An Unambiguous Statement of Interest Rate Parity

Kuntara Pukthuanthong*

University of Missouri

Lee R. Thomas III

Independent

Abstract: Interest rate parity (IRP) states that the difference between the interest rates of two currencies equals the expected percentage change in the associated exchange rate. Unfortunately, the IRP is ambiguous. There are two, inconsistent values of the expected percentage change in the exchange rate, depending on which currency is treated as the numeraire. We derive a corrected IRP condition that is unambiguous but subtly different. The equilibrium interest rate differential equals the percentage difference between the spot exchange rate and the geometric mean of the future exchange rate distribution, not its expected value.

Keywords: Interest rate parity; Siegel's paradox; growth investing; dynamic risk neutrality; universal hedging

*Please contact Kuntara Pukthuanthong at pukthuanthongk@missouri.edu and Lee R. Thomas III at lee.r.thomas.3@gmail.com

Uncovered interest rate parity (IRP) is ubiquitous in international finance models. According to the IRP condition, if investors are risk neutral, then the difference between the interest rates on two currencies equals the expected percentage change in the exchange rate.

Suppose that the two currencies in question are called dollars and pounds. Let X_t be the dollar per pound exchange rate at time t, and suppose X_o is known, but X_1 is a random variable. Further, let x_t= X_t^{-1} be the contemporaneous pound per dollar exchange rate and let $i_\$$ and $i_£$ represent the continuously compounded dollar and pound interest rates, respectively. Then, based on continuous compounding, the conventional statement of the IRP relationship (derived subsequently) is as follows.

$$i_\$ - i_£ = \log E\,[X1/X0], \tag{1a}$$

or

$$i_£ - i_\$ = \log E\,[x1/x0]. \tag{1b}$$

Unfortunately, (1a) and (1b), both representing IRP, are inconsistent; they cannot hold simultaneously. Specifically, as Siegel (1972) points out, because of Jensen's inequality, $EX_1 \neq (\mathrm{E}x_1)^{-1}$, so the expected future exchange rate, and the expected percentage change in the exchange rate, depend on which currency is treated as the numeraire. The implication is that there is no equilibrium interest rate differential. Involving the covered interest rate parity condition, there is no equilibrium forward foreign exchange rate when investors are risk neutral.

We will demonstrate that the corrected statement of IRP is as follows.

$$i_{\$} - i_{£} = \log G\,[X1/X0] \tag{2a}$$

or

$$i_{£} - i_{\$} = \log G\,[x1/x0] \tag{2b}$$

where G denotes the geometric mean operator. These equations are consistent, since

$$G(X1) = [G(x1)]^{-1}.$$

To derive this corrected statement for the IRP, we rely on the following assumptions:

(i) In each period, t=1,...T, an investor can buy or sell foreign currency assets or choose instead to invest in a riskless domestic asset. The proceeds produced by one period's round of investing are reinvested in the next period so that the investor realizes compound returns on their wealth.

(ii) The rate of return produced by holding foreign currency in each T period depends on the foreign interest rate and the capital gain or loss associated with random and independently distributed changes in the exchange rate.

(iii) We assume that the investor's objective is to maximize the expected growth of wealth after T periods. Investors do not demand a risk premium to take risky speculative positions. Specifically, they are concerned only with their expected compound growth rate of wealth, not its riskiness. Such an investor is sometimes referred to as “growth optimal.”

Then, (2a) and (2b) follow.

In Section 2, we reproduce 'Siegel's paradox' (1972) and demonstrate the resulting problem to derive the IRP relationship. Section 3 describes a growth optimal model of currency speculation and applies it to deriving an unambiguous statement of IRP. In Section 4, we discuss assumption (iii), comparing the growth optimal model to a model that assumes in a multi-period context that investors maximize their expected terminal wealth. Section 5 explicitly considers what it means to say that an optimal foreign exchange speculator for growth is 'risk neutral', defining a novel concept that we call geometric risk neutrality. Section 6 briefly considers some implications of the analysis for hedging foreign exchange risk exposures. Section 7 summarizes the conclusions.

