega > \min(t, \tau). \end{cases} \tag{4}$$

These stochastic processes are measurable with respect to the filtration just defined.

The fact that ω occurs with probability $2^{-\omega}$ implies that the distribution of z_t^τ, the random variable with realization $z_t^\tau(\omega)$ in state ω, is

$$z_t^\tau = \begin{cases} 1 & \text{with probability } 1 - 2^{-\min(t,\tau)} \\ \\ -2^{\min(t,\tau)} + 1 & \text{with probability } 2^{-\min(t,\tau)}. \end{cases} \tag{5}$$

It follows that $E(z_t^\tau) = 0$. Further, we have $z_t^\tau = E(z_{t+1}^\tau | z_t^\tau)$, implying that z^τ is a martingale.

Under the doubling strategy condition (2) on the portfolio weights, with the integral replaced by a summation, is violated for $t = \tau = \infty$. Therefore the Ito integral is not defined. However, the gain z_t^τ converges in probability to 1 as $t = \tau$ rises, which therefore is the gain on the infinitely repeated doubling strategy under the usual definition. Thus despite the fact that truncated versions of the doubling strategy are increasingly risky, and therefore are increasingly unattractive to any agent who is strictly risk averse, under the usual treatment the infinite version of the doubling strategy is an arbitrage.

It appears natural to take the state ω as the date at which the agent first wins ($\omega = 1, 2, ...$), since then the state probabilities are $1/2, 1/4, ...$, which sum to 1. Doing so in conjunction with identifying the set of states with N, however, is seen to have the unfortunate consequence that by definition the event that the agent never wins cannot occur. In other words, the event that gives the doubling strategy its interest—unlimited possible losses—is defined out of existence in the formal treatment. Obviously the specification that agents literally cannot lose under the doubling strategy plays a central role in generating the conclusion that the doubling strategy is an arbitrage.

It is likely that analysts find it acceptable to rule out the event that the agent never wins because this event occurs with probability zero, and zero-probability events are routinely ignored in all applications of probability. Neglecting zero-probability events is acceptable when, as is usual, the payoff that occurs with zero probability is finite. However, it is not so obvious that infinite (or, as under the doubling strategy,

3

negative infinite) payoffs that occur with zero probability can be ignored without consequence.[1]

Some analysts, to their credit, appear to be uneasy about labeling the infinitely-repeated doubling strategy an arbitrage, particularly in the absence of any characterization of investors' preferences. For example, Delbaen and Schachermayer [4] concluded from the fact that potential losses are unbounded under the doubling strategy that "[e]verybody, especially a casino boss, knows that [the doubling strategy] is a very risky way of winning 1€. This type of strategy has to be ruled out: there should be a lower bound on the player's loss" ([4] p. 130). This passage is interesting on several levels. Most obviously, under the received treatment the doubling strategy, as an arbitrage, in fact has no risk—there exists no state in which the agent can lose. In labeling the doubling strategy as risky, Delbaen and Schachermayer make clear, apparently without realizing it, that they have doubts about the implication of the usual definition of the gain on infinite portfolio strategies that the infinitely-repeated doubling strategy is a riskless arbitrage.

In any case, as a general matter there is no need to rule out very risky portfolio strategies. Assuming that agents are risk averse, they will not adopt very risky portfolio strategies even if such strategies are available unless they are well compensated in terms of expected return. Just the opposite: the (apparent) necessity for ruling out the doubling strategy arises precisely from the representation of the doubling strategy as not being risky.

Despite Delbaen and Schachermayer's ambivalence, a case can be made that the treatment of the doubling strategy is correct, at least as applied to some situations. Under the doubling strategy, after all, there is no finite date at which the investor must realize his loss. Therefore it can be argued that there is in fact no possibility of loss since he can always continue playing. The situation is similar to a Ponzi scheme, in which an investor borrows money and then rolls over the indebtedness forever. Because there is no date at which the borrower must repay the loan, it is effectively never repaid.

