

The Evolution of Security Designs

THOMAS H. NOE, MICHAEL J. REBELLO, and JUN WANG*

ABSTRACT

We consider a competitive and perfect financial market in which agents have heterogeneous cash flow valuations. Instead of assuming that agents are endowed with rational expectations, we model their behavior as the product of adaptive learning. Our results demonstrate that adaptive learning affects security design profoundly, with securities mispriced even in the long run and optimal designs trading off underpricing against intrinsic value maximization. The evolutionary dominant security design calls for issuing securities that engender large losses with a small but positive probability, but that otherwise produce stable payoffs, almost the exact opposite of the pure state claims that are optimal in the rational expectations framework.

THE REAL OPERATIONS OF THE ECONOMY PRODUCE RANDOM CASH FLOWS. If investors have varying attitudes toward risk, they will have different valuations of these risky cash flows. For this reason, the received wisdom is that investor welfare can be improved by sharing rules that allocate risky cash flows among investors. Financial markets can serve as a mechanism for implementing such welfare-improving risk sharing. These markets enable agents with different risk preferences to trade securities and, through their trades, establish equilibrium prices for risky assets. In equilibrium, state prices are determined by the marginal investor's valuation; and the marginal investor may be different in different states of the world.

Although the welfare-enhancing effects of dividing and marketing cash flows have long been recognized, Allen and Gale (1988) are the first researchers to consider the matrix of risk-sharing securities not as an exogenous datum but rather as an endogenous product of design. In particular, Allen and Gale develop a model of financial markets in which corporations profit by designing securities. They show that, in the absence of transactions costs, the optimal security design calls for the issuance of pure state securities, that is, firms should issue a bundle of many securities, with each security paying off only in a single state of the world and only one security paying off in any given state. Such a design captures the highest marginal valuation for each security across

*Noe and Rebello are from the A. B. Freeman School of Business, Tulane University; Wang is from Baruch College and Louisiana State University. The authors would like to thank the seminar participants at Louisiana State University and the Third International Workshop on Computational Intelligence and Finance for their comments. Rebello would like to thank the Robinson College of Business for research support. We also extend special thanks to Rob Stambaugh (the editor) and two anonymous reviewers for very insightful comments. Finally, we thank Susan Bergman for editorial help. All errors are our own.

the heterogeneous pool of investors. However, in actuality, corporations tend to issue only a few securities that often provide positive payoffs simultaneously by splitting the cash flows associated with individual states across multiple securities. Indeed, Allen and Gale show further that given transaction costs that are proportional to the number of securities a firm issues, a firm will reduce the number of securities it issues by issuing bundles of pure securities and thereby bundling together cash flows from multiple states. Nevertheless, splitting a given state's cash flow among multiple securities is still suboptimal.

By relaxing the assumptions of no transactions costs, free information, and competitive financial markets, the extant literature attempts to rationalize the stylized fact of firms issuing a narrow range of securities in a world with a broad range of investor preferences. For instance, Madan and Soubra (1991) show that marketing costs can make both the bundling and splitting of cash flows optimal. In their analysis, marketing costs interfere with the firm's ability to match its cash flows with the investors who have the highest marginal valuations for these cash flows. Madan and Soubra start by assuming that security prices are not determined in competitive markets by Bertrand competition, but rather by firms approaching investors one at a time with first-and-final offers. When the firm approaches an investor, the firm is uncertain of the structure of that investor's preferences. Attempting to sell a security that is likely to appeal to only a narrow range of investors is risky; an investor is likely to reject such a claim, in which case more sales effort is required and in turn more marketing costs are sustained. For this reason, firms design securities that appeal to a broad range of investors, with claims that pay off in multiple states and split cash flows from a given state among securities.

While Madan and Soubra rationalize bundling of cash flows from different states and splitting of cash flows from a given state based on a firm's uncertainty about investor preferences, Gale (1992) rationalizes bundling and splitting based on investor uncertainty with respect to their own preferences. In particular, he develops a model of security design in which the investor, because of uncertainty about the realization of state-dependent future income, does not know his own marginal valuations across states and determining these valuations requires costly investment in information. Gale shows that, in this setting, coordination-failure equilibria exist in which only a narrow range of securities is issued even if a broader range of claims is optimal and efficient.

