The Valuation of Nonsystematic Risks and the Pricing of Swedish Lottery Bonds

Richard C. Green
Carnegie Mellon University

Kristian Rydqvist
Norwegian School of Management

Swedish government lottery bonds have coupon payments determined by lottery. They offer a unique opportunity to study a security with uncertain payoffs having a known, observable distribution. The risk associated with the lotteries is idiosyncratic by construction and should not command a risk premium in equilibrium. The bonds are traded in two forms, allowing us to evaluate the rewards to bearing extra lottery risk. Despite its idiosyncratic nature, we find prices appear to reflect aversion to this risk. We evaluate the empirical determinants of this differential pricing and possible explanations for it.

One of the most basic insights offered by modern financial theory concerns the distinction between

The authors wish to thank Erik Ringqvist and Malin Hedman for help in obtaining data, along with Svante Johansson of Carnegie Fondkommission, Sture Lundén of Swedbank, Per Nygren, Ann Sjöberg, Håkan Tiderman of the Swedish Treasury, Niels Madsen of Stora Finance, Berry Sandström of Svenska Penniglotteriet, and Sune Rydqvist for information about the lottery bond market. Seminar participants at Arizona State, British Columbia, Carnegie Mellon, Columbia, Duke, Indiana, Göteborg, Northwestern, Texas at Dallas, Tulane, Washington, Washington at St. Louis, the 1995 Winter Finance Conference, the 1994 CEPR Summer Institute in Financial Markets, the 1995 WFA, and the Swedish Treasury provided numerous helpful comments and suggestions, as have Ralph Bernstein, Mark Levy, Dan Quan, Neil Pearson, Eduardo Schwartz, Neal Stoughton, Suresh Sundaresan, Kirk Vandezande, two anonymous referees and David Hirshleifer (the editor).
Address correspondence and requests for reprints to Richard C. Green, GSIA, Carnegie Mellon University, Pittsburgh, PA 15213.

The Review of Financial Studies Summer 1997 Vol. 10, No. 2, pp. 447–480

systematic and nonsystematic risk. Systematic risks are those that drive returns on well-diversified portfolios. To the extent that investors succeed in diversifying their positions, theory predicts risks that are independent of the sources of systematic risk should be uncompensated.

This idea is central to research methods, policy recommendations, and pedagogy in finance. Though often exposited through an appeal to the mean-variance capital asset pricing model (CAPM), the distinction between systematic and nonsystematic risks is an important characteristic of other asset pricing models, both static and dynamic. If those risks can be avoided simply through diversification, at equilibrium they will be independent of investors' marginal rates of substitution, and thus unpriced.

Unfortunately, observed economic outcomes rarely afford an opportunity to empirically evaluate these principles at a level anywhere close to their full generality. Determining which risks are nonsystematic requires an asset pricing model that restricts investors' marginal rates of substitution in terms of observables. If it appears that nonsystematic risks are priced, this can always be attributed to misspecification of the model. Since the probability distributions of the asset returns are unknown, additional structure must often be imposed to facilitate estimation, adding another potential source of misspecification.

The market for Swedish lottery bonds provides a unique laboratory to evaluate the degree to which these basic principles are reflected in actual market prices. On these bonds, which are issued by the Swedish Treasury, the total coupon payment on any given date is fixed, but the allocation of the payment across bonds within the issue is determined by lottery. In other situations where asset payoffs are risky, the distribution cannot be observed. It must be estimated. Here the probability distribution of the payments is, in fact, observable. Further, the risk associated with the payments is by construction independent of other random variables in the economy. It is purely idiosyncratic. A hypothetical investor purchasing an equally weighted portfolio of all the bonds in the issue would have a deterministic payoff, so the risks are by construction diversifiable.

Evaluating whether these risks are priced would be difficult if what were required were an otherwise identical riskless Swedish government bond with deterministic payoffs. There is not a sufficiently rich array of government bonds in the Swedish market to permit reliable estimation of the term structure. Moreover, coupon payments on the lottery bonds are tax exempt, further complicating comparison with the available riskless bonds, which are taxable.

Lottery bonds, however, are traded simultaneously as individual securities, and in packages that include enough sequentially numbered bonds to guarantee a portion of the periodic coupon payment. Thus we have the same security trading in one case as a "diversified" portfolio and in the other case as a mean-preserving spread of the diversified portfolio, where the risk embodied in the spread is idiosyncratic by construction. A positive reward for bearing the extra risk would certainly appear to be anomalous relative to traditional asset pricing models, which would predict all investors hold portfolios that are diversified. We argue further that it would be anomalous even relative to the diversification opportunities available in the Swedish lottery bond market. The marginal distributions created by holding the portfolio of sequenced bonds differ very little from those created by holding a similar portfolio of bonds out of sequence, because the effects of diversification are quantitatively much more important than the partial guarantee.

The prices at which the bonds trade in sequence and out of sequence thus provide a unique opportunity to evaluate whether investors in an organized secondary market understand the impact of diversification on their portfolios and price assets accordingly. Moreover, by investigating the cross-sectional and time-series determinants of the relative prices, we can evaluate whether such factors as tax-based trading and relative liquidity influence the degree to which prices deviate from the theoretical predictions.

We examine a dataset of daily transaction prices for 39 different lottery bonds over the period from November 1986–February 1994. We find that the individual bonds are on average underpriced relative to the portfolios that ensure a minimal riskless payoff with the same expected return. This premium is difficult to justify as a rational response to the risks faced by investors given available diversification opportunities. We examine several possible explanations for the mispricing, including barriers to trade by outside investors, the relative liquidity of the bonds trading in different forms, the failure of investors in the market to understand the relative risks of the bonds from a portfolio theoretic perspective, and the role of the bonds in tax-motivated transactions. The price differential is empirically related to the effects of the partial guarantee, and thus to the differential risk investors (irrationally) ascribe to the bonds in the two forms. Liquidity factors also appear to have a role in explaining some features of the price differential.

A number of theoretical arguments appear in the asset pricing literature that predict idiosyncratic risk would be priced. Typically, however, these arguments appeal to some market friction to exclude a subset of investors from trading an asset. As a result, the first-order conditions that price the asset will hold with equality only for that subset of investors trading in it. The risk premium on the asset will reflect its covariance with the aggregate consumption of the market participants, but this, as it is only part of the payoff on aggregate wealth, includes idiosyncratic risk. Hirshleifer (1988), for example, uses fixed costs of participation in futures markets to obtain risk premia for futures contracts that deviate from the CAPM's predictions depending on the direction producers are hedging. This result has received empirical support in Bessembinder (1992). Such arguments cannot explain nonzero prices in the market for lottery bonds for a very simple reason. Aggregate holdings of lottery bonds, and hence the aggregate consumption of those investors participating in this market, carry no lottery risk. The Swedish Treasury's liability for the bonds at each coupon date is deterministic.

We conclude that the price differential is anomalous. Investors either misunderstand the nature of nonsystematic risk or have nonstandard preferences that value the minimum payout per se. Because the bonds are tax exempt, this market is dominated by individuals rather than institutions, and the marginal investor may be unsophisticated. Thus, while the lottery bond market provides a relatively clean laboratory in which to consider the valuation of nonsystematic risk, and while there is nothing in asset pricing theory that excludes retail investors in Sweden, still, our conclusions may not generalize to markets where competitive forces have freer rein.

In the next section we describe the lottery bonds and the markets where they are traded. Section 2 clarifies the claim that the bonds should command the same price irrespective of the form in which they are traded. Section 3 presents empirical results on the price differences between the bonds with partial guarantees and those without. In Section 4 we summarize possible explanations for the anomalous evidence and present auxiliary empirical results that bear on these explanations. Section 5 concludes and briefly considers the policy implications that follow from the results.

1. Swedish Lottery Bonds

Lottery bonds are obligations of the Swedish Treasury. The coupon payment, which like municipals in the United States are tax exempt, is determined with a lottery. We first describe how the lotteries are structured. We then provide some background on the markets in which the bonds are traded and discuss the data employed in our analysis. Finally, we summarize the tax treatment of the bonds.

1.1 The structure of the coupon lottery

Over the period we studied the Swedish Treasury floated one to three issues of lottery bonds per year. The bonds are, by convention, denoted by the year of issuance. For example, the second issue of 1991 is

Table 1
Payout distribution for 1987:1

No. of payments	Amount	Prob. of receipt	Expectation	Variance
2	1,000,000 SEK	.000000167	.17	167,000
24	100,000 SEK	.000002	.20	20,000
240	10,000 SEK	.00002	.20	2,000
1,920	5,000 SEK	.00016	.80	4,000
4,900	2,000 SEK	.00041	.82	1,633
12,200	1,000 SEK	.00102	1.02	1,011
120,000	100 SEK	.01	1.00	92
—	0 SEK	0.988	.00	17
Sum equals mean and variance			4.20	195,753

The probability a payment is received in column 3 is computed as the number of such payments (column 1) divided by the total number of bonds outstanding in the issue (12 million). The contribution to the expectation in column 4 is this probability times the payment in column 2. The contribution to the variance in column 5 is computed as the squared difference between the sum of column 4 times the probability in column 3.

denoted 91:2. Within each issue, bonds are identified by two numbers, a series number and an order number. The lotteries through which the coupons are distributed are random drawings of these numbers. Table 1 details the distribution of payments associated with a single coupon payment date of the 87:1 issue. This issue has 12,000 series with 1,000 order numbers each, or 12 million bonds, with a face value of 200 Swedish kronor (SEK). At each of three payment dates per year, 50.4 million SEK are paid out; 6.3% of face value per year.