However, there is another way to look at the matter. Infinite portfolio strategies

---

[1]The assumption that justifies the usual treatment corresponds in mathematical terms to the presumption that $0 \times \infty$ can be taken to be equal to 0. Mathematicians have not adopted a uniform treatment of this question. Wolfram MathWorld and Wikipedia, "Extended Real Number Line" state that $0 \times \infty$ is usually left undefined, like $0/0$. Royden [11] listed the assumption $0 \times \infty = 0$ separately from the others, identifying it as an "arbitrary *convention*" (emphasis in original), suggesting that he was not entirely comfortable with the assumption.

In probability and measure theory, in contrast, $0 \times \infty$ is sometimes defined as equaling 0. For example, let $f(x)$ equal 1 on $[0, 1]$ and 0 on $(1, \infty)$. Then if λ denotes Lebesgue measure, the Lebesgue integral $\int_0^\infty f(x)d\lambda$ equals $\lambda([0, 1]) \times 1 + \lambda((1, \infty)) \times 0$. Since $\lambda((1, \infty)) = \infty$, this leads to the standard result of 1 only if $0 \times \infty$ is defined to equal 0.

Thanks to Hrishikesh Singhania for conversations on this material.

are often rationalized as approximations to finite portfolio strategies. For example, the Black-Scholes model of option pricing requires that agents who wish to hedge the risk of holding an option will trade at each of an uncountably infinite number of dates. Actually conducting such a set of trades, of course, is impossible. The model is appropriately viewed as an approximation to a model in which transactions costs are low, but not zero, so that agents can conduct a large finite number of trades, but not an infinite number. Often, as with the Black-Scholes model, it is easier to model settings where agents can trade an infinite number of times than where they can trade a large finite number of times. Now, if infinite portfolio strategies are to be viewed as approximations of finite portfolio strategies, it is important that there be no discontinuities in going from the finite to the infinite case. Under the conventional representation of the doubling strategy, as we have seen, the gain is increasingly unattractive in the finite case, but is an arbitrage in the infinite case if convergence is defined as convergence in probability. This is clearly a discontinuity.

This raises the question whether there are alternatives to convergence in probability in characterizing the payoffs of infinite portfolio strategies as limits of finite portfolio strategies. The desired alternative would have the property that utility and value can be defined to be continuous in τ, so that there are no jumps in utility and value going from finite τ to ∞. We will propose such an alternative definition of convergence below. Because valuation is continuous under that definition, absence of arbitrage on finite portfolio strategies implies absence of arbitrage on infinite portfolio strategies. Thus admitting infinite portfolio strategies does not necessitate imposing trading restrictions to exclude the doubling strategy or, for that matter, other arbitrages such as Ponzi schemes. Thus if gains of infinite portfolio strategies are defined so that they are approximations to gains of finite portfolio strategies, considerable simplification in the analysis become available.

The point is that it is inappropriate to impose a dictum either that the infinitely repeated doubling strategy is necessarily an arbitrage, or that it cannot be properly viewed as an arbitrage. Instead, whether or not such portfolio strategies are arbitrages should be seen as an implication of a modeling decision that is at the discretion of the analyst. The appropriate choice depends on the context and the desired interpretation of the model. Thus the analyst who wishes to treat portfolio strategies like the doubling strategy as exemplifying continuous valuation can achieve this by adopting our modified definition of the portfolio gain in place of the usual definition; the analyst who prefers to view the infinitely repeated doubling strategy as an arbitrage and is willing to impose trading restrictions to rule such trading strategies out, of course, is welcome to do so.

2   Gains at Infinity as Weak Limits

The alternative definition of the gain on an infinite portfolio strategy as the limit of the gains on finite portfolio strategies obviously requires that the sequence of gains

on truncated versions of the infinite portfolio strategy converge. Under most of the familiar characterizations of convergence, such as convergence under a sup norm or a mean-square norm, gains on such portfolio strategies as the finitely repeated doubling strategy do not converge. The way around this is to define payoffs at ∞ and adopt a convergence definition such that gains on finite portfolio strategies for which condition (2) fails at ∞ converge to a payoff that includes a nonzero component at ∞. Suppose, then, that we add ∞ to the set of dates.[2] Further, in defining the set of states we add to the natural numbers (which represent the event of winning at the finite dates) the state $\omega = \infty$, representing the event of never winning. Thus the state space is now $N_\infty \equiv N \cup \{\infty\}$, and sample paths are defined on $N_\infty \times N_\infty$. We can extend the definition of the filtration to $N_\infty \times N_\infty$ by assuming that at date ∞ agents know the state for all states, including ∞. It is assumed that the gain at the state ∞ of any finite portfolio strategy is 0.