2. Siegel's Paradox

To understand the problems deriving the international interest rate the parity condition, consider the following exchange rate puzzle. If investors expect the dollar per pound exchange rate to be 2 next year, what would they expect the future pound per dollar rate to be? Even though intuitively, 0.5 appears to be the obvious answer, the mathematical expectation of the pound per dollar cannot be 0.5 if the expected dollar per pound rate is 2 because of Jensen's inequality. Accordingly, when we calculate the expected future exchange rate, we must first stipulate which currency we use as the numeraire, pounds, or dollars.

This perplexing conundrum manifests Siegel's paradox (Siegel, 1972, Siegel, 1975, McCulloch, 1975, Roper, 1975). Siegal's paradox is intriguing and has curious implications. For example, Siegel's paradox implies that exchange rate volatility can be a source of expected speculative profits from holding foreign currency (Dumas and Solnik, 1995), and that those expected speculative profits accrue to both sides of a gamble on foreign exchange simultaneously. Accordingly, foreign and domestic investors may each want to hold foreign exchange-rate-exposed assets, even if both investors share the same expectations concerning prospective exchange rate changes.

There is another curious implication of Siegel’s paradox. Siegel's paradox appears to make conventional uncovered interest rate parity (IRP) impossible. According to IRP, in the absence of a risk premium, the international interest rate difference equals the percentage difference between the current spot rate and the expected future rate. Siegel's paradox observes that two inconsistent values represent the expected future exchange rate. The international interest rate differential cannot meet both expectations. There is no equilibrium value of the interest rate differential under risk neutrality. Factoring in covered interest rate parity, this means that there is no equilibrium forward exchange rate.

Because interest rate parity is so essential to international finance, many researchers have been researching this problem and have tried to explain, question, or utilize Siegel's paradox. Although some authors hold that the paradox is valid and significant for investment selection (e.g., Black, 1989; Sinn, 1989; Black, 1990; Dumas and Solnik, 1995), others argue that Siegel's paradox is an avoidable illusion (Levich and Thomas, 1993) or generates too tiny of an effect to be significant (McCulloch, 1975). Other models try to explain the paradox by introducing new assumptions, such as investor risk aversion or currency diversity preference (Roper, 1975). These and other attempts are inadequate because none explicitly offers an intuitively satisfying explanation of the paradox or a corrected version of IRP. We attempt to fill in this gap.

For simplicity, we refer to two currencies: dollars and pounds. These currencies could be interpreted instead as distinct consumption baskets; one could be “apples” and the other “oranges”, as in Black (1989), so the exchange rate is the relative price of apples in terms of oranges. Then there is no need to introduce currencies or distinguish between domestic and foreign currencies. Alternatively, we could assume that there are two consumption good baskets, and one can only be purchased with dollars and the other only with pounds, perhaps because the baskets contain non-tradable goods.

Closely following Siegel (1972), we now formally demonstrate the paradox. Suppose that investors are risk neutral and share the same information relevant to forecasting the exchange rate that will prevail at t=1. At t = 0, investors with initial wealth W_0 pursue either investment strategies to convert their initial wealth endowment into their preferred currency at t = 1. They can hold their wealth in riskless domestic currency-denominated assets, or they can hold their wealth in default-free, foreign currency-denominated assets and trade foreign currency for domestic currency at the end of the period, whatever exchange rate then prevails.

The risk-free interest rate on dollars, expressed as an exponential growth rate, is $i_{\$}$, and the risk-free rate on pounds is $i_{£}$. Speculators can borrow or lend at those interest rates. Following Siegel (1972), suppose that a risk-neutral investor maximizes expected terminal wealth at the end of one period, EW_1. From a dollar investor's perspective, the terminal wealth associated with investing a dollar in a riskless dollar asset is:

$$e^{i_{\$}} \text{ dt.}$$

The uncertain terminal wealth from investing in pounds and then converting pounds to dollars at the ending exchange rate, is

$$(X_1/\, X_0)e^{i_{£}} \text{ dt.}$$

The end wealth from investing in pounds, measured in dollars, is uncertain since X_1 is a random variable.