The total payment on each date is deterministic. By holding a portfolio of all the bonds in the issue, the lottery risk could be "diversified away" and thus should be unpriced in equilibrium. The prices of the lottery bonds should be the same as otherwise identical Swedish Treasuries without the lottery feature. Making such a comparison, however, is very difficult for several reasons. First, the bonds are tax exempt. Second, they are bearer bonds, and have value for hiding wealth and money laundering. Finally, the maturity structure of the Swedish Treasury market is not sufficiently rich to allow one to reliably estimate a term structure or identify close matches to a given lottery bond.

Fortunately, the lottery has a two-tiered structure. Larger payments are based on draws, without replacement, from the entire pool of bonds in the issue. Smaller payments are distributed through drawings only of order numbers, one from each sequence of 50, 100, or 1,000 bonds, depending on the issue in question. To use the 87:1 issue as an example, a number between 1 and 100 is drawn, another between 101 and 200, and so on to a number between 901 and 1,000. Then a payment of 100 SEK is assigned to the holder of any bond with an order number that has been drawn. To any investor holding 100 bonds in a series with consecutive order numbers covering one of these intervals, this smaller payment is therefore guaranteed.

The bonds are issued to original subscribers in sequenced blocks. As they begin to trade in the secondary market, some sequences of order numbers remain intact, while others are broken up. The bonds then continue to trade actively in both forms, and prices for both types are reported daily in the Swedish financial press. Following the conventions in the Swedish market, we refer to the prices of the bonds out of sequence as "mixed," as opposed to bonds that are traded "in sequence." The recorded prices for the mixed bonds are not necessarily prices for small lots or single bonds. They may represent very substantial trades, but the bonds involved in the trade will not all come from the same series and have sequential order numbers.

1.2 Characteristics of the Swedish lottery bond market

Lottery bonds have been in use in Sweden since 1918 and now constitute some 8% of the Swedish government debt.[1] Treasury officials argue the lotteries are attractive to retail investors and increase the dispersion of ownership. Prior to the advent of modern data processing, the lotteries also offered substantial cost savings to the issuer due to the reduction in the total number of payments that had to be distributed on a given coupon date.

The bonds are issued by the Swedish Treasury through direct sales at a fixed price or through exchange offers to refinance older issues. Lottery bonds are traded on the Stockholm Stock Exchange, and the Swedish financial press reports on a daily basis transaction prices and bid-ask quotes for the bonds trading in mixed form and for various sizes of blocks of bonds with sequential order numbers. These prices are reported for sequences of 50, 100, and, for some issues with low face value, 500 and 1,000 bonds. Trading in the bonds is suspended for a period of 6 days prior to and 6 days after each lottery to facilitate distribution of the lottery.[2] Our data consists of end-of-day bid prices for the 39 different lottery bond issues outstanding over the period November 1986–February 1994. These are the highest bids in the open book of the (computerized) Stockholm Stock Exchange. Thus, unlike U.S. bond price data that comes from phone surveys, these bid prices represent a commitment to trade.[3]

---

[1] Our source for this and other descriptive information about the bonds is *Riksgäldskontorets Faktabok om Premieobligationer 1991*, a publication of the Swedish Treasury.

[2] Since many of the bonds are held in accounts at Swedish banks (according to Swedish Treasury officials, about one-third), this allows the bank time to register the series and order numbers of any bonds they hold with the treasury authorities before the lottery, and to collect the lottery payments before ownership changes.

[3] Nevertheless, it seems likely that the bid-ask spread would be narrower for the larger blocks, and since we examine only one side of the spread, a part of any price difference between the mixed and sequenced bonds could be due to the spread. Casual checks using transactions prices

From 1977–1980, and from 1982–1985 there were three lottery bond issues per year,[4] and these bonds made two coupon payments per year. Since 1986 there have been two issues per year, except for 1990 when there was only one. These bonds all have three lotteries per year. All the bond issues since 1990 have been variable rate, tied to a short-term Swedish Government rate, with caps and floors.[5] All the bonds in our dataset except one have some portion of the coupon payment guaranteed to holders of sequences of 50 (2 issues), 100 (21 issues), or 1,000 (15 issues) bonds.[6] Through the 1970s and again in recent years, lottery bonds have been issued with 8 to 10 years to maturity. From 1981–1985, though, the Swedish Treasury issued bonds with 5- and 6-year maturities. Bonds issued between 1984 and 1987 also had some blocks of bonds that could not be "unbundled" and sold as individual or mixed bonds. This feature may explain the unusual price behavior for these bonds.

Table 2 reports daily volume statistics for the bonds in our dataset. This may be important in evaluating the relative liquidity of the bonds in their several forms. The largest amount of trade, in terms of face value, takes place in 1,000-bond blocks, but such transactions are relatively infrequent. A more appropriate measure of liquidity may well be in the lower rows of the table, which report the number of days with no reported transactions. Here we find that the most active markets are those for the mixed bonds and the 100-bond blocks.

1.3 The tax environment

The proceeds from the lotteries are exempt from income tax.[7] Capital gains and losses are handled in the same way as with other bonds, except that, because the coupon is random, interest is not accrued. Price

from the Swedish financial press strongly suggest a narrower spread is not the source of price differences between mixed bonds and sequenced bonds. For example, on February 1, 1994, bid prices for the 90:1 issue were 935 (mixed) and 1,100 (100-sequence), while transactions prices were 950 (mixed) and 1,110 (100-sequence). For the 91:1 issue, bid prices were 935 and 1,030 and transactions prices were 940 and 1,040 for mixed and sequenced bonds, respectively.

[4] In 1981 there was one such issue. We do not, however, have price data for this bond, which had matured prior to the start of our sample period.

[5] Five days prior to each coupon date, the total interest payment for the coupon date following the upcoming one is determined as 65% of the 6-month Swedish Treasury bill rate. The result is then rounded to fit a discrete grid of, for example, 6.0%, 6.3%, 6.5%, 6.7%, 7.0%.

[6] For the currently outstanding bonds, only the smallest of the lottery payments is subject to the partial guarantee. Bonds issued from 1977–1978 and 1981–1986 had partial guarantees that applied to the two smallest payment levels, while those issued in 1979 and 1980 had partial guarantees for the three smallest payments.

[7] Lottery proceeds are subject to a special "lottery tax" at a flat rate of 20%, but do not contribute to ordinary income. Thus, apart from the requisite accounting entries on the part of the treasury, this tax simply serves to reduce the lottery prizes by 20%. Consistent with prevailing convention, we quote coupon rates net of this special tax.

Table 2
Daily volume statistics

	All trades	Mixed	50	100	500	1,000
Average volume (face)	2,619	319	122	492	25	1,661
Trading days	25,197	25,197	23,379	23,562	4,209	16,562
Days with No Trade (%)	—	15.8	19.2	7.0	52.5	33.9

Daily volume is measured in face value traded per bond and reported in thousands of SEK. The bonds trade in unsequenced form (mixed), or in sequenced blocks, where the numbers of bonds in each block are the column headings.

appreciation prior to the lottery is treated as a capital gain and any drop in price associated with the distribution of the lottery proceeds is a capital loss.[8] Formerly the lottery bonds were heavily employed in tax-arbitrage strategies. By purchasing bonds prior to a coupon at a price that reflected its expected value, and then selling after the coupon is paid, an investor earned a tax-exempt coupon, while the associated capital loss could be used to offset gains from other securities. Prior to 1980, such losses could be used to offset any income. In 1980 this was changed. On bonds already issued, capital losses could offset only gains on equities and lottery bonds. For subsequently issued bonds, only gains from lottery bonds could be offset. Since 1991, investors can use 70% of losses from lottery bonds to offset any investment income up to a limit of 100,000 SEK.

To be effective these positions must be highly leveraged. Banks reportedly required 50- and 100-bond sequences for collateral in these short-term arbitrage positions, because of the lower handling costs when the winning numbers had to be verified by hand, although mixed bonds were allowed as collateral for other types of trades. Modern data processing has reduced the difference in handling costs, and banks now accept mixed bonds as collateral. According to practitioners, however, all such tax-arbitrage trading, and the forward markets in the bonds that facilitated it, ended with the final lottery on the pre-1980 bonds in 1990. Apparently this is due to a combination of lower tax rates, the 100,000 SEK cap on capital loss offsets, and a surplus of losses from the stock and real estate market collapses in Sweden.

Our sample includes some bonds issued prior to 1981, which could be used for tax arbitrage, bonds issued after 1980, which could not, as well as data after 1990, when losses on all bonds could partially offset other income, allowing us to empirically evaluate the importance of tax trading in the relative pricing of the mixed and sequenced bonds.

[8] Capital gains on fixed income securities in Sweden prior to 1991 were taxed at 100% of the ordinary marginal rate if the asset had been held for less than 2 years, 75% if held for 2 to 3 years, 50% if held for 3 to 4 years, 25% if held for 4 to 5 years, and 0% if held for more than 5 years. Since 1991, all capital gains, both short and long term, have been taxed at a flat rate of 30%.

2. Why Bonds in Mixed and Sequenced Form Should Command the Same Price

From the perspective of formal asset pricing models, different prices for mixed and sequenced bonds would be anomalous. Since the total coupon is deterministic, lottery risk does not alter aggregate outcomes and would be unpriced in these models. It is also curious from a more practical point of view. In this section we will discuss these arguments in detail before moving on to the empirical evidence.