In contrast to the above literature, we explain bundling and splitting of cash flows based on a *behavioral* deviation from the rational expectations paradigm, where this deviation is rooted in the fact that real investors learn how to price securities through experience. Rather than assuming any fixed behavioral bias, we assume initially random agent behavior together with an evolutionary dynamic that weeds out suboptimal strategies. Specifically, we follow the approach used in Noe, Rebello, and Wang (2003) and model the evolution of security design as a genetic algorithm. However, in contrast to Noe, Rebello, and Wang, who assume homogeneous risk-neutral investors, we consider a world similar to that of Allen and Gale (1988) in which investors have different state-specific marginal utilities for income. In this world, firms, which are artificial agents

whose behavior is governed by a genetic algorithm, select a security package to sell to individual investors, who are also artificial agents governed by a genetic algorithm. These investors compete for securities using classical Bertrand price competition, that is, the highest bidder wins the security. After the results of a given round of competition, agents update the strategy sets, with low payoff strategies weeded out of and high payoff strategies selected into the strategy set. Each run of each simulation lasts at least 10,000 rounds; the results we report in this paper correspond to the outcomes of the simulations at the end of these runs. These results aim to capture the long-run behavior of financial markets. At first glance, performing 10,000-round simulations may seem excessive, especially if one thinks of each investor as starting his investment career with a blank slate and observing only the securities that he purchases. However, we take the perspective of von Hayek (1945), formalized by Young (1998), which argues that each agent starts his career armed with decision heuristics that impound the learning of earlier agents, and the agent modifies this heuristic based not only on the outcomes of his own decisions but also on his observations of many other agents' decisions. Thus, in this context, our long-run simulations do not seem unrealistic.

The results of our simulations indicate that even in a competitive, zero-transactions cost environment, adaptive learning has a profound effect on the design of securities by corporations. In general, even after the learning process has reached convergence, securities are underpriced. Further, the dominant security designs differ substantially from a basket of pure state securities: Dominant designs do not consist of risk-free securities or securities that pay off only in a few states of nature, rather, dominant designs bundle cash flows from different states and split cash flows from the same state to ensure that investors incur losses in only a few states of the world. Note that while the dominant security designs we obtain differ from the zero-transactions cost designs of Allen and Gale in this respect, both share the property that they complete the securities market.

To understand these results, first consider the pricing of a fixed security in a world of adaptive investors with heterogeneous valuations of cash flows across states. Adaptive learning is both (a) conservative, that is, large positive profits from a given strategy in a single round will not radically modify the investor's strategy profile, and (b) forgetful, that is, because of random mutation, strategies that are selected because of large profits in past rounds tend to decay unless reinforced by further experience. If the firm issues a pure state security, the security will generate a large payoff in a single state and zero payoff in all other states. Thus, buying the security, at any price, will generate losses most of the time. Because the regularity of these losses and the conservatism of the selection mechanism put considerable downward pressure on investor bids, underpricing obtains. Moreover, the forgetfulness of the adaptive learners ensures that the impact of an infrequent large payoff in a single state is not sufficient to counter the downward pressure on pure security prices. Extending this logic to other security designs shows that one can minimize underpricing by issuing securities that, even when correctly valued, confer gains upon the purchasing

investors most of the time. Such claims are the mirror image of pure securities, requiring both the bundling of cash flows across states and the splitting of cash flows from a single state.

Given the learning dynamic for investors, the payoffs to the firm from choosing different financing strategies are governed by the following trade-off. On the one hand, narrow securities with payoffs in only a few states maximize the profits associated with heterogeneous investor valuations. This is especially important if investors place radically different marginal valuations on cash flows in different states of the world. On the other hand, securities that pay off in most states and isolate losses to only a few states minimize underpricing. When the average variance of investors' marginal valuations across states is not too large, the underpricing effect dominates. Thus, optimal designs involve a bundle of securities, each member of which produces losses in only a few states of the world. Because the firm earns cash flows in all states of the world, this pattern of losses can be produced only by splitting cash flows from individual states across multiple securities. Thus, the dominant claim design with adaptive learning features both bundling and splitting of cash flows.

These results relate to a large body of emerging research that explores behavioral deviations from rational expectations pricing.¹ Like most of this research, our analysis produces asset prices that lie outside the rational expectations framework. However, our work is distinct in two respects. First, we do not endow our agents with any particular bias based on psychological considerations. Instead, our agents start with random initial strategies and learn from experience. In this sense our work most closely resembles that of Arifovic (1996), who simulates exchange rate determination, Arifovic (1994), and Chen and Yeh (1996), who simulate the evolution of market clearing demand and supply functions, Allen and Karjalainen (1999), who simulate the evolution of stock trading rules, and Routledge (2001), who simulates information acquisition in financial markets. Second, in contrast to this literature and also to behavioral finance literature in general, we do not take the matrix of market securities as given and simply analyze the investors' demand-side problem. Rather, we examine the corporate supply side as well. In this respect, our work resembles the study of security issuance by corporations by Noe et al. (2003). However, Noe et al. model investors as risk-neutral agents with identical marginal valuations. Thus, in Noe et al. security markets have no role for providing investor diversification, rather, their only role is to fund corporate growth. In the present paper, we abstract from the real financing needs of firms and focus on the surplus that firms can extract from investors based on variation and heterogeneity in their valuations across states.