Portfolio gains as defined above are random variables on N_∞. Let z be the gain on a portfolio strategy that is not a finite portfolio strategy (because it cannot be generated by a finite number of trades of the primitive securities). The infinitely-repeated doubling strategy is the obvious example. Consider the portfolio strategy that coincides with z up to date τ, but at τ the agent terminates the strategy by closing out all positions. For any finite τ this a finite portfolio strategy, so its gain is as defined above. We will define the gain on the infinite portfolio strategy z as the weak limit as τ goes to ∞ of the gains on the z_τ so defined.

To do this we convert the mathematical representation of portfolio gains from functions to measures. Fix a finite portfolio strategy with gain z. The fact that z_t has value 0 at ∞ implies that for any z_t there exists a measure μ_t, continuous with respect to probability (because $z_\infty = 0$ and $\phi_\infty = 0$), such that z_t is the Radon-Nikodym derivative of μ_t with respect to ϕ :

$$ z_t = \frac{d\mu_t}{d\phi}. \tag{6} $$

Thus we can characterize the gain z_t at t on any finite portfolio strategy either as a function z_t on N_∞ or as a measure μ_t on the partition of N_∞ at t implied by the information filtration described above.

Next we characterize the continuous functionals on N_∞, since these are the test functions required for weak convergence. Define a metric on $N_\infty \times N_\infty$ by

---

[2] Adding ∞ to the sets of states and dates means that we cannot view our analysis of the doubling strategy as embedded within a sequential equilibrium (that is, an equilibrium in which agents have a separate budget constraint at each event, with that constraint involving only consumption and security transactions at that event), since payoffs at ∞ would not enter such constraints. Rather, we are assuming an Arrow-Debreu equilibrium concept, where each agent trades subject to a single budget constraint that is integrated over time. In that case the relevant stochastic processes can involve components at the date and state ∞.

The distinction between sequential and Arrow-Debreu equilibrium is discussed further below.

$$d(\omega, \omega') = |2^{-\omega} - 2^{-\omega'}|, \qquad \omega, \omega' \in N_\infty, \tag{7}$$

which corresponds to Euclidean distance when $1, 2, 3, ...\omega, ..., \infty$ are identified with $1/2, 3/4, 7/8, ..., 1 - 2^{-\omega}, ..., 1$. In the topology induced by this metric, ∞ is a limit point of N_∞, and it is the only limit point, so we have that adding ∞ to N to obtain N_∞ constitutes the Alexandroff one-point compactification of N. Under the definition of convergence implied by the metric (7) the continuous functions c on N_∞ are those for which $\lim_{\omega \to \infty} c(\omega) = c(\infty)$.[3] The purpose of this specification is to implement the idea that the high elements of N are increasingly close to ∞. This assumption will make possible the result that gains at t converge to gains at ∞.

A sequence of truncated portfolio strategies $\{\mu_\tau\}$ weakly converges to μ_∞ if $\sum_\omega c d\mu_\tau$ converges to $\sum_\omega c d\mu_\infty$ for every continuous function c. We will define μ_∞ as the gain on the assumed infinite portfolio strategy. We will assume that the utility of the limiting gain is the limit of the gains of the associated truncated portfolio strategies, and similarly for valuation. Thus utility and valuation are continuous.

These definitions are readily applied to the doubling strategy. Interpreted as a measure, the gain at τ of the doubling strategy truncated at τ is

$$\mu_\tau = \sum_{\omega=1}^{\tau} 2^{-\omega} \delta_\omega + \sum_{\omega=t+1}^{\infty} (2^{-(\omega+1)} - 1) \delta_\omega, \tag{8}$$

where A is any subset of N_∞ and δ_ω is the point mass associated with ω:

$$\delta_\omega(A) = \begin{cases} 1 \text{ if } \omega \in A \\ \\ 0 \text{ if } \omega \notin A \end{cases} \tag{9}$$

for $A \subset N_\infty$. The sequence of measures $\{\mu_\tau\}$ weakly converges to μ_∞, defined by

$$\mu_\infty = \sum_{\omega=1}^{\infty} 2^{-\omega} \delta_\omega - \delta_\infty, \tag{10}$$

which is therefore the gain on the infinitely repeated doubling strategy. This occurs because, under the metric (7), the point-masses δ_ω are very close to δ_∞ for high values of ω.