2.1 Portfolio separation and nonsystematic risks

Formal asset pricing models achieve lower-dimensional pricing by restricting the nature of risks that move agents' marginal rates of substitution. For example, in the Lucas (1978) ICAPM or in Breeden's (1979) continuous-time ICAPM, the consumption of every agent is perfectly correlated with aggregate consumption. Therefore any risk that is statistically independent, or mean independent, of the aggregate is uncorrelated with agents' marginal rates of substitution and will command a risk premium of zero. In the arbitrage pricing model of Ross (1976) or Connor (1984), risks that are uncorrelated with the factors driving the return generating process are not borne in equilibrium and are unpriced. In the mean-variance CAPM, and international versions of it, any risk that is uncorrelated with the "market portfolio" has an expected excess return of zero.

An equally weighted portfolio of mixed bonds has payments that have exactly the same distribution as those associated with a sequenced portfolio, with one exception. The lowest payment level for the mixed bonds is random, while that earned on the sequenced portfolio has the same mean but is deterministic. While the *absolute* prices of each type of bond reflect systematic risks associated with currency and interest rate movements, the *relative* prices of the two types can only reflect their differences, and they differ only in terms of lottery risk. In fact, the bonds are remarkable in that they offer a perfect example of the conditions enumerated by Ross (1978) as sufficient to prove portfolio separation and a zero price for residual risk.

The following assumptions allow us to make the formal arguments clear. They also serve as a useful way of organizing the empirical evidence.

Assumptions 1–5:

1. Trade in lottery bonds takes place in frictionless markets with no transactions or handling costs.
2. Investors are risk-averse expected utility maximizers.
3. There are no taxes.

4. Investors have rational expectations. They understand the distinction between systematic and unsystematic risk and, based on this understanding, trade to their preferred position.
5. Lottery bond markets are perfectly liquid.

Our arguments follow the classic portfolio separation reasoning of Ross (1978), applied to a dynamic setting. We first show that if the lottery risk is unpriced, then the mixed and sequenced bonds, which have the same expected payoffs, can be ordered by second-order stochastic dominance. We then show that, because the lottery risk is independent of other sources of economic uncertainty, any consumption allocation that includes lottery risk can be second-order stochastically dominated by another, with the same cost, that is free of lottery risk. Finally, since the lottery risk is mean independent of each agent's consumption allocation, it will be uncorrelated with their marginal rates of substitution and thus unpriced.

Note that this argument implies the systematic risks of the mixed and sequenced bonds are identical. The payoffs on the mixed bonds can be written as equal in distribution to the payoffs on the sequenced portfolio plus a zero-mean error that is mean independent of each agent's consumption allocation. Therefore the covariances between payoffs and marginal rates of substitution are identical for the mixed and sequenced bonds. Note also that we do not require that investors be perfectly diversified, or trade friction free, in other markets. Because the lottery payoffs are independent of other variables driving consumption allocations, our arguments are stronger than appeals to classical asset pricing results that would require all investors to hold "the world wealth portfolio."

In the following propositions we consider a lottery bond issue with N series numbers, M order numbers, and requiring K order numbers in sequence to qualify for the partial guarantee. Proofs for this and the subsequent propositions are in the Appendix.

Proposition 1. *Suppose all bonds in an issue are priced the same, whether mixed or sequenced. Consider the following positions in a lottery bond issue, all with equal face value of unity:*

(i) An equally weighted portfolio of K mixed bonds, which has J_i bonds with order number i, $i = 1, \ldots, M$.

(ii) An equally weighted sequenced portfolio of K bonds with sequential order numbers.

(iii) An equally weighted portfolio of all NM bonds in the issue.

These positions have the same expected payoff at each date, and they can be ordered by second-order stochastic dominance. Consumption *allocations, which are the same except that they hold the bonds in the above forms, can be similarly ordered.*

The next proposition follows from the fact that there is no aggregate lottery risk and the preference all agents have for second-order stochastic dominance.

Proposition 2. *If equilibria exist, then identical prices for mixed and sequenced bonds is an equilibrium outcome.*

Without further assumptions we cannot, in a dynamic setting, rule out equilibria in which lottery uncertainty is priced. Indeed, we know from the "sunspot" literature that rational expectations equilibria can respond to purely extrinsic uncertainty [see, e.g., Cass and Shell (1983)]. Under either of the following assumptions, however, mixed and sequenced bonds will be priced identically in *any* equilibrium.[9]

Assumptions 6 and 7:

6. Equilibrium allocations are Pareto optimal.
7. In all equilibria, at least one agent has a dynamic consumption allocation with no lottery risk.

Proposition 3. *Under Assumptions 1–5 and either 6 or 7, mixed and sequenced bonds must sell at the same equilibrium price.*

2.2 Are mixed and sequenced bonds close substitutes?

From a more practical perspective mixed and sequenced bonds might be viewed as close substitutes simply in terms of the diversification opportunities immediately available to a small investor. The payoffs on the individual bonds are independent, aside from the fact that the lotteries are drawn without replacement. Therefore, diversifying across mixed bonds dramatically reduces variance and has a similar effect on the higher moments of the distribution. These effects would appear to be much more significant than the impact of the partial guarantee. A diversified portfolio of mixed bonds will have a marginal payoff distribution that is very similar to that of a sequenced portfolio, even though it fails to carry the same partial guarantee. Accordingly, a large price difference between mixed and sequenced bonds would appear anomalous, not just relative to the esoteric criteria of asset pricing theory, but also relative to the more fundamental notion that close substitutes should be priced similarly. More sophisticated practitioners in the lottery bond market appear to be well aware of these

[9] In a one-period setting, these extra assumptions would not be needed.

arguments. Akelius (1980) advocates the use of mixed bonds in certain tax arbitrage strategies for precisely this reason.

To get an idea of the relative magnitudes involved, we ignore the issue of replacement and assume the bonds have *i.i.d.* payoffs. We decompose the payoff per unit of face value on bond h, Y_h, as follows:

$$Y_h = X_h + Z_h, \tag{1}$$

where X_h is the payoff on the lottery over the subset of K bonds, of which bond h is a member, which carries the partial guarantee, and Z_h is the payoff on the lotteries common to all the bonds in the issue. Since X_h and Z_h are independent,

$$\mathrm{var}(Y_h) = \mathrm{var}(X_h) + \mathrm{var}(Z_h). \tag{2}$$

A portfolio of K mixed bonds with the same face value would have variance

$$\mathrm{var}\left(\sum_{h=1}^{K} \frac{1}{K} Y_h\right) < \frac{1}{K}\mathrm{var}(X_h) + \frac{1}{K}\mathrm{var}(Z_h), \tag{3}$$

where the inequality results from sampling without replacement, which induces small amounts of negative correlation, increasing the power of diversification. If on the other hand the portfolio is composed of bonds with sequential order numbers, then the variability in the payoff received from the lottery unique to the order numbers will be eliminated and the portfolio variance will be bounded by $\frac{1}{K}\mathrm{var}(Z_h)$. The difference is just $\frac{1}{K}\mathrm{var}(X_h)$, which is extremely small relative to the total variance of the portfolio. For the bonds in our sample, $\mathrm{var}(X_h)$ is generally less than 0.5% of the total variance. The standard effects of portfolio diversification swamp the risk reduction due to the partial guarantee. For example, for the 87:1 issue the partial guarantee applies to the smallest lottery payment of 100 SEK, or 150% of face value when annualized. This is a large portion of the expected return, but as Table 1 shows, it represents only $92/195,637 = 0.05\%$ of the variance. In contrast, holding the 100 bonds needed for a complete sequence in an equally weighted portfolio reduces variance by a factor of 100, even if the bonds have mixed order numbers.

3. The Price Differential

Figure 1 plots the average of the price of the sequenced bonds less the price of the mixed bonds as a fraction of face value for each bond issue in our sample. Also depicted are two standard errors adjusted for first-order autocorrelation above and below the mean. The second panel provides the analogous information for the yield differentials.

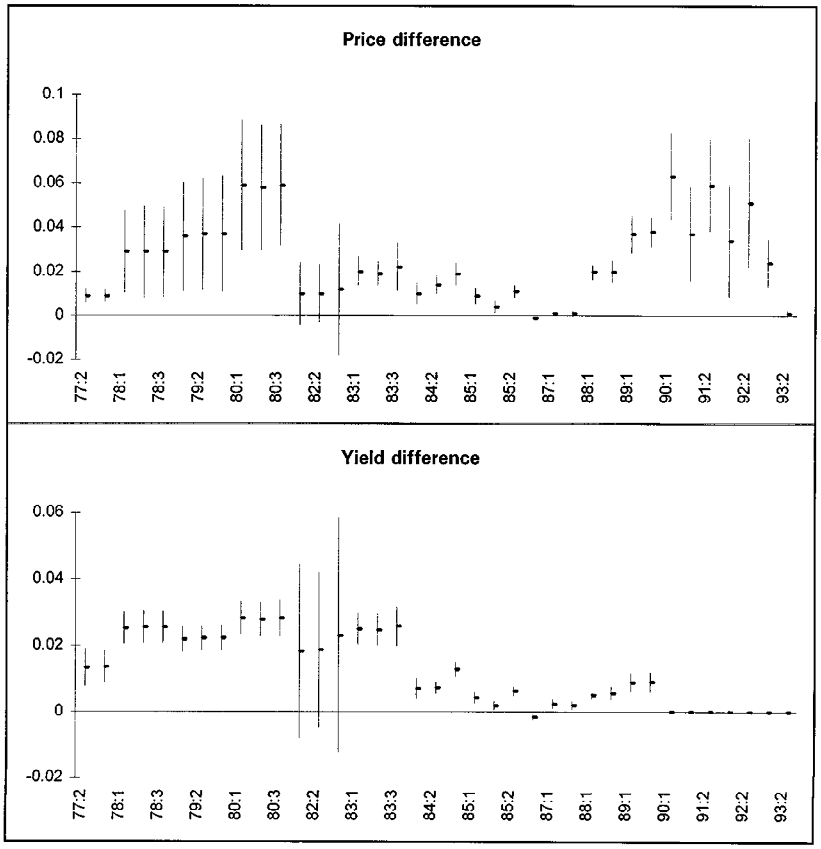

Figure 1
Average price and yield differentials
The dots represent the average difference between the price of the sequenced bonds and the mixed bonds as a fraction of face value for each bond in the sample. The vertical lines describe two standard deviations above and below the means. Standard errors are adjusted for first-order autocorrelation using the transformation $y_t - \rho y_{t-1} = \alpha_0 + e_t - \rho e_{t-1}$. Because of unequal numbers of observations, the bonds issued in the middle years have smaller standard deviations. The second panel provides the same information for the yields on mixed bonds less the yields on the sequenced bonds.