The remainder of paper is organized as follows. Section I describes the market environment modeled in this paper. In particular, Section I presents the genetic algorithm we employ to model firm and investor learning. Section II provides a brief description of the central results. We discuss the baseline results in Section III. In Section IV, we examine the robustness of our results. Finally, we

¹ For a recent review of this literature see Hirshleifer (2001) and Lo (2002).

conclude the paper with a summation of our results and their implications in Section V.

I. The Capital Market

Consider a firm that is designing its capital structure. The firm operates in a single-period setting with three possible end-of-period states of the world, which are indexed by $j \in \{1, 2, 3\}$. The firm generates a cash flow of \$1 in each of the three states and sells claims on this cash flow to investors. Each investor has state-dependent preferences (i.e., investors' marginal valuations of cash flows depend on the state of the world). This scenario is consistent with risk-averse investors buying claims that are an infinitesimal fraction of their portfolio of holdings (see Allen and Gale (1988)).² Type- j investors prefer payoffs in state j to payoffs in the other two states, thus, they value \$1 received in state j at par and they discount payoffs received in each of the other two states by $\alpha < 1$. This structure of preferences could arise, for instance, in an economy in which type- j investors have a higher marginal utility for income in state j because of a shock to their nonmarketable labor income in this state. The distribution of investor types is uniform. Note that, under this specification, the average investor valuation of income is constant across states. Moreover, the firm's income is constant across states. Thus, in this scenario, the potential welfare improvement from security design rests not in sharing macroeconomic risk, but rather from insuring against the idiosyncratic risks that individual investors face. Admittedly, this structure of preferences is rather stylized. However, the very symmetry and simplicity of the structure are what make our result easy to interpret. Through a number of robustness checks that we discuss below, we verify whether our results are qualitatively dependent on these assumptions.

The firm has to select a financing strategy, that is, it has to issue a set of securities that pays out all its future cash flows to investors. The firm can choose from a menu of 10 securities. Three of these securities are pure securities, which we denote as S_j , where $j = 1, 2, 3$. Security S_j pays \$1 if state j obtains and \$0 if the other states obtain. Three of these securities are portfolios of two pure securities each, or "two-state" securities, which we denote as S_{ij} , where $ij \in \{12, 23, 13\}$. Security S_{ij} pays \$1 if either state i or state j obtains and \$0 in the third state. Three of these securities are "split securities," that is, securities that pay off only \$0.5 in two states and \$0 in the third state, which we denote as S_{ij}^* , where again $ij \in \{12, 23, 13\}$. The final security, S_{123} , is a riskless bond that pays \$1 in all three states.

Given the menu of securities that the firm can issue and the requirement that all cash flows be paid out to investors, the firm can choose one of six financing strategies. The first strategy, the state-independent strategy, ST_I , calls for the issuance of the riskless bond, S_{123} . The second strategy, which we refer to as the state-dependent strategy, ST_D , requires that the firm issue the three pure

² Other frameworks could also give rise to this preference structure; for example, state-dependent utility, or variations in marginal valuation caused by tax effects.

securities, S_1 , S_2 , and S_3 . The next three strategies, which we label hybrid strategies, ST_{Hj} , where $j = \{1, 2, 3\}$, result in the issuance of two securities each, namely, a pure security that pays off only in state j and a two-state security that pays off in the other two states. For example, the firm issues securities S_1 and S_{23} if it adopts the hybrid strategy ST_{H1} . The final strategy, ST_S , calls for the firm to issue the following set of three split two-state securities: $\{S_{12}^*, S_{23}^*, S_{13}^*\}$. This strategy differs from the others because the firm's cash flow is split across two securities in every state, and thus it is the only strategy whereby multiple securities provide investors with positive payoffs in every state. We refer to this strategy as the "splitting strategy."

Investors respond to the firm's security offering by engaging in Bertrand-type competition to acquire securities. That is, investors bid separately for each security issued by the firm and each security in the package is then sold to the highest bidder for that security at the price offered by that bidder. Ties are broken by splitting the security into equal shares and assigning one share to each of the high bidders; that is, all high bidders for a security share equally in the payment for the security and its payoff.

We specify the setting above such that we achieve a number of objectives. First, we select the preference structure to ensure that in the perfect market, rational agent environment of Allen and Gale (1988), pure state-dependent securities are optimal. Next, we specify the number of bidders so that at least two bidders with identical preferences are available to bid on every possible claim. Finally, we select the menu of securities to ensure that the options for the firm include both bundling and splitting of cash flows across states. This design allows us to compare equilibrium security designs under adaptive learning with equilibrium designs under rational expectations while holding all other market characteristics constant.