The investor is risk neutral. Therefore, investment in domestic and foreign deposits must have the same expected ending payoff. Equating expected ending wealth for the two investing strategies, equilibrium requires that for dollar-based and pound-based investors, respectively, the following conditions hold:

$$e^{i_\$} = e^{i_£}\ \mathrm{E}(X_1/\ X_0) \tag{3}$$

$$e^{i_£} = e^{i_\$}\ \mathrm{E}\ (x_1/\ x_0) \tag{4}$$

Equations (3) and (4) hold simultaneously iff $EX_1 = (Ex_1)^{-1}$. This is not true, as it violates Jensen's inequality. Compared with investors who take their profits in dollars and investors who take their profits in pounds, the expected future exchange rate, and therefore the expected percentage change in the exchange rate, are inconsistent. As a result, $i_\$ - i_£$ equals either log[E($X_1/\ X_0$)] or -log[E($x_1/\ x_0$)]. Uncovered interest rate parity, therefore, is ambiguous.

3. Growth Optimality and IRP

Foreign currency speculators can make what is essentially the same bet repeatedly. Moreover, they can roll the proceeds of one period’s investment into the next period’s investment and earn a compound return. We assume that an investor with starting wealth of W_o takes speculative foreign currency positions or instead invests in the domestic asset, T times in a row, reinvesting the intermediate proceeds each period, where T>1.

Subsequently, we will expand upon the concept of risk neutrality in a multi-period context, explicitly defining what we call geometric risk neutrality. For the moment, we assume that an investor is indifferent between a strategy that produces a specific growth rate of wealth and one that produces an uncertain growth rate of wealth with the same expected value. Expressed differently, the investor we model is concerned only with the expected growth rate of their wealth and not, for example, with the growth rate's variability or the variance of terminal wealth.

An investor who maximizes his expected growth of wealth while taking serial investments is called growth optimal. The growth optimal investment model, first proposed in a gambling context in a paper dealing with information theory by Kelly (1956), advanced in a financial investment context by Latane (1959) and developed further by Breiman (1986) and Hakansson (1971), is a framework used to select investments in a multiperiod setting that incorporates the effects of the compounding of return. Optimal growth investments have several desirable properties. For example, it not only maximizes the expected long-run growth rate of an investor's wealth but also minimizes the expected time required to achieve an arbitrary wealth target. In the future, the optimal growth portfolio will inevitably outperform any other portfolio. We discuss the growth optimal model and compare it to an alternative model based on the assumption that investors maximize expected terminal wealth in Section 4.[1]

The representative investor's objective is to maximize E log(W_T/W_o), where wealth evolves as a multiplicative random process

$$W_T = W_o \prod_{t=1}^{T} R_t, \tag{5}$$

and $R_t = W_t/W_{t-1}$ is a random variable representing investment return and on capital during period t.

From (4), the expected growth rate of wealth is

$$\text{E Log}[\text{W(t)/W(0)}] = \sum \text{E}\,[\log \text{R(t)}], \tag{6}$$

and the investor's objective is to maximize the expectation of (6). Notably, decision-making is myopic in the growth optimal model of investment selection. That is, although the expected growth rate of wealth of the investor depends on R(t) at every time point t=1,2,...,T, at each decision point, the investor is only concerned with the information available at that time, specifically the value of E [log R(t)] for all available investments.

[1] Vander Weide (2010) offers an excellent comparison of the growth optimal model and two other popular models of investor behavior, the Markowitz mean-variance model and the lifetime consumption-investment model. Estrada (2010) describes the debate concerning growth optimality between Markowitz (1976) and Samuelson (1971).