For all but one of the bonds the average price premium is positive, and for most it is significantly so. Four bonds have very few observations associated with them. The bonds issued in the early 1980s matured soon after our sample period began, and 93:2 was issued only a few months before our sample period ended. As a result, their standard errors are large and the mean is insignificantly different from zero. Three bonds from the mid-1980s have means very close to zero: 86:2, 87:1, and 87:2.

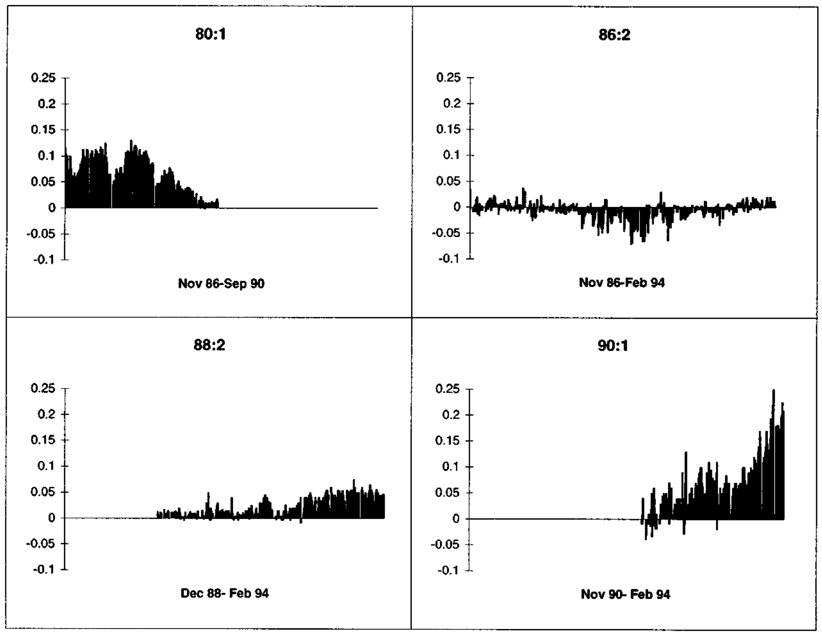

Figure 2
Daily price differentials
Each plot shows the price of sequenced bonds less the price of the mixed bonds as a fraction of face value for four different bond issues. Gaps in the plots are due to suspensions of trading around coupon dates. The horizontal axis represents calendar time over the sample period. The 80:1 issue matured halfway through the sample period, while the 88:2 and 90:1 were issued during it.

In Figure 2 we plot the daily time series of price differences as a fraction of face value for each of four bonds, which are either broadly representative or are exceptional in some important way. The figure makes clear that the averages understate the magnitudes of the premia in certain situations. From the plot of the 80:1 bond, which matured during our sample period, it is clear that the premium falls close to maturity. Similarly, the premium only becomes large some time after a bond has been issued, as can be seen in the plots for the 88:2 and 90:1 issues. From all four plots it is evident that the premium is very persistent and that once it becomes large and positive it hardly every reverses.[10]

The 90:1 issue has the largest premia in our sample, approaching 25% of face value. The 86:2 issue has a negative but insignificant

[10] For the three fixed rate bonds in the figure, the largest yield differentials (in absolute value) are 5.35% (80:1), −1.91% (86:2), and 2.3% (88:2).

average premium over its life, but the plot reveals a substantial period of time when the negative premium was quite large. This is the only bond of the 39 in our sample that, when the sample is broken into years, shows significant negative premia in any year. From 1/89–12/92 the average premium was −1.04%, with a t-statistic adjusted for first-order autocorrelation of −7.9.

The natural places to look for the source of the premium for sequences involve differences in risk and in liquidity of the bonds trading in mixed and sequenced form. To evaluate the quantitative and statistical significance of these factors, we conduct a pooled time-series and cross-sectional regression of the price difference against the following variables, which we have identified below as pertaining to either risk or liquidity, although there is necessarily some overlap between the two categories.

Risk-related variables

GUAR. The fraction of the total expected payoff subject to the partial guarantee. If the price differences are attributable to the perceived risk differences between the bonds, this variable should have a positive coefficient.

END2–END0. The variable END_i is equal to one if there are exactly i coupon payments remaining. With fewer lotteries remaining, the mixed bonds and the sequences are closer substitutes. After the final coupon they are perfect substitutes from the standpoint of lottery risk. If differences in risk determine the premium, these variables should enter the regression with negative coefficients.

Liquidity-related variables

PACK. The portion of the issue sold in sequences that cannot be broken up. This variable is nonzero for 9 of the 39 issues in our sample. Since this tends to decrease the relative supply of individual and mixed bonds, it may have a negative impact on the magnitude of the premium for the sequenced bonds.

SEAS. The minimum of 180 and the number of the days since the bond began trading. Less-seasoned bonds are likely to be held in dealer inventories and thus are probably more liquid. Further, many trades recorded as mixed early in the life of the bonds involve full sequences being broken into partial sequences, where the opportunity cost to the seller is clearly the value of the bonds as sequences.

VOLAT. An estimate of the conditional volatility of interest rates. In periods of high volatility, liquidity should be more valuable, as has been shown to be the case with other liquidity premia such as the bills-to-note spread in the U.S. Treasury market [see

Kamara (1994)]. We estimated the conditional volatility with an ARCH(1, 1).

RELVOL. Volume in the mixed bonds of the issue as a fraction of total volume in this issue, averaged over the previous 5 days. If lower liquidity for the mixed bonds is the source of the discount, it should be larger in periods of relatively low volume for the mixed bonds.

LOTT. A dummy variable equal to one for the five trading days prior to and following a lottery, and zero otherwise. The lotteries appear to be periods of relatively active trading, and this liquidity may decrease the effects of smaller, risk-averse traders on the pricing. Further, tax-arbitrage activities are concentrated in the larger blocks with partial guarantees. This may lead to the sequenced blocks being bid up relative to the mixed bonds prior to the lotteries, leading to a negative coefficient for this variable.

OLD. A dummy variable that equals 1 if the bond was issued prior to 1980. Such bonds were valuable for tax-arbitrage activities, and the sequences were easier to employ as collateral in these trades. Thus the sequences may sell at a greater premium for these bonds.

NEW 91. A dummy variable equal to 1 in the post-1990 period. During this time losses on lottery bonds could again be used to offset gains on other assets, although at much lower rates than for the *OLD* bonds.

OLS regressions of the price discount for the mixed bonds against these variables produce autocorrelated and heteroskedastic residuals. We use the Cochrane and Orcutt transformation to adjust for first-order autocorrelation, with a separate estimate of the autocorrelation parameter for each bond. We also report robust standard errors obtained using the GMM.[11]

Regression results are reported in Table 3. The variables associated with the early and final periods of a bond's life show the premium grows over the early days when the bond is traded and drops close to maturity. Two of the cross-sectional characteristics of the bonds are especially important: the portion of the coupon subject to the partial guarantee, and the portion of the issue that cannot be broken up into mixed bonds. Volume and the proximity of the lottery are insignifi-

[11] The autocorrelations for the Cochrane and Orcutt transformation, which are computed from the residuals of a first-pass OLS regression on unadjusted levels, generally lie between 0.85 and 0.95. The GMM standard errors were computed using the Newey and West weighting matrix estimated with 30 lags. Robustness checks (using lags of 0, 3, and 100 days) with respect to the lag structure did not alrter the significance of any of the coefficients, or the relative magnitudes of their t-statistics, although these generally declined as the number of lags increased. For example, the t-stastic for *PACK* ranged from -22.42 (3 lags) to -7.08 (100 lags).

Table 3
Regression of price difference between mixed bonds and 100-bond sequenced blocks on characteristics of the bonds.

Variable	Coefficient estimate	t-statistics OLS	GMM
INTERCEPT	.0001	−1.36	−0.68
GUAR	.0479	9.89	3.83
END2	−.0133	−11.06	−5.35
END1	−.0241	−21.63	−10.03
END0	−.0323	−23.64	−10.04
PACK	−.0509	−25.05	−18.73
SEAS	.0001	10.55	4.45
VOLAT	.0002	2.52	2.06
RELVOL	−.0003	−0.44	−0.47
LOTT	−.0003	−0.27	−0.23
OLD	.0021	1.85	1.43
NEW91	.0054	4.84	2.67
N		22995	
R^2		.076	

The t-statistics are adjusted for first-order autocorrelation using the Cochrane and Orcutt transformation to adjust for first-order autocorrelation, with a separate estimate of the first-order autocorrelation for each bond. The GMM t-statistics are computed using the Newey and West covariance matrix with 30 lags.

cant, but the marginally significant coefficient for volatility suggests the premium for the sequences is higher in periods where liquidity might be more highly valued.[12] These behaviors will be discussed in more detail in the next section as we consider possible explanations for the price difference.