II. Genetic Algorithm

Genetic algorithms were first developed by Holland (1962, 1975). A genetic algorithm is an optimization technique based on the principles of natural selection. A population of solution candidates is evaluated according to their "fitness," that is how much of a benefit they generate. Next, the solution candidates are recombined through selection, crossover, and mutation to generate new solution candidates, or "offspring." Candidates with higher fitness measures are more likely to reproduce offspring and keep their "genes" alive. The new generation of solution candidates are evaluated and they in turn reproduce their offspring following the same selection process. This evolution process stops whenever certain terminating criteria are met.

The success of genetic techniques as an optimization tool is well documented. Genetic methods conduct a parallel search of many points in the solution space. This leads to an optimal search in the stochastic environment, balancing between a promising search direction and a less-visited direction. Indeed, this parallelism makes the genetic method an efficient optimization technique for a number of complex problems. For instance, several economists use genetic

algorithms to study the learning dynamics of agents in different game settings, and find that agents using genetic algorithms are able to find the optimal strategy in these games.³

We simulate the behavior of agents in the economy we describe in the previous section by playing a “virtual” game between a virtual firm and virtual investors. A genetic algorithm determines agent actions. Each run of the genetic-algorithm simulation consists of a number of rounds. Figure 1 illustrates the sequence of events for each round. In each round, the virtual agents select strategies and receive payoffs. First, the firm selects a security issuance strategy from the menu of six available strategies. Let the strategies $ST_I, ST_D, ST_{H1}, ST_{H2}, ST_{H3}$, and ST_S be represented by the numbers 0, 1, 2, 3, 4, and 5, respectively. Each of the firm’s 80 chromosomes encodes one of the six possible financing strategies. In each round, one chromosome is drawn at random from the pool of 80 chromosomes, and the firm plays the strategy encoded by the selected chromosome.

At the beginning of each run, the composition of the firm’s chromosome pool is uniform, with each strategy having an equal probability of being included in the pool.⁴ After each round, the fitness of each of the strategies is evaluated based on the profitability of the strategy in the last round in which it is used, where profitability is measured by the proceeds from the security sale. The chromosome pool is then updated via a process of selection and mutation. Selection takes the form of a lottery in which chromosomes are selected at random to form a new pool of chromosomes. The probability that a chromosome will be selected into the new pool is proportional to the profitability of the chromosome. More specifically, after each round, the probability that chromosome n is selected into the new pool is given by

$$\frac{\pi_n^F}{\sum_m \pi_m^F},$$

where π_n^F represents the profitability of chromosome n . This process continues until the new pool consists of 80 chromosomes. After selection, each chromosome is subject to possible mutation. The probability of mutation for each chromosome in each round is 0.001. Mutation results in a chromosome’s being replaced by a new mutant chromosome. The mutant chromosome is equally likely to be encoded with each one of the six possible strategies.

For each security, there is a distinct investor pool. This pool contains six investors, each of which has the preference configuration that places the highest

³ Genetic techniques are used in many different areas, such as aircraft ground traffic optimization, automobile design, protein structure prediction, and so on (see Spector and Goodman (2001)).

⁴ Under this setup, chromosomes that encode hybrid strategies, ST_{Hj} , where $j = 1, 2, 3$, are likely to make up a disproportionately large fraction of the firm’s initial chromosome pool. This introduces a bias against the system’s converging to the other financing strategies. The qualitative nature of our results does not change in unreported simulations that are designed to control for this bias by adjusting the mechanism for selecting the firm’s original chromosome pool.




Figure 1. Time line of genetic algorithm learning in each round.

value on that security. Investor preferences vary as follows: For each of the pure securities, S_j , every one of the six investors is of type j . For securities that pay off in two states, S_{ij} and S_{ij}^* , there are three type- j and three type- i investors. Finally, for the security S_{123} , there are two investors of each type $j = 1, 2, 3$. This design ensures Bertrand competition, which, under rational expectations,

would exhaust the investor surplus from security purchases. By varying the investor pool in this manner we ensure that the market size is constant across all securities, as there is always a constant number of investors who have a high valuation for a security who are bidding for the security. In the absence of this market size assumption, the market size for pure state securities would be much smaller than for other claims, in which case we would observe even less issuance of pure state securities.⁵