The assumption that valuation is continuous is seen to imply that the value of the infinitely-repeated doubling strategy equals the values of the finitely-repeated portfolio strategy, which are zero. The utility of the consumption streams generated by infinite portfolio strategies can be defined similarly as the limits of the utilities generated by the corresponding truncated portfolio strategies. The utility of a consumption stream consisting of the infinitely-repeated doubling strategy is zero for a

---

[3]A basis for the topology that has these continuous functions is the singleton sets of N together with the unions of ∞ and the tails of N_∞. Note that in this topology the singleton set ∞ is not open.

risk-neutral agent, so that the agent is indifferent as to whether to initiate either the finite or the infinite version of the doubling strategy. An agent with quadratic utility would assign negative infinite utility to the infinitely repeated doubling strategy. Thus in contrast to the usual treatment, which holds that the infinitely repeated doubling strategy is an extremely good outcome, under the revised treatment it is seen to be an extremely bad outcome.

Note that gains that have a nonzero component at ∞ cannot be converted from measures to functions on N_∞. This is so because such gains are represented by measures that are discontinuous with respect to probability, implying that their Radon-Nikodym derivatives with respect to probability are not defined. Observe that here there is no violation of the condition that zero probability events do not determine values. The probability of the state ∞ is 0, but this means only that the gains on infinite portfolio strategies, as measures, may not be continuous with respect to the probability measure. This is the case with the infinitely-repeated doubling strategy.

Using weak convergence is seen to enlarge the set of well-defined infinite portfolio strategies to include portfolio strategies for which the condition $\int_{t=0}^{T} E(\theta_t^2)dt < \infty$ is not satisfied, as with the doubling strategy. The values of gains on truncated infinite portfolio strategies, including the doubling strategy, is the same as the initial costs of the infinitely repeated doubling strategy, so there is no arbitrage in passing from the truncated to the infinite doubling strategy. [4]

3 Bubbles

The *fundamental value* of the payoff of a portfolio strategy is the summed value of its finite-date dividends. The bubble is the component of the gain that occurs at ∞, and the *bubble value* is the value of this term, equal to the difference between the initial cost of the portfolio strategy and the fundamental value. This definition indicates that portfolio strategies that have bubbles are exactly those with gains that are nonzero at ∞. It is seen that this paper is about how to model bubbles, although we have not used that term up to now. Of course, if one is willing to rule out bubbles

---

[4]Of course, even with the revised definition of the gain on infinite portfolio strategies, there still exist candidate portfolio strategies for which the gain is undefined.

The Alaoglu theorem guarantees existence of a limit point for any net of portfolio strategies that is bounded. Unbounded portfolio strategies may not have limit points, implying that the gains on the infinite versions of such strategies may be undefined. In this case we must reject the association of θ with an infinite portfolio strategy.

As the doubling strategy indicates, there are many portfolio processes θ that are associated with bona fide portfolio strategies under the definition involving weak limits even though they are not so associated under the usual extension.

at the outset, as is frequently done in financial analysis, then the issues discussed in this paper do not arise.

The analysis of bubbles goes in different directions depending on whether one is assuming Arrow-Debreu markets, as here, or sequential markets. What we are calling the Arrow-Debreu specification is basically an extension of the original static general equilibrium models of Arrow and Debreu to settings in which there are an infinite number of states, dates, and in which agents can conduct an infinite number of trades. Under the Arrow-Debreu specification, as in the classic finite general equilibrium model, it is assumed that each agent has a single budget constraint that applies to transactions at all states and dates. One can assume that all trading occurs before time begins; as in the finite case this treatment involves no loss of generality because trades are state- and date-contingent. Bewley [2], Gilles [6], Gilles and LeRoy [7], [8] and others have conducted analyses of bubbles in Arrow-Debreu settings. Santos and Woodford [12] and especially Huang and Werner [9] have shown that the analysis of bubbles in sequential markets has little in common with the Arrow-Debreu case.