4. Possible Explanations for the Price Differential

In Section 2 we provided sufficient conditions for lottery bonds carrying a partial guarantee to have the same price as lottery bonds in mixed form. Empirically it is quite evident the bonds in the two forms do not command the same price, which suggests one or several of Assumptions 1–7 must be violated. In this section we consider each of them in turn and examine evidence concerning their potential in explaining the anomaly.

Our discussion distinguishes between factors that might cause the bonds to be valued differently versus factors that might keep outsiders from eliminating the differential. Investors purchasing the bonds as part of their portfolios always have the choice between mixed and

[12] This does not appear to be a consequence of multicolinearity. The correlation between *VOLAT* and any of the other independent variables is less than 2%. The correlation between *RELVOL* and *LOTT* is 2.7%.

sequenced bonds. Evidently they are willing to pay more for the latter. Sources of this willingness might be perceived risk, liquidity, or tax differences. Other investors, who are not concerned with these differences, can trade profitably if a price differential emerges. Barriers to such trading will allow the mispricing to persist, but we would still expect the mispricing to be related to the characteristics that initially led investors to prefer the bonds in one form over another. The first subsection discusses the difficulties in exploiting the price differential through arbitrage or quasi-arbitrage transactions. The remaining subsections consider possible explanations for investors' evident preference for the sequenced bonds.

4.1 Assumption 1: no market frictions

Whatever the ultimate source of investors' preference for the sequenced bonds, it is quite evident that there are frictions in this market that prevent outsiders from exploiting it. Traders, who could purchase the bonds in mixed form and use them to reconstruct sequenced blocks, would earn arbitrage profits. Large investors, who could hold the bonds as a small part of a very well-diversified portfolio, could earn quasi-arbitrage profits by selling the sequenced bonds and purchasing the mixed bonds.

Transaction costs include commissions, which are generally 0.30% of face value, and the bid-ask spread, which varies between 1 and 10 SEK per 1,000 SEK face value. From January 1989–December 1992 there was also a transaction tax of 0.15% of price.[13] Since the mispricing certainly exceeds these bounds, other barriers must exist.

With regard to opportunities for true arbitrage, reconstructing sequenced bond portfolios is problematic. To complete a sequence, bonds with specific order numbers are required. There are always more order numbers than the number of bonds required for a sequence, so it is difficult to find a match. For example, the 87:1 issue had 1,000 order numbers. An investor holding, say, bonds with order numbers 1–50 would have to find another 50-bond block with order numbers 51–100, which is one of the 20 possible 50-bond sequences that could, in principle, be available. A 50-bond sequence with order numbers 151–200, for example, would not provide the partial guarantee. On the other hand, sequenced blocks can be purchased and immediately sold in smaller blocks as mixed bonds. The price differential reflects this asymmetry. With the exception of the bonds issued in 1986–1987, discussed below, there are only occasional observations

[13] See Akelius (1987), Chapter 13. We have verified through conversations with brokers that these costs are roughly accurate.

where the price of the sequences is less than the price of the bonds in mixed form, and such negative price differentials are invariably reversed in a matter of days.

It is also difficult to exploit the quasi-arbitrage opportunities of selling sequences and buying mixed bonds. International investors are effectively shut out of the market because the bonds are tax exempt and their yields are unattractive to investors not taxed in Sweden. Domestic investors face difficulties short-selling the bonds because of the awkward tax treatment of such positions and the difficulty collateralizing them. Interest payments made by the borrower of the securities to the original owner are not tax exempt. They are, of course, tax deductible to the borrower, so that if the borrower and the lender have symmetric tax rates a quasi-arbitrage position in the two types of bonds could be constructed. Every such position, though, would be nonstandard in the sense that the terms of the agreement would require either passing through a "tax-adjusted" lottery payment to the original holder, or substituting random, tax-exempt lottery income for a deterministic, taxable interest payment. Short positions must also be collateralized against the possibility that one of the shorted bonds receives the highest payments in the lottery. Relatively few domestic investors are in a position to assume as a liability a payoff distribution that is so extremely skewed.

Given these barriers to trading against the anomaly, the question of why the price difference persists becomes one of why primary purchasers of the bonds do not view the mixed bonds as a superior value. The rest of this section addresses that question.

4.2 Assumption 2: concave, expected utility preferences

It may be that investors in this market are not maximizing expected concave, state-independent utility. Thus we might look upon the evidence presented here in the context of experimental results, such as the Allais paradox and preference reversals [see Machina (1987) for a review]. In laboratory settings subjects' choices systematically violate the "linearity in probabilities" characteristic of expected utility. There is also considerable experimental evidence of "framing" effects. How choices between lotteries are evaluated depends on how the choices are described, even when objective probabilities are the same.

Violations of concavity may also be in evidence here. Convexity of the utility index over some range has long been appealed to by economists to explain a preference for lotteries with highly skewed payoffs. Friedman and Savage (1948) argued that preferences that were concave for low levels of wealth and convex over high levels of wealth could rationalize purchasing lotteries simultaneously with life insurance. Similar reasoning might motivate the packaging of coupon lotteries with a partial guarantee. This would suggest that positive value for the partial guarantee need not imply the bonds would earn higher revenue for the treasury were the lottery simply eliminated. The negative coefficient on *PACK*, and the reversal of the premium for the 86:2 issue, are direct evidence that some investors either value the skewness from the lotteries from individual bonds, which is reduced by diversification, or face wealth-related liquidity constraints, which limit their participation in the market through sequences.

This does not explain, however, why investors would buy the sequences when portfolios of mixed bonds offer close substitutes at much lower prices. This would also be awkward to explain with the unexpected-utility preferences that have been used to rationalize the experimental results. In most of our data the relative prices imply that a portfolio with a partial guarantee provides a more desirable payoff distribution than either an individual bond or a portfolio of mixed bonds. The classic central limit theorems for bounded random variables imply that an equally weighted portfolio of mixed bonds converges in distribution to an equally weighted portfolio with the partial guarantee [see, for example, Chung (1974), Theorem 5.1.1 and 4.4.5]. Given the rest of an investor's intertemporal consumption allocation, and the statistical independence of the lottery payoffs from that allocation, this in turn implies that total consumption allocations associated with the two portfolios converge.[14] Therefore, any preferences that are continuous in the probability distributions of consumption will ascribe close to equal utility to sufficiently diversified portfolio positions in mixed and sequenced bonds.

The unexpected-utility preferences used to rationalize choices that are nonlinear in probabilities all have the property that they are continuous in the distribution of payoffs. While there is no natural distinction between "systematic" and "idiosyncratic" risk for these more general preferences, there is no reason they should value portfolios of mixed and sequenced bonds very differently. What would be required would be preferences that are "discontinuous at certainty" in gambles that represent only a part of an asset's payoff and only a part of an agent's wealth.

Such discontinuities could, of course, be viewed as a "framing" effect. Something that is described as a "guarantee" is percieved as very different than something close, but still risky. Framing effects, however, are simply evidence that subjects do not appear to understand the objective probabilistic structure of lotteries they are confronted with, a possibility we discuss later in this section.

[14] This is a consequence of Chung (1974), p. 92.

4.3 Assumption 3: no taxes

Lottery proceeds from the bonds were, over the period of our study, either untaxed or taxed at a flat rate that applies to all lotteries in Sweden. Linear, symmetric taxation has no effect on the portfolio separation arguments in Section 2.[15] Capital gains are taxed on the bonds, and taxed in a nonlinear fashion. For the capital gains tax liability to differ across the mixed and sequenced bonds, however, there has to be some reason for the prices of mixed bonds to differ from the prices of the sequenced bonds in the first place. If future prices are the same, the capital gains taxes paid on bonds in each form must, of necessity, also be the same. This argument rules out "tax yield" effects, due to differential taxation of coupon and capital gain, as a source of the anomaly. While coupons on lottery bonds are untaxed, and capital gains are, there is no asymmetry in how a portfolio of mixed bonds would be taxed relative to a portfolio of sequenced bonds as long as their prices are equal. Thus, even in the presence of such taxation, equal prices are an equilibrium.

4.4 Assumption 4: investors have rational expectations and understand nonsystematic risk

One possible explanation of the premium for sequenced bonds is that investors in this market simply fail to understand that lottery risk will have little effect on their wealth or consumption if their positions in the bonds are diversified, and that the partial guarantee is largely irrelevant in its effect on the marginal distribution of their portfolio of lottery bonds. Discussions with practitioners in the market indicate very clearly that this is what traders and Swedish Treasury officials believe, and there are a number of regularities in the data consistent with risk as a source of the premium.

First, there is the fact that the larger the effect of the partial guarantee on the distribution of the portfolio's payoff, the larger is the premium. This is visually evident in Figures 1 and 2 in the high premium for the sequences in the bonds issued in the early 1990s, which have variable coupon rates. The rates on these bonds are adjusted by altering the higher payoffs in the lottery, leaving the payoff subject to the partial guarantee unchanged. The time period where the premium for the bond is extremely high is one of falling interest rates, and thus a larger and larger fraction of the expected payoff is subject to the partial guarantee. The premium for the sequenced bonds appears to grow accordingly. At the points in time where the premia are

[15] Let d_{it} be the payoff on asset i at date t. Then if d_{it} is distribed as $d_{jt} + \epsilon$ where $E_{t-1}\{\epsilon \mid d_{jt}\} = 0$, it will also be the case that $(1-\tau)d_{it}$ is distributed as $(1-\tau)d_{jt} + \epsilon^*$ where $E_{t-1}\{\epsilon^* \mid (1-\tau)d_{jt}\} = 0$, and $\epsilon^* = (1-\tau)\epsilon$.

largest, these bonds were paying rates close to their floors of 5.1% or 5.4%, while the partial guarantee applied to payments equaling 4.5% of face value. In our regressions this effect is apparent in the explanatory power of *GUAR*, the fraction of the expected coupon that can be guaranteed by holding the sequence. In both cross section and in time series the discount is larger when the partial guarantee applies to a larger portion of the payout. It is not apparent why this variable would be related to differences in liquidity of the two types of bonds. Neither is the difference in risk the partial guarantee generates economically meaningful from the standpoint of an investor taking a diversified position. Nevertheless, it clearly has a significant effect on prices.