Each investor makes a single bid for each security issued by the firm. Each bid is determined by the investor's pricing strategy, where this pricing strategy for each security is drawn at random from the strategy pool associated with that security. For each of the 10 securities that the firm can issue, each of the six bidding investors has a strategy pool that consists of 80 chromosomes. Each chromosome is encoded with a strategy. Strategies are binary strings of 0s and 1s that correspond to a 10-place, base-2 representation of a number, n , between 0 and 1,023. For example, the string 0000000010 represents the number 2. If the strategy encoded by the selected chromosome decodes to the number n , then the bid of the investor is $A \frac{n}{1,023}$. The coefficient A , which regulates the tick size and the range of bid values available to investors, is set to one. The security is awarded to the investor that submits the highest bid, with ties being broken by splitting both the payment for the security and the security's payoffs among the highest-bidding investors. The price paid for each security equals its winning bid.⁶

At the beginning of each run of the genetic algorithm, the 80 chromosomes in each pricing pool are randomly encoded according to a uniform distribution over strategies. After each round, chromosome pools are updated. This process involves selection, crossover, mutation, and election. Selection, which we perform first, involves estimation of the profitability of each strategy. Strategy k 's estimated profitability for investor i is the investor's profit from using that strategy, assuming that other investors continue to choose the strategies they selected in the previous round. We transform investors' profits by adding one to the profit of each strategy. Because investor profits are bounded from below by negative one, the transformed profits, π_k^i , are always positive. Following this transformation of investor profits, we use a lottery mechanism similar to that employed in the firm's selection process to generate a new pool of strategies by

⁵ One alternative is to keep the investor pool constant regardless of the security issued with equal numbers of the three types of investors. This biases the results against the state-dependent strategy, ST_D , because there are only one-third as many investors with a high valuation bidding for a security that pays off in only a single state as there are bidding for security S_{123} . We report simulations that employ this specification of the investor pool later in the paper and demonstrate that the qualitative nature of our results is unaffected by the change in the composition of the investor pool.

⁶ Other auction mechanisms could be employed to simulate the security sale. For instance, we modify the genetic algorithm and employ a second-price auction to determine the price paid by the winning bidder. While this change affects the winning bids, it does not alter the qualitative results on security choice that we present in this paper.

randomly selecting chromosomes from the old pool. The probability of a chromosome k being selected into the new chromosome pool for investor i is given by

$$\frac{\pi_k^i}{\sum_l \pi_l^i}.$$

This process continues until the new pool consists of the same number of chromosomes as the old one.

The new chromosome pool above is then subject to crossover, mutation, and election. First, two chromosomes are randomly selected from the pool, and, with a probability of 0.6, they are crossed.⁷ A crossing of two chromosomes involves splitting each chromosome into two parts at a randomly selected position and then swapping the parts between the pair of chromosomes. For example, suppose that two selected chromosomes are encoded as 0000010101 and 1000111000, respectively. If the randomly selected position for splitting the chromosomes is before the third digit, two new chromosomes are generated such that the chromosome encoded by 0000010101 is recoded as 0000011000, and the chromosome encoded by 1000111000 is recoded as 1000110101. The two new chromosomes are then subject to mutation with probability 0.003. Mutation involves replacing a chromosome with a new chromosome that is randomly encoded according to a uniform probability measure over the 1,024 permitted strategies. The new chromosomes that result from crossover and mutation, together with the two original chromosomes, are then subject to election; that is, only the two chromosomes that encode the most profitable strategies are added to the new chromosome pool. With each iteration of this process of crossover, mutation, and election, two new chromosomes are added to the new pool. The process continues until the new pool contains 80 chromosomes, with the completion of the new chromosome pool marking the end of a round. A run of the genetic algorithm starts with the initial encoding of chromosome pools for the firm and the investors. A run ends when 10,000 rounds have elapsed. Each experiment consists of 100 runs.

III. Results of the Baseline Simulations

In this section, we characterize the results of simulations that implement the algorithm we specify above. Our strategy for analyzing the results is as follows. First, we address the central concern of this paper, namely, which security designs are optimal in the presence of two-sided evolutionary learning? Next, we try to determine why these designs are optimal by considering both the outcomes from one-sided learning and the pricing dynamic for securities.

⁷ Note that only chromosomes representing investor strategies are crossed. Because we permit only six different strategies for the firm, using crossover to generate novel strategies for the firm is not necessary.

Table I
Frequency Summary

This table presents the frequency with which the firm adopts each of six financing strategies available to it in the final round of three sets of simulations. The table also presents the sum of frequencies for the three generically similar hybrid strategies. Each set of simulations consists of 100 runs and every simulation lasts for 10,000 rounds. The intensity of investor state preferences, α , varies across the three sets of simulations.