From a dollar investor's perspective, holding a dollar asset produces a logarithmic return of

$$\log R_t = i_{\$}\, dt,$$

and holding a pound asset produces a logarithmic return of

$$R_t = \log\,[e^{i_{£}}\,(X_t/\,X_{t-1})] = i_{£}\ \mathrm{dt} + \log\,[X_t/\,X_{t-1}].$$

In equilibrium, the expected log returns to dollars and pounds must be equal. Accordingly, in equilibrium.

$$i_{\$}\text{-}i_{£} = \mathrm{E}\,\log\,[X_t/\,X_{t-1}].$$

Using the fact that for an arbitrary random variable Z,

$$\mathrm{E}\,[\log \mathrm{Z}]=\log \mathrm{G}[\mathrm{Z}],$$

where G is the geometric mean operator, this can be rewritten

$$i_{\$}\text{-}i_{£} = \log(\mathrm{G}[X_t/\,X_{t-1}]),$$

which is given by Eq. (2a). Repeating this from a pound-based investor's perspective produces Equation (2b):

$$i_{£} - i_{\$} = \log(\mathrm{G}[x_t/x_{t-1}]).$$

Importantly, unlike the expectation operator the geometric mean operator has the property that it inverts.[2] That is, GX_1 = (G$x_1)^{-1}$. Accordingly, (2a) and (2b) are consistent. The equilibrium interest rate differential is the percentage difference between the current spot exchange rate and the geometric mean of the future exchange rate distribution.

4. Modeling Considerations

Interest rate parity relies on the assumption that speculators are risk neutral. This assumption means there will be no risk premium embedded in the return they demand holding foreign currency-denominated assets. Consequently, the observed interest rate differential reflects only the anticipated change in the exchange rate. Conventional risk-neutral investors maximize expected ending wealth, so their implied utility function is U(W) = W. However, we assume investors are growth optimal. The implied utility function of a growth optimal speculator, one who maximizes expected log wealth, is U(W) = log W. The logarithmic utility function represents absolute risk aversion, not risk neutrality. This appears to conflict with the assumption that speculators are risk neutral.

The speculators we model are risk-neutral within the context of our model. Examining Equations 2a and 2b, we observe that no term represents a risk premium. However, the assumption of conventional risk neutrality must be replaced by a new definition, which we call geometric risk neutrality.

[2] The geometric mean is the only measure of the first moment of a distribution that has this property, says the math books I looked at. The paper you sent me proved that the geometric mean is the only "aggregator" that has this property – which apparently is well known by mathematicians. I assume an "aggregator" means a summary of the data, i.e.., something that produces a first moment. That means that no operator other than the geometric mean is going to give a unique equilibrium. (Any other "aggregator" will run into the same problem that the expectation operator does, i.e. a Jensen's inequality type problem.) So our statement using the geometric mean operator is the only statement of IRP that works.

Conventionally, a risk-neutral investor is indifferent to a statistically fair bet, that is, with the expectation of his stochastic cash flows equal in magnitude to its cost. Thusly, upon accepting a fair bet, the bettor's expected change in wealth is zero. A growth optimal investor's objective is to maximize the expected growth rate of wealth, not expected future wealth. For a growth optimal investor, a just fair bet leaves the expected growth of wealth unchanged. We call this a geometrically fair bet. It is easy to demonstrate that a bet cannot be statistically and geometrically fair. If an investor accepts a statistically fair bet, then his or her wealth is expected to decline.

An investor is conventionally risk-neutral when indifferent to a statistically fair bet. We call investors geometrically risk neutral when they are just indifferent to a geometrically fair chance. That is, a growth-optimal investor is indifferent to an investment opportunity that expects to produce a compound growth of wealth of zero.

A growth optimal investor accepts any investment for which $E \log (W_T / W_0) > 0$, rejects any investment for which $E \log (W_T / W_0) < 0$, and is indifferent when $E \log (W_T / W_0) = 0$. That means that growth optimal investors are geometrically risk-neutral. That is why there is no risk-dependent terms in equations (2a) and (2b).