Particularly striking, and strong evidence that the guarantee is priced, is the maturity effect apparent in the price difference. The bonds continue to trade for half a coupon period after the final coupon lottery. At this point the mixed bonds and the 100-bond sequences are obvious perfect substitutes, and they are priced as such. Figure 3 shows the average price difference plotted against time to the final coupon. Date zero is the last day of trading before the final lottery. The plot in the first panel includes all the bonds that were exchanged for newly issued bonds at maturity. In these exchanges, holders of maturing bonds receive new bonds in sequenced form. The plot in the second panel includes all the bonds that were retired for cash. For both groups the average price difference drops dramatically immediately following the last lottery, and appears to fluctuate randomly around zero thereafter.[16] This behavior is not unique to the last lottery. In Table 3 the variables *END0–END1* all have significant negative coefficients. Neither do these variables show high degrees of multicollinearity with obvious liquidity proxies.[17]

Further evidence that it is the difference in lottery risk that is being priced is found in the behavior of 50-bond sequences, depicted in Figure 4. Recall that the exchange quotes prices for blocks of 50 bonds, in sequence, as well as blocks of 100. The 50-bond sequences, with the exception of the 93:1–93:2 issues, do not qualify for the partial guarantee.

The figure shows, for each bond issue, the average premium for 50-bond blocks over mixed bonds as a fraction of the premium for 100-bond blocks over mixed bonds. In most cases the 50-bond blocks

[16] For all the 1,077 days in our sample after a final lottery, the average value for $\frac{P_{100}-P_{mix}}{F}$ is -0.00062, with an associated t-statistic of -3.1. While different from zero in a statistical sense, this price difference is not economically significant. It is well within transactions costs.

[17] The correlations between the *END* dummy variables and *RELVOL* and *LOTT* are, respectively, 0.12 and 0.05 (*END0*), 0.10 and 0.08 (*END1*), and 0.01 and 0.05 (*END2*).

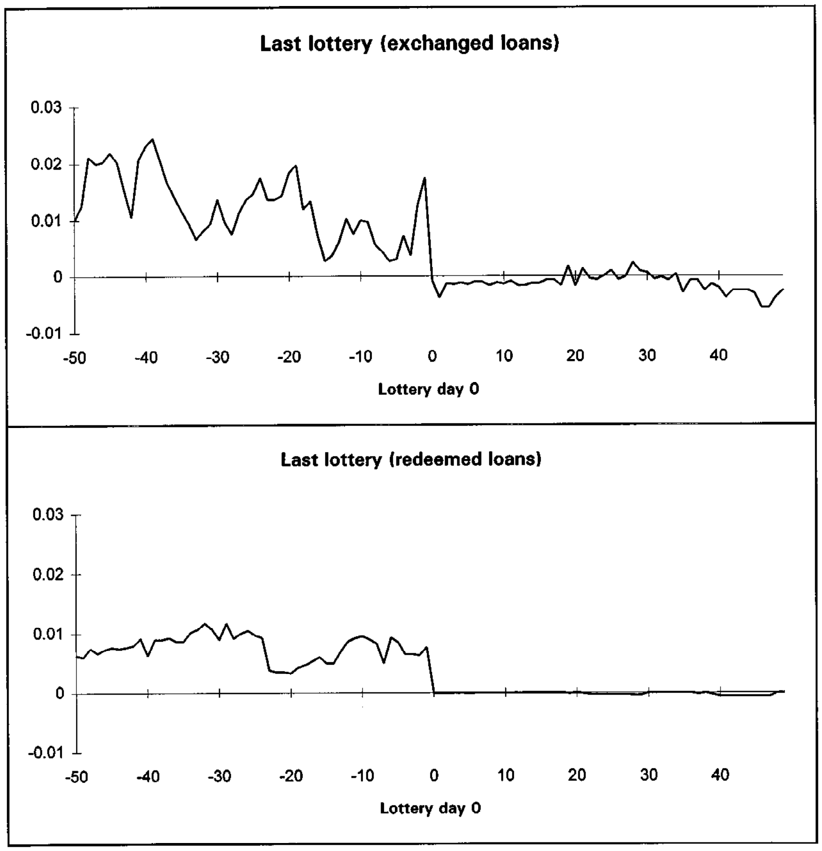

Figure 3
Premium behavior around final coupon
The two panels plot the average premium for sequenced bonds as a percentage of face value on trading days surrounding the final coupon lottery. The averages plotted in the first panel are for all bonds that were exchanged at maturity for new issues. In these exchanges, holders of mixed or sequenced bonds both received sequences. The second panel plots averages for bonds that were retired for cash at maturity.

sell at average premia that are above 50% of the total premia. This suggests perceived risk as a source of the premium. Holding a 50-bond sequenced block rather than a single bond increases the chances of receiving the smallest payment from 0.01 to 0.5, and, accordingly the 50-blocks are priced to reflect half of the risk premium. Moreover, the premia for the 50-bond blocks manifest the same behaviors in cross section and time series as the premia for the 100-bond blocks. Regressions, not reported in detail, of the 50-bond block premia on the variables in Table 3 show exactly the same patterns

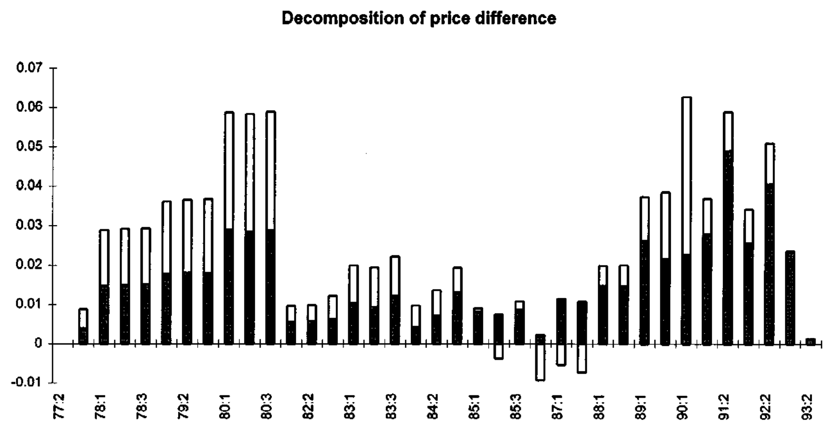

Figure 4
Premium for 50-bond and 100-bond sequenced blocks
Sequenced blocks of 50 bonds increase the probability of receiving the smallest payment. For each bond issue, the black bar represents the average premium of the 50-bond blocks over mixed bonds. The white portion of the bar represents the average premium of 100-bond sequences over 50-bond sequences.

of sign and significance as those reported for the 100-bond block premia.

4.5 Assumption 5: perfect liquidity

Brokers and dealers report considerable variation in liquidity across bond issues and through time. Based on volume data, the most liquid market for Swedish lottery bonds is the market for 100-bond sequences. Relative liquidity could therefore be a source of the premium.

Some features of the behavior we see here constrasts with liquidity effects that have been observed in other markets. Kamara (1994) reports price differences for notes and bills, despite the fact that they are default-free perfect substitutes. Price differences between coupon strips and principle strips, which appear to be due to liquidity, also persist despite the perfect substitutability of the cash flows. In the lottery bond market, on the other hand, at the point when the bonds become perfect substitutes in terms of relative risk, the price differential disappears.

Evidence for a liquidity-based explanation of the price difference is conflicting. The premium for the sequences does not appear to respond dramatically to the variables researchers have shown are important in explaining other liquidity premia. Volume is insignificant in our regressions. Volatility measures have marginally significant explanatory power, however, which would be consistent with a view of

Table 4
Descriptive statistics for large blocks with same coupon guarantee

Issues	Variable	N	Mean	Min	Max
93:1–93:2	$\frac{P_{100}-P_{50}}{F}$	217	.000	−.035	.020
82:1–87:2	$\frac{P_{1000}-P_{100}}{F}$	10410	.000	−.055	.060
86:2–87:2	$\frac{P_{500}-P_{100}}{F}$	4173	−.002	−.055	.030
86:2–87:2	$\frac{P_{1000}-P_{500}}{F}$	4176	.000	−.040	.060

The price differences are fractions of face value for blocks of bonds with sequential order numbers carrying the same partial coupon guarantee. Observations are pooled for all bonds trading in the forms indicated by the subscripts in column 2. For the 93:1 and 93:2 issues the partial guarantee is achieved with 50 bonds, while for other issues in the table, sequences of 100 are needed.

the sequences as safer when liquidity is most important. These findings are robust to the choice of proxies for volume and volatility.[18]

For many of the bonds in our dataset, the Stockholm Stock Exchange reports prices of larger blocks than the minimum number needed in sequence to qualify for the partial guarantee. The larger blocks are reportedly less liquid than the 100-bond blocks, and this is borne out in the volume statistics in Table 2. Since these blocks carry the same guarantee as the 100-bond blocks, their prices are informative about the importance of liquidity versus risk in pricing. Further, the larger blocks offer more pure diversification than do the 100-bond blocks. If diversification itself, rather than the partial guarantee, were priced in this market, we would expect to see that here. Table 4 reports statistics concerning the price differences between differently sized sequential blocks that carry the partial guarantee. The mean price differences are very close to zero, and even the extreme values are small relative to the premia at which the sequences sell over the mixed bonds.