Bubbles in Arrow-Debreu settings can be analyzed directly using general equilibrium methods. This is done in the papers just cited. This involves directly defining utility functions on infinitely nonzero payoff streams. Here, in contrast, we have instead assumed that utility was defined initially on payoffs of finitely nonzero portfolios, following which utility and value were extended to the limit points that represent payoffs of infinite portfolios. The relation between these two methods is far from clear. It can be verified that they are consistent in the simplest examples, but no results of a general nature are available. This is a worthwhile subject for future research.

4   Conclusion

The simplest—and least controversial—contribution of this paper is to raise questions about the easy characterization of the infinitely repeated doubling strategy as an arbitrage that has to be ruled out by imposing trading restrictions. On the contrary, that characterization emerges as a consequence of modeling decisions that are at the discretion of the analyst. We demonstrated this by presenting an alternative modeling procedure under which the infinitely repeated doubling strategy is not an arbitrage. At least for some purposes the alternative procedure is attractive, inasmuch as it avoids a discontinuity at infinity that does not have any obvious counterpart in actual human behavior.

From an analytical point of view this conclusion is good news. Under our treatment, if agents are strictly risk averse (and have continuous preferences) our analysis implies that they will avoid the doubling strategy even if it is available (and priced fairly). Whether or not agents are risk averse, there is no need to impose portfolio restrictions to exclude portfolio strategies like the doubling strategy. Analytically this conclusion is convenient: without trading restrictions choice sets are linear spaces

and valuation is linear; with trading restrictions matters are more complicated, even in settings that would otherwise be easy to analyze (see, for example, LeRoy and Werner [10] for analysis of asset pricing in the presence of trading restrictions; see Huang and Werner [9] for conditions under which trading restrictions rule out bubbles in sequential equilibrium).

There exists some empirical evidence—relying on data prior to the recent financial crisis, to be sure—to the effect that bubbles do not occur in real-world financial markets (Abel *et al.*, [1]). As usual, this evidence is not conclusive, and in any case it only tests for the existence of the simplest sort of bubbles. We take the view that any strong conclusion on questions involving bubbles is premature. Many phenomena occur in financial markets that are difficult to reconcile with the simplest rational-agent model: one thinks of asset price volatility, the equity premium puzzle, and the periodic occurrence of liquidity crises. These phenomena do not appear to produce obvious profit opportunities that agents are irrationally ignoring. Accordingly, it should be possible to analyze them using the orthodox methods of financial economics. If some or all of these phenomena turn out to be connected to bubbles, the analytical techniques developed here may be useful.

References

[1] Andrew B. Abel, N. Gregory Mankiw, Lawrence H. Summers, and Richard. J. Zeckhauser. Assessing dynamic efficiency: Theory and evidence. *Review of Economic Studies*, 56:1–20, 1989.

[2] Truman F. Bewley. Existence of equilibria in economies with infinitely many commodities. *Journal of Economic Theory*, 4:514–540, 1972.

[3] K. L. Chung and R. J. Williams. *Introduction to Stochastic Calculus*. Birkhauser, Boston, 1990.

[4] Freddy Delbaen and Walter Schachermayer. *The Mathematics of Arbitrage*. Springer, 2006.

[5] Darrell Duffie. *Dynamic Asset Pricing Theory, Second Edition*. Princeton University Press, Princeton, N. J., 1996.

[6] Christian Gilles. Charges as equilibrium prices, and asset bubbles. *Journal of Mathematical Economics*, 18:155–167, 1989.

[7] Christian Gilles and Stephen F. LeRoy. Bubbles and charges. *International Economic Review*, 33:323–339, 1992.

[8] Christian Gilles and Stephen F. LeRoy. Rational bubbles as payoffs at infinity. *Economic Theory*, 9:261–281, 1997.

[9] Kevin X. D. Huang and Jan Werner. Asset price bubbles in Arrow-Debreu and sequential equilibrium. *Economic Theory*, 15:253–278, 2000.

[10] Stephen F. LeRoy and Jan Werner. *Principles of Financial Economics*. Cambridge University Press, Cambridge, 2001.

[11] H. L. Royden. *Real Analysis: Second Edition*. Macmillan, New York, 1968.

[12] Manuel Santos and Michael Woodford. Rational asset pricing bubbles. *Econometrica*, 65:19–57, 1997.