Turning to a second modeling question, we consider speculators' objectives in our model. We assume speculators maximize the expected growth rate of their wealth.

In a single-period model, an investor maximizes end-of-period wealth, W_1. A natural extension of this to a multi-period model assumes that speculators maximize expected terminal wealth, W_T. However, when wealth evolves as a multiplicative random process, maximization of expected terminal wealth can lead to poor investment decisions.

When returns compound, an investor's wealth evolves as a multiplicative random process. Accordingly, the distribution of ending wealth is positively skewed, driving a wedge between the distribution's mean and median. Under these circumstances, the expected value of ending wealth misleads as a measure of central tendency. It does not represent the typical values of ending wealth observed in a small sample. It makes little sense for an investor to maximize expected ending wealth if the expected value of ending wealth is not a good predictor of what the investor expects to observe in the long run if he or she invests serially. Moreover, as we will demonstrate using an example famous for economic thought development, maximizing expected terminal wealth implies that investors make dubious investment decisions.

An extreme example of this phenomenon was posed as a question in decision theory about 300 years ago, designed to test the limit of the expected value's usefulness in representing anticipated terminal wealth under certain circumstances. The problem posed by Nicholas Bernoulli, important in the development of economic thinking, is called the St. Petersburg paradox.

Toss a coin. If it comes up a tail, you are paid $1.00, and the game ends. However, if you toss a head, then toss again. If you toss a tail on the second toss, your payoff is multiplied by two; you are paid $2.00, and the game ends. If the first tail appears on the third toss, you receive $4, and so on. The question is, what would you pay for a ticket to play the Saint Petersburg lottery?

Observe that the player's wealth evolves as a multiplicative random process. Therefore, it is unsurprising that the payoffs to the St. Petersburg lottery ticket have a right-skewed distribution (Figure 1).

In this histogram, we generate the results with the following assumptions.

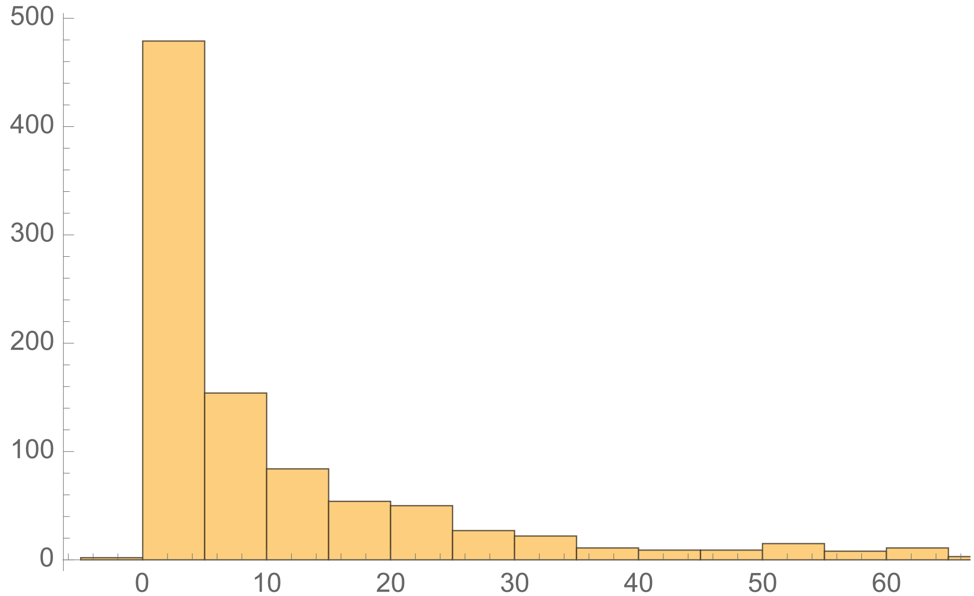

This figure shows that after 30 years, the mean and median continuously compounded returns are 9.5% and 5.3%, respectively. This histogram shows that the distribution of returns becomes more and more non-normal as time passes and returns compound. Mean-variance requires returns to be normal unless the utility function is quadratic. Thus, it makes little sense for investors to have as their objective maximizing expected ending wealth if the expected value of ending wealth is not meaningful. Moreover, we will show examples that maximizing expected terminal wealth would lead investors to make dubious investment choices.