In summary, the larger blocks, while they may differ in liquidity from the 100-bond sequences, do not show persistent, economically significant price discounts, in marked contrast to the behavior of the mixed bonds and 50-bond sequences, which do not qualify for the partial guarantee and sell at large discounts.

[18] We ran the regressions in Table 3 with a number of alternative specifications for the conditional volatility, including more naive ones such as simple averages of lagged squared deviations from the mean, as well as less parsimonious ARCH and GARCH models. None of these produced different results. We have also run the regressions in Table 3 using total volume, with no consequences for the results.

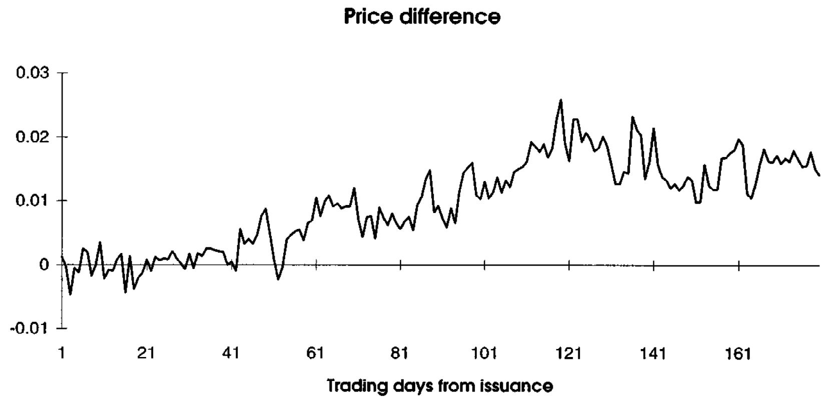

Figure 5
Seasoning effects
The plot shows the average price difference between mixed and sequenced bonds as a fraction of face value against days from the start of trading. The premium for sequenced bonds increases as the bond becomes seasoned.

There are, however, some characteristics of the relative pricing of the mixed and sequenced bonds that appear to be due to "liquidity," in the broad sense that they are related to difficulties trading the bonds in one form versus another.

First, there is a seasoning effect, shown clearly in Figure 5, which plots the average premium as a percentage of face value against days from issuance. The price differences are initially small and centered around zero. The regression results in Table 3 support the evidence in the figure. The variable *SEAS*, which measures days since issuance, has a significant positive coefficient. This behavior appears to be due to the way the bonds are distributed. The initial owners of the bonds hold them in large sequential blocks. When they sell pieces of these blocks to investors as mixed bonds, the opportunity cost to them is the price they would receive by selling the bonds in sequenced form. They appear to demand, and to receive, this value in the initial trades. This has obvious implications for the incidence of the discount. Neither the government nor the dealers who initially buy the bonds appear to bear this cost. It is borne by investors who first purchase the bonds in mixed form.

There is one bond in our sample, the 86:2 issue, with a negative average price differential. Several issues in the mid-1980s included sequences that investors were prohibited from breaking apart. The fraction of the issue packaged this way, the variable *PACK*, has a significantly negative coefficient in our regressions. For the other bonds with such prohibitions, *PACK* varies between 26% and 54%, while for the 86:2 issue, 78% of the sequences cannot be transformed to mixed bonds. This could be evidence that, in some cases, the marginal investor values the skewness from the lottery, and so holds individual bonds. Or it could be a wealth-related liquidity effect. Investors with high valuations for the bond lack the wealth to buy sequences.

The evidence presented in Figure 4 suggests that these effects offset a positive valuation of the partial guarantee that is still at work in the pricing of these particular bonds. The two 1987 issues, which have positive but insignificant price differentials, also had substantial portions issued in unbreakable blocks, and were close substitutes for the 86:2 issue in other respects. Blocks of 50 for these issues sell at a *premium* relative to the mixed bonds, while blocks of 100 sell at a *discount* relative to blocks of 50. The blocks of 50 can, of course, be broken apart to sell as individual bonds in the retail market, while the blocks of 100 cannot. Sequences of 50, nevertheless, have less variance associated with the lowest payment level than does a portfolio of 50 mixed bonds or an individual bond. This lower lottery risk commands a premium in pricing, just as is the case for the 100-bond blocks for the other bonds in our sample.

Another liquidity-related phenomena involves the role of the bonds as collateral in the tax-arbitrage strategies discussed in Section 2.3. While the higher handling costs for mixed bonds have obviously been mitigated by modern data processing, and banks are now willing to accept the mixed bonds as collateral, it is an empirical issue whether this advantage explains some or all of the higher price for the sequences. The Swedish government has attempted to limit this tax strategy by changing the treatment of capital losses on lottery bonds. Our sample includes data from three regimes: bonds issued before 1980, for which losses could be fully utilized; bonds issued after 1980, when loss utilization was limited, and bonds trading after 1990, generally generating tax savings of 21% of the value of the loss. In addition, over our sample period the tax rates on investment income, and hence the marginal value of this tax-arbitrage trading, has changed.[19]

The effects of these tax-arbitrage strategies is very evident in volume statistics. Table 5 reports on abnormal volume and relative volume around the lotteries.

[19] The top marginal capital gains rate was 88% in 1979. This had dropped to 80% by 1986, when our sample period begins. It continued to drop year by year until it reached 65% in 1990. The tax reform of that year introduced a flat rate of 30% for the taxation of capital gains. Capital losses generally generate tax savings at a rate 21%. Up to 70% of a capital loss on lottery bonds can be used to offset other investment income, which is taxed at 30%. If capital losses exceed investment income and capital gains, then as long as the taxpayer has other taxable income, they generate a tax credit of 21% of the loss for the first 100 thousand SEK, and 14% of any additional losses.

Table 5
Average volume around lottery dates

	Daily volume		
	Pre-1980	Post-1980	Post-1990
10 days around lottery	9.3	6.5	2.2
All other days	2.4	3.1	1.3
	Mixed to Total Volume		
	Pre-1980	Post-1980	Post-1990
10 days around lottery	.353	.162	.268
All other days	.396	.167	.280

The top panel reports averages of daily volume in millions of SEK. The lower panel reports averages of this volume which consists of trades in mixed form. Abnormal volume around lottery dates suggests trading for purposes of tax arbitrage. Tax treatment of losses on bonds issued before 1980 differed from those issued after 1980. The treatment of losses for all bonds, regardless of date of issuance, was changed in 1990.

The old bonds show dramatic increases in volume around the lottery, but trading in mixed bonds also increases. The post-1980 bonds also show abnormal volume around lotteries. This appears to be a spillover effect. Buyers without losses from other sources would participate as buyers in the ex-coupon periods and sellers prior to the lottery, thus generating capital gains that could be used against losses from any of the bonds. The evidence of abnormal volume is much less pronounced in the recent period, consistent with claims by dealers and brokers that this activity has ceased.

Higher collateral value for the sequences is consistent with several features of the price data. Figure 1 suggests the premium for sequences drops between the 80:3 issue and the 82:1 issue, and increases again with the 90:1 issue. Second, the fact that the premium drops as the bonds approach maturity would also be consistent with the differential handling costs diminishing, as these costs are only incurred on lottery dates.

It is also evident, however, that this is not the only source of the pricing differential. The bond issues from 81:1 through 85:3 had all matured by the end of 1991. As can be seen in Figure 1 these bonds still show positive premia for the sequences. (Because these bonds are close to maturity in our sample, the premia are more evident in yields than in prices.) In addition, there are very large premia for the bonds issued in the early 1990s, despite the fact that in terms of volume the tax-arbitrage trading appears to be a much less important factor for these bonds. This is also a period in which banks allow both types of bonds to be used as collateral. The variable *LOTT*, which measures the days to the next lottery, is insignificant in our regressions. This suggests that the source of the average premium is not due to tax arbitraguers bidding up the price of the sequences immediately around the lotteries.

Perhaps the strongest evidence that the premium cannot be fully explained through differences in handling cost is in the pricing of the 50-bond blocks, shown in Figure 4. While 50-bond blocks do not qualify for the partial guarantee, and thus carry more perceived lottery risk, the handling costs for blocks of 50 bonds differ very little from the handling costs for blocks of 100, and they have always been acceptable as collateral in tax-arbitrage positions. Nevertheless, the 50-bond blocks are priced at a substantial discount to the 100-bond blocks, and at a premium over the mixed bonds.

To evaluate the importance of the differential value in tax arbitrage, our regressions include dummy variables for the two periods in which the bonds could be used to generate losses that offset other income. These variables do enter with the predicted sign. Curiously though it is the dummy for the post-1990 regime, rather than the dummy for the older bonds, that is significant, despite the fact that during the post-1990 period the forward market that supported the arbitrage trading is no longer active and reports by dealers that the tax-arbitrage trading has ceased.

4.6 Assumptions 7 and 8: extrinsic versus intrinsic uncertainty

Following Cass and Shell (1983), a large literature in economics has shown that rational expectations equilibria can respond to "sunspots"—purely extrinsic uncertainty not reflected in technology, preferences, or endowments. Dynamic equilibria may exist in which lottery risk is priced, and because it is priced investors would optimally hold undiversified positions in equilibrium.

While possible in principle, lottery risk seems an unnatural candidate for a sunspot. Unlike events such as sunspots, or who wins the Super Bowl, the outcomes of the coupon lotteries are not simple public events that would be easy to coordinate on. The only uncertainty involves who gets which payment.[20] Sunspot equilibria would also have the characteristic that other economic outcomes were correlated with lottery outcomes, and it is very difficult to imagine how one might evaluate this, even in the most casual empirical manner.

[20] The "sunspot" would have to be something like, "If the winner's name is Sven, it will be a good year for the economy, and if the winner's name is Eric, that suggests a bad year."