Strategy	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 0.9$
State-dependent	38	1	0
Hybrid (ST_{H1})	11	13	3
Hybrid (ST_{H2})	8	12	1
Hybrid (ST_{H3})	12	9	1
All hybrid (ST_H)	31	34	5
Splitting	24	65	95
State-independent	7	0	0

A. Equilibrium Strategies

We begin by examining the firm's strategies that result from the simulations. We vary the degree of investor's state preferences by varying the level of α . Lower values of α reduce the value of securities that pay off in multiple states. Table I presents the frequency with which the firm adopts each of the six available financing strategies in the final round of the simulations. A clear pattern is evident in the firm's strategy choices. When $\alpha = 0.1$, that is, when state-dependent preferences are most marked, the state-dependent financing strategy, ST_D , is most frequently employed. At the other extreme, when $\alpha = 0.9$, that is, when the differences in the values that investors place on cash flows across states are at their lowest, strategy ST_S , which splits the cash flows in every state between two securities, dominates. This strategy also dominates in the intermediate case in which $\alpha = 0.5$.

These results establish that, even when investors have state-dependent preferences, state-dependent financing strategies may not characterize equilibria in a market populated by adaptive agents. Rather, state-dependent financing strategies are adopted in equilibrium only when the state dependence of investors' preferences is high. When these state-dependent preferences are moderate, equilibrium financing strategies call for the issuance of securities that (i) pay off in multiple states (i.e., cash flow bundling) and (ii) split cash flows from a single state across multiple securities. The remainder of this paper explains these results and demonstrates their robustness to a variety of modifications in the genetic algorithm.

B. One-Sided Learning

In this subsection, we investigate the underlying forces that produce the results we describe above. First we present simulations that examine the

Table II
Outcomes of Simulations with Fixed Investor Responses

This table presents the frequency with which the firm adopts each of the six financing strategies available to it in the final round of three sets of simulations. The table also presents the sum of frequencies for the three generically similar hybrid strategies. Each set of simulations consists of 100 runs and every simulation lasts for 10,000 rounds. Investor bids for each security issued by the firm are fixed at their intrinsic values for the duration of each simulation. The intensity of investor state preferences, α , varies across the three sets of simulations that employ a mutation rate of 0.001 for the firm.

Strategy	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 0.9$
State-dependent	100	98	48
Hybrid (ST_{H1})	0	1	17
Hybrid (ST_{H2})	0	1	13
Hybrid (ST_{H3})	0	0	11
All hybrid (ST_H)	0	2	41
Splitting	0	0	4
State-independent	0	0	7

outcomes that obtain when only one set of agents, investors, or firms, learns. We start by describing outcomes when only the firm learns, with investor bids fixed at the intrinsic values for the securities; we then examine outcomes when the firm's issuance strategy is fixed but investors learn to price securities through experience.

Table II contains the outcomes of simulations in which only the firm's financing strategy evolves. Investor responses to the issuance of each security are constrained to equal the security's intrinsic value. The table presents the frequencies with which each of the six financing strategies available to the firm is employed in the final round of three sets of simulations. Each set of 100-run simulations employs a different level of α . The state-dependent strategy dominates in simulations in which α is set equal to 0.1 and 0.5. For $\alpha = 0.9$, while the state-dependent strategy remained the most frequently employed strategy, it is no longer dominant. The decrease in the usage of the state-dependent strategy when investors' state preferences are weak is understandable given that, when $\alpha = 0.9$, there is little difference in the value of the firm's financing alternatives. What emerges from these simulations is that, when investors price securities at their no-arbitrage rational expectations prices, the firm gravitates to the perfect markets rational agent solution suggested by Allen and Gale (1988) (i.e., the firm employs state-dependent financing).

In Table III we present the outcomes of simulations in which the firm continually employs the same financing strategy for 10,000 rounds while investors learn to price. The table presents outcomes for four securities—one of each generic type.⁸ Several regularities are apparent. First, security prices are increasing in expected cash flows and decreasing in the strength of investors'

⁸ That is, one security is used to represent each permutation-equivalence class of securities.

Table III
Outcomes of Simulations Fixing Firm Issue Strategies

This table presents outcomes of the final rounds of three sets of simulations. Each set of simulations consists of 100 runs and every simulation lasts for 10,000 rounds. The firm's financing strategies are fixed for the duration of each simulation. The intensity of investor state preferences, α , varies across the three sets of simulations that employ crossover and mutation rates for investors of 0.6 and 0.003, respectively. For each of the four generic classes of securities that we employ in our paper, the table presents an intrinsic value (THEV), the average of winning bids (WIN), the average value of chromosome pools across investors (AVEBID), undervaluation of the winning bid relative to intrinsic value (UWIN), and bid shading as captured by the difference between the intrinsic value and chromosome pool average scaled by the intrinsic value (UAVE).