As a result, we anticipate the expected value to be a value that few players will enjoy. No investors can be paid as much as the expected value. The player receives a finite payoff in all cases, but the expected payoff is infinite. Strictly speaking, an investor should be willing to pay any finite sum for a St. Petersburg lottery ticket. What if the prizes were only half as much in every state of the world? An investor is still willing to spend the same for a ticket, precisely, any finite amount.

This absurd conclusion results from assuming an investor maximizes expected terminal wealth. That is inappropriate when wealth dynamics are multiplicative.

Consider the value of a ticket to a growth optimal investor. Since growth optimal investors are geometrically risk neutral, in equilibrium the St. Petersburg lottery must be a geometrically fair game. That is, the ticket must be priced so that the investor's expected growth rate of wealth is zero. If N is the number of tosses it takes to observe a tail, then the payoffs are 2^{N-1} with probability 0.5^N. Accordingly, in equilibrium for a growth optimal investor, the return on his investment is the following:

$$\frac{\sum(0.5^N)\log(2^{N-1})}{W_0} = 0$$

Solving, the equilibrium ticket price, W_0, is $2.00.

** and ** (****) performed an experiment to estimate the equilibrium ticket price for the St. Petersburg lottery, with the experimental subjects incentivized by an actual monetary payment. The mean of their sample was a ticket price of $1.75, albeit with a large sample variance. Accordingly, we cannot reject the hypothesis that the subjects of the experiment priced the ticket as if they were growth optimal.

5. Risk Neutrality

The derivation of IRP relies on the fact that a growth optimal investor is concerned only with the expected growth rate of wealth and does not demand a risk premium to compensate if the growth rate is uncertain rather than known. However, since the investor maximizes log [W(T)/W(0)] the growth optimal model of investor behavior implicitly assumes that the utility of wealth is given by U(W)=log W. Conventionally, a logarithmic utility function is considered to represent risk aversion, seemingly contradicting our assumption of risk neutrality. We now expand the concept of risk neutrality within the context of the optimal growth model of investor behavior.

A 'fair bet' is conventionally defined as one with expected net cash flows of zero. Defining a fair bet in this way is useful in a one-period model but fails when an investor may take speculative bets multiple times, earning compound returns. In general, an actuarily fair bet in one period is unfair when repeated with compounding. To demonstrate, suppose that an investor takes a conventionally fair bet, flipping a coin for “double or nothing.” In other words, suppose that an investor has the opportunity to pay $1, receiving $2 if they toss a head or $0 if they toss a tail, using a fair coin. This clearly is what is conventionally called a fair bet. The expected value of the risky cash flow is $1, equal to the $1 required to purchase the risky payment. Conventionally, by definition, a risk-neutral investor is indifferent to the choice of an actuarially fair but risky bet or a certain thing with the same expected value.

Now, assume that the risk-neutral investor is confronted by this gambling opportunity T times in a row. Should repetition be a factor in the investor's decision whether to take a chance or play it safe? If the investor chooses to hold on to their initial $1 round by round, without betting, their ending wealth is $1. However, if the investor chooses to gamble each time, double or nothing on every round of betting, they eventually will lose all their wealth because in the limit, as the number of coin tosses approaches infinity, the probability of tossing a tail approaches 1. In fact, the investor would have no wealth left after only 5 tosses with a probability of 97% and has only a 0.1% probability of avoiding ruin after ten tosses.

It is difficult to see why a multi-period investor would be indifferent to taking sure of something or accepting a serial double-or-nothing gamble. Even though the game is conventionally fair in one period, in the long run, the specific outcomes are $1 if they refuse to gamble or $0 if they gamble serially.