5. Conclusions and Implications

This article documents the behavior of price differences between two securities that, in principle, should be perfect substitutes in investors' portfolios: mixed and sequenced portfolios of lottery bonds. The two securities are nevertheless priced quite differently.

If one concludes from this evidence that investors in the market are averse to lottery risk, and that prices reflect this, then this may have policy implications. The Swedish Treasury is one party in a position to exploit the anomaly by restructuring or even eliminating the lottery feature. Indeed, the lottery bond issue most recently announced, which is not in our dataset, has a face value of 5,000 SEK, and sequences of only 10 bonds qualify for the partial guarantee, which represents almost 60% of the total expected coupon. Thus the treasury appears to be issuing bonds that our results suggest would be more attractive to investors. It must be kept in mind, though, that this involves extrapolating outside the context of the experiment examined here. It is possible that investors are averse to the lottery risk associated with the smaller payment levels, yet still value the chance at very high payoffs. This, in turn, might suggest guaranteeing the lowest coupon level for all bonds. But the fact that the price differential does not emerge until well after the bonds have been issued may imply the cost of the loss of the partial guarantee is not borne by the treasury in any case.

Appendix A: Proofs of Propositions

A.1 Proof of Proposition 1

First consider portfolios (1) and (2). Suppose d_0 is the payment subject to the partial guarantee. Let $Z(K)$ be a random variable with a payoff equal to the receipts to K bonds from the lotteries, drawn without replacement, involving all the bonds in the issue. In the mixed bond portfolio, let $X(J_1, \ldots, J_M)$ represent the random payoff to the bonds from the lotteries based on order numbers alone. Per unit of face value, the mixed bond portfolio has the same marginal distribution as

$$Y_K^M = \frac{1}{K} X(J_1, \ldots, J_M) + \frac{1}{K} Z(K), \tag{4}$$

while the sequenced portfolio is distributed as

$$Y_K^S = \frac{1}{K} d_0 + \frac{1}{K} Z(K). \tag{5}$$

Next, define $\epsilon \equiv \frac{1}{K} X(J_1, \ldots, J_M) - \frac{1}{K} d_0$. The probability that order number i is drawn is $\frac{1}{K}$, and if J_i bonds with this order number are held, they will receive $J_i d_0$. Thus,

$$E\{X(J_1, \ldots, J_M)\} = \sum_{i=1}^{M} \frac{1}{K} J_i d_0 = \frac{1}{K} d_0 \sum_{i=1}^{M} J_i = d_0 \tag{6}$$

as $\sum_{i=1}^{M} J_i = K$. It follows that ϵ has unconditional mean of zero. It also has mean zero conditional on $Z(K)$, since the lotteries over the order numbers alone are independent of the lotteries for the higher payments involving all the bonds. Therefore, $E\{\epsilon \mid Y_K^S\} = 0$, and $Y_K^M \cong Y_K^S + \epsilon$ where $\cong$ denotes equality in distribution. This proves the first part of the proposition.

Now consider the sequenced portfolio and an equally weighted portfolio of all the bonds in the issue. Since the payoff on the latter is deterministic, second-order stochastic dominance follows if we can simply show that the two portfolios have the same expected payoff. There are $\frac{NM}{K}$ sequenced portfolios in the issue. Let Y_{Ki}^S denote the payoff on the ith such portfolio. The aggregate portfolio has a payoff of

$$Y_A = \sum_{i=1}^{\frac{NM}{K}} \frac{Y_{Ki}^S}{\frac{NM}{K}}. \tag{7}$$

Since the payoffs Y_{Ki}^S are identically distributed (although not independent, because drawings are without replacement),

$$E(Y_A) = \left(\frac{NM}{K}\right) \frac{E(Y_{Ki}^S)}{\frac{NM}{K}} = E(Y_{Ki}^S). \tag{8}$$

To prove the final part of the proposition, let c_t denote an agent's consumption at date t, exclusive of the proceeds from any lottery position. Suppose $Y_a \cong Y_b + \epsilon$ with $E(\epsilon \mid Y_b) = 0$. Then we must show that

$$c_t + Y_a \cong c_t + Y_b + \epsilon \tag{9}$$

with

$$E(\epsilon \mid c_t + Y_b) = 0. \tag{10}$$

This will follow as long as $E(\epsilon \mid c_t) = 0$, an immediate consequence of the fact that, in our case, the ϵ term is simply lottery risk, and the lotteries are by construction statistically independent of other variables in the economy. ■

A.2 Proof of Proposition 2

Let c_t be an agent's equilibrium allocation at date t, let β be the discount factor, and let $u(\cdot)$ be the periodic utility index. Then the equilibrium price of any asset must satisfy

$$P_{it} = E_t \left\{ \beta \frac{u'(c_{t+1})}{u'(c_t)} (x_{i,t+1} + P_{i,t+1}) \right\}, \tag{11}$$

where x_{it} is the periodic payoff on the asset. Suppose then that

$$x_{j,t+1} = x_{i,t+1} + \epsilon_{t+1} \tag{12}$$

where

$$E_t\{\epsilon_{t+1} \mid x_{i,t+1}\} = 0. \tag{13}$$

In our case, ϵ is lottery risk, with zero mean, and assets i and j are a portfolio of sequenced bonds and a portfolio of mixed bonds, respectively. We know that the lotteries are drawn independently over time so that the error is mean independent of time-t information, and

$$E_t\{\epsilon_{t+1}\} = 0. \tag{14}$$

Thus we are only requiring that ϵ_{t+1}, in addition, be mean independent of asset i's payoff, which was demonstrated in Proposition 1.

We now conjecture that (1) the agent's consumption allocation is free of lottery risk, and (2) $P_{i,t+1} = P_{j,t+1}$. We will show this implies $P_{it} = P_{jt}$. If aggregate holdings of the two assets are free of the residual risk, then $E_t\{\epsilon_{t+1} \mid c_{t+1}\} = 0$. Under the conjecture that $P_{i,t+1} = P_{j,t+1}$ we obtain

$$\begin{aligned} P_{jt} - P_{it} &= E_t \left\{ \beta \frac{u'(c_{t+1})}{u'(c_t)} \epsilon_{t+1} \right\} \\ &= E_t \left\{ \beta \frac{u'(c_{t+1})}{u'(c_t)} E_t\{\epsilon_{t+1} \mid c_{t+1}\} \right\} \\ &= 0. \end{aligned} \tag{15}$$

At maturity, the sequenced and mixed bonds have the same payoff. Treating this as date $t + 1$, we can argue recursively that the prices are equal at each previous date.

We now return to the conjecture that the agent's consumption allocation has no lottery risk. Given prices for mixed and sequenced bonds that are equal at each date, this follows from the additivity of the utility function and Proposition 1. It is also consistent with market clearing since the aggregate allocation has no lottery risk. Thus it is optimal for agents to bear no lottery risk and markets clear, when the risk is unpriced, and when agents behave this way, lottery risk is unpriced by their first-order conditions. ■

A.3 Proof of Proposition 3

In any general equilibrium that is Pareto optimal, individual consumptions are measurable with respect to aggregate consumption. [See Chamberlain (1988) for a formal demonstration of this point.] In aggregate there is no lottery risk. Thus, individual consumptions have no lottery risk, and the result follows from the first-order conditions as in the proof of Proposition 2.

If at least one agent bears no lottery risk, then, again, the result follows from a consideration of the agent's first-order condition as in Proposition 2. ■

References

Akelius, R., 1987, *Skatt och Placeringar*, Akelius Publishing, Jönköping, Sweden.

Akelius, R., 1980, *Tjäna Mycket Mera på Premieobligationer*, Akelius Publishing, Jönköping, Sweden.

Bessembinder, H., 1992, "Systematic Risk, Hedging Pressure, and Risk Premiums in Futures Markets," *Review of Financial Studies*, 5, 637–668.

Breeden, D. T., 1979, "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 7, 265–296.

Cass, D., and K. Shell, 1983 "Do Sunspots Matter?" *Journal of Political Economy*, 91, 193–227.

Chamberlain, G., 1988, "Asset Pricing in Multiperiod Securities Markets," *Econometrica*, 56, 1283–1300.

Chung, K. L., 1974, *A Course in Probability Theory*, Academic Press, New York.

Cochrane, D., and G. H. Orcutt, 1949, "Application of Least Squares Regression to Relationships Containing Auto-Correlated Error Terms," *Journal of the American Statistical Association*, 44, 356–372.

Connor, G., 1984, "A Unified Beta Pricing Theory," *Journal of Economic Theory*, 34, 13–31.

Friedman, M., and L. J. Savage, 1948, "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy*, 56, 279–304.

Hirshleifer, D., 1988, "Residual Risk, Trading Cost, and Commodity Futures Risk Premia," *Review of Financial Studies*, 1, 173–193.

Kamara, A., 1994, "Liquidity, Taxes, and Short-Term Treasury Yields," *Journal of Financial and Quantitative Analysis*, 29, 403–417.

Lucas, R. E., 1978, "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.

Machina, M. J., 1987, "Choice Under Uncertainty: Problems Solved and Unsolved," *Journal of Economic Perspectives*, 1, 121–154.

Newey, W. K., and K. D. West, 1987, "A Simple, Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica*, 55, 703–708.

Ross, S., 1976, "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory*, 13, 341–360.

Ross, S., 1978 "Mutual Fund Separation in Financial Theory—The Separating Distributions," *Journal of Economic Theory*, 17, 254–286.

Swedish Treasury, 1991, *Riksgäldskontorets Faktabok om Premieobligationer 1991*, Swedish Treasury, Stockholm, Sweden.