Security	THEV	WIN	AVEBID	UWIN	UAVE
Panel A: $\alpha = 0.1$					
Pure	0.333	0.127	0.066	0.619	0.801
Two-state	0.367	0.141	0.082	0.617	0.776
Split two-state	0.183	0.070	0.040	0.617	0.779
Riskless	0.400	0.145	0.088	0.638	0.781
Panel B: $\alpha = 0.5$					
Pure	0.333	0.127	0.066	0.618	0.802
Two-state	0.500	0.482	0.334	0.035	0.331
Split two-state	0.250	0.245	0.168	0.021	0.326
Riskless	0.667	0.507	0.442	0.239	0.337
Panel C: $\alpha = 0.9$					
Pure	0.333	0.125	0.069	0.625	0.793
Two-state	0.633	0.725	0.472	-0.145	0.254
Split two-state	0.317	0.412	0.272	-0.302	0.140
Riskless	0.933	0.898	0.799	0.037	0.144

state preferences. Second, as the strength of state preferences diminishes, underpricing diminishes, and diminishes at a faster rate for the two-state security and the split two-state security. Finally, as in Noe et al. (2003), winning bids exceed the average bid in investors' chromosome pools.

Prices display systematic deviations from no-arbitrage prices as a result of the learning dynamic. The workings of the dynamic are transparent in the case of pure securities. The basis for the dynamic is the asymmetric effect of losses and gains. First consider losses, which, for a pure state security, occur two-thirds of the time. When losses occur, bidding less than the market-clearing bid generates the highest payoff to the buyer. All bids that meet this criterion produce the maximum profit of zero and thus are equally likely to be selected into the chromosome pool. It follows that when losses occur (two-thirds of the time), updating of the chromosome pool uniformly increases the proportion of strategies that feature bids between zero and the market-clearing bid. Next consider gains, which, for a pure state security, occur one-third of the time.

When gains occur, bids that beat the winning bid by just one tick generate maximum profit. Thus, updating is asymmetric as losses, which occur two-thirds of the time, uniformly increase the use of bids from zero to the winning bid, while gains, which occur one-third of the time, increase bids very close to but slightly above the winning bid. Thus, when pure securities are fairly priced based on intrinsic value, there is excess selection pressure in the direction of lowering bids, as losses are frequent. This asymmetric dynamic results in underpricing of pure securities by about 60%.

The pricing of the two-state and split two-state securities is governed by a similar dynamic. When state preferences are extreme, that is, $\alpha = 0.1$, adding cash flow in another state provides investors little benefit, in which case two-state securities and split two-state securities behave like pure securities and result in significant underpricing. When $\alpha = 0.5$, state preferences are less extreme. Investors reap a large gain from the cash flow to the security in the second state, and the strength of the asymmetric downward pressure on prices that we discuss above is proportional to the probability of losses. Thus, the amount of underpricing required to counter the losses on the two-state security is less than that for the one-state security, which pushes the price of the two-state security close to investor valuations of cash flows in the added state. Hence, the two-state and the split two-state securities are subject to less severe underpricing.

Now consider the riskless security. Interestingly, the addition of the cash flow from the final state produces only a marginal increase in the price of the security. This follows because the learning dynamic ensures that prices can never be low enough to permit profits in every state. Because investors have low valuations for cash flows in two of the three states, it must be the case that they realize losses from holding the bond in two of the three states. However, because of the asymmetric nature of the dynamic, losses cannot be large in these states, otherwise there would be too much pressure for downward revision. Thus, the price settles close to the investor valuation for the least preferred states, α . This effect ensures that the price of the riskless security is close to the price of the two-state security. This pricing similarity, combined with the higher intrinsic value of the three-state claim, results in larger underpricing for the three-state security.

C. Pricing of Securities with Two-Sided Learning

We now describe the pricing of securities when both investors and the firm learn. Table IV describes the final-round outcomes of three sets of simulations in which both firms and investors learn over the course of 10,000 rounds in each run of the simulations. The intensity of investors' state preferences varies across the three simulations.

Consistent with the results from the one-sided learning simulations described in Table II, there appears to exist a positive correlation between the pricing of security packages and their usage by the firm. When investors' preferences across states are relatively extreme ($\alpha = 0.1$), financing strategy ST_D is the

Table IV
Outcomes of Simulations with Two-Sided Learning

This table presents outcomes of the final rounds of three sets of simulations. Each panel of the table presents outcomes for a set of simulations that employ one calibration of the intensity of investor state preferences, α . Each set of simulations consists of 100 runs and every simulation lasts for 10,000 rounds. Both investors and firms learn during the simulations. The simulations employ crossover and mutation rates for investors of 0.6 and 0.003, respectively, and a mutation rate for the firm of 0.001. For each of the six financing strategies available to the firm, the table presents the frequency with which the strategy is adopted by the firm in the final round (FREQ), the intrinsic value of the strategy (STTV), average firm proceeds in the last round from the strategy (FGAIN), and the undervaluation of the strategy relative to its theoretical value (UVAL). For each security issued under each financing strategy, the table presents its intrinsic value (THEV), the average of winning bids (WIN), the average value of chromosome pools across investors (AVEBID), undervaluation of the winning bid relative to intrinsic value (UWIN), and bid shading as captured by the difference between the intrinsic value and chromosome pool average scaled by the intrinsic value (UAVE). The table also contains averages of this information across the three generically similar hybrid strategies.