To generalize, let P_o be the cost of purchasing a random cash flow P_i, i ~{1,2,…,N}, with each state having the associated probability p_i. Further assume that P_i is identically and independently distributed, and that a bet is characterized by EP_i=P_o, so that the bet is actuarially fair. Subsequently, log E(P_i/P_o)=0. An investor’s wealth, W_t, evolves so that

$$W_t = W_o \ \pi[(P_i)_t / P_o],$$

where $(P_i)_n$ is the realized value of P_i in period t. Accordingly,

$$\text{E} \log (W_t/W_o) = \text{T E} \log (P_i/P_o).$$

Since

$$\log \text{E} (P_i/P_o) > \text{E} \log (P_i/P_o),$$

$$\text{E} \log (W_t/W_o) < 0.$$

An actuarially fair bet is expected to produce a negative growth rate i wealth if the bet is repeated serially.

In this study, we redefine a fair bet in the context of a gamble made repeatedly, with the reinvestment of the proceeds of prior bets at every round of betting. For what we call a "geometrically fair bet,"

$$\text{E} \log (W_T/W_o)=0.$$

Since

$$\text{E}[\log (W_t/W_o)] = \log \text{G}(W_t/W_o),$$

this implies that

$$\text{G}(W_t/W_o) = 1.$$

That is, for a "geometrically fair bet"

$$\text{G}P_i = P_o.$$

This compares to

$$\text{E}P_i = P_o$$

for “an actuarially fair bet.

Suppose that an investor is offered a dynamic fair bet. If they are just indifferent to accepting it we characterize the investor as being dynamically risk neutral. A dynamically risk-neutral investor accepts a bet when G $P_i > P_o$, rejects it when G $P_i < P_o$ and is just indifferent when G $P_i = P_o$. Importantly, dynamic risk neutrality is consistent with the Bernoulli utility function of wealth, U(W)=log W, since a dynamically fair bet leaves log W unchanged. The currency speculators we model are geometrically risk-neutral. Implicitly, they behave as if they have logarithmic utility functions.

6. “Universal Hedging”

One practical implication of our solution of Siegel's paradox pertains to exchange rate hedging. Black (1990) proposes the following formula for international investors' optimal hedge ratios when they have currency exposures associated with their investment portfolios.

He asserts that this relationship is obtained for all investors, regardless of which consumption basket they consume (equivalently, regardless of their base currency.) The hedge ratio is

$$1-\lambda=\frac{\mu_m-\sigma_m^2}{\mu_m-\frac{1}{2}\sigma_e^2} \tag{7}$$

where μ_m and σ_m are the mean return and volatility of the market portfolio, σ_e^2 is the volatility of the exchange rate, and λ is an aggregate measure of risk aversion. The most obvious feature of interest in (7) is that, in general, not all foreign currency exposure embedded in an international equity portfolio should be hedged away. In fact, implicitly it suggests that investors who do not have any foreign currency exposure in their investment portfolio should buy foreign currency to secure such exposure. This result is a consequence of the fact that in equilibrium foreign currency exposures can have positive expected return, viewed from the perspective of either the domestic or the foreign investor, because of Siegel's paradox. The effect is non-trivial: Black (1990) estimates that 30% to 70% of foreign currency exposures should not be hedged.

In Black (1990) the expected return associated with holding foreign currency is a consequence of the expected return produced by Siegel's paradox, a fact that Black acknowledges (1989, p. 905.) Absent Siegel's paradox, the optimal hedge ratio is 100% for all investors since foreign currency exposures are risky but do not bear compensating expected return.

7. Conclusions

Conventional IRP asserts that the international interest rate differential equals the expected percentage in the exchange rate if investors are risk neutral. However, Siegel's paradox implies that this intuitive relationship cannot be true for all investors because Jensen's inequality means that the expected future exchange rate is not unique. We show that if investors are dynamically risk-neutral, they always accept (reject) a bet that increases (decreases) the expected growth rate of their wealth and take serial speculative foreign exchange positions, reinvesting proceeds each period, then an exact IRP condition holds. Measuring interest rates on a continuously compounded basis, the international interest rate differential equals the logarithmic difference between the spot exchange rate and the geometric mean of the distribution of the future exchange rate.

References

Bessembinder, Hendrik. "Do Stocks Outperform Treasury Bills?." *Journal of financial economics* 129, no. 3 (2018): 440-457.

Bessembinder, Hendrik. "Extending Portfolio Theory To Compound Returns." *Available at SSRN 3875870* (2021).

Black, Fischer, "Universal hedging: optimizing currency risk and reward in international equity portfolios," Financial Analysts Journal, 1989, 45 (4), 16-22.

Black, Fischer, “Equilibrium exchange rate hedging," The Journal of Finance, 1990, 45 (3), 899-907.

Breiman, L. (1961), “Optimal Gambling Systems for Favorable Games”, Fourth Berkeley Symposium on Probability and Statistics I, 66–78

Dumas, Bernard and Bruno Solnik, “The world price of foreign exchange risk," The Journal of Finance, 1995, 50 (2), 445-479.

Estrada, Javier, “Geometric mean maximization: an overlooked portfolio approach?" Journal of Investing, 2010, 19 (4), 134-147.

Grauer, Robert R, “A comparison of growth optimal and mean variance investment policies," Journal of Financial and Quantitative Analysis, 1981, 16 (1), 1-21.

Hakansson, Nils H, “Capital growth and the mean-variance approach to portfolio selection," Journal of Financial and Quantitative Analysis, 1971, 6 (1), 517-557.

Kelly, J., "A New Interpretation of information rate," Bell System Technical Journal (Aug. 1956), 917-926.

Latané, Henry Allen, “Criteria for choice among risky ventures," Journal of Political Economy, 1959, 67 (2), 144-155.

Latané, Henry Allen, 1959. “Rational decision-making in portfolio management,” Journal of Finance, American Finance Association, vol. 14(3), pages 429-430, September.

Levich, M. and R. R. Lee. Thomas, “The significance of technical trading-rule profits in the foreign exchange market: a bootstrap approach," Journal of International Money and Finance, 1993, 12 (5), 451-474.

Markowitz, Harry M, “Investment for the long run: new evidence for an old rule," Journal of Finance, 1976, 31 (5), 1273-1286.

Markowitz, Harry M. 1952, Portfolio Selection, The Journal of Finance 7, 77–91.

Markowitz, Harry M., 1959, Portfolio selection: efficient diversification of investments (Yale University Press, New Haven).

Mossin, Jan, “Optimal multiperiod portfolio policies," The Journal of Business, 1968, 41 (2), 215-229.

Nalebuff, Barry. “The other person's envelope is always greener,” Journal of Economic Perspectives, 3 (1), winter 1989, 171-181.

Roper, Don E, “The role of expected value analysis for speculative decisions in the forward currency market," The Quarterly Journal of Economics, 1975, pp. 157-169.

Samuelson, P. (1963). “Risk and Uncertainty: A Fallacy of Large Numbers.” Scientia, 1-6.

Samuelson, Paul (1971). “The ‘Fallacy’ of Maximizing the Geometric Mean in Long Sequences of Investing or Gambling.” Proceedings of the National Academy of Sciences, 68, 2493-2496.

Siegel, Jeremy, “Reply: Risk, interest rates, and the forward exchange," The Quarterly Journal of Economics, 1975, 89 (1), 173-175.

Siegel, Jeremy J, “Risk, interest rates and the forward exchange," The Quarterly Journal of Economics, 1972, pp. 303-309.

Sinn, Hans-Werner, “Expected utility and the Siegel paradox," Journal of Economics, 1989, pp. 257-268.

Vander Weide, James H., “Principles for lifetime portfolio selection: Lessons from portfolio theory," in John Guerard, ed., Handbook of Portfolio Construction, Springer, 2010, pp. 153-177.