Strategy	FREQ	STTV	FGAIN	UVAL	Security	THEV	WIN	AVEBID	UWIN	UAVE
Panel A: $\alpha = 0.1$										
State-dependent	38	1.000	0.394	0.606	S_1	0.333	0.122	0.067	0.634	0.798
					S_2	0.333	0.125	0.067	0.625	0.800
					S_3	0.333	0.147	0.069	0.559	0.792
Hybrid (ST_{H1})	11	0.700	0.331	0.528	S_1	0.333	0.146	0.067	0.561	0.798
					S_{23}	0.367	0.184	0.108	0.497	0.704
Hybrid (ST_{H2})	8	0.700	0.375	0.465	S_2	0.333	0.130	0.067	0.609	0.800
					S_{13}	0.367	0.244	0.108	0.334	0.706
Hybrid (ST_{H3})	12	0.700	0.351	0.498	S_3	0.333	0.118	0.069	0.647	0.792
					S_{12}	0.367	0.234	0.118	0.363	0.679
All hybrid (ST_H)	31	0.700	0.350	0.500	S_i	0.333	0.131	0.068	0.607	0.797
					S_{jk}	0.367	0.219	0.111	0.403	0.696
Splitting	24	0.550	0.378	0.312	S_{12*}	0.183	0.109	0.064	0.407	0.652
					S_{13*}	0.183	0.134	0.062	0.268	0.661
					S_{23*}	0.183	0.135	0.063	0.262	0.658
State-independent	7	0.400	0.266	0.334	S_{123}	0.400	0.266	0.157	0.334	0.606
Panel B: $\alpha = 0.5$										
State-dependent	1	1.000	0.927	0.073	S_1	0.333	0.252	0.088	0.243	0.737
					S_2	0.333	0.299	0.087	0.103	0.738
					S_3	0.333	0.375	0.089	-0.126	0.732
Hybrid (ST_{H1})	13	0.833	0.634	0.239	S_1	0.333	0.140	0.088	0.580	0.737
					S_{23}	0.500	0.494	0.291	0.012	0.418
Hybrid (ST_{H2})	12	0.833	0.627	0.247	S_2	0.333	0.155	0.087	0.536	0.738
					S_{13}	0.500	0.472	0.278	0.055	0.444
Hybrid (ST_{H3})	9	0.833	0.674	0.191	S_3	0.333	0.160	0.089	0.520	0.732
					S_{12}	0.500	0.514	0.283	-0.027	0.433
All hybrid (ST_H)	34	0.833	0.642	0.229	S_i	0.333	0.151	0.088	0.548	0.736
					S_{jk}	0.500	0.492	0.284	0.017	0.432
Splitting	65	0.750	0.714	0.048	S_{12*}	0.250	0.236	0.151	0.056	0.397
					S_{13*}	0.250	0.240	0.153	0.041	0.386
					S_{23*}	0.250	0.238	0.154	0.046	0.385
State-independent	0	0.667			S_{123}	0.667		0.319		0.521

(continued)

Table IV—Continued

Strategy	FREQ	STTV	FGAIN	UVAL	Security	THEV	WIN	AVEBID	UWIN	UAVE
Panel C: $\alpha = 0.9$										
State-dependent	0	1.000			S_1	0.333		0.129		0.614
					S_2	0.333		0.132		0.605
					S_3	0.333		0.132		0.605
Hybrid (ST_{H1})	3	0.967	0.992	−0.026	S_1	0.333	0.278	0.129	0.167	0.614
Hybrid (ST_{H2})	1	0.967	0.837	0.134	S_{23}	0.633	0.714	0.423	−0.128	0.333
					S_2	0.333	0.258	0.132	0.226	0.605
Hybrid (ST_{H3})	1	0.967	0.944	0.023	S_{13}	0.633	0.579	0.423	0.086	0.332
					S_3	0.333	0.252	0.132	0.243	0.605
All hybrid (ST_H)	5	0.967	0.951	0.016	S_{12}	0.633	0.692	0.450	−0.093	0.290
					S_i	0.333	0.269	0.131	0.194	0.608
Splitting	95	0.950	1.223	−0.287	S_{jk}	0.633	0.683	0.432	−0.078	0.318
					S_{12*}	0.317	0.410	0.265	−0.296	0.163
					S_{13*}	0.317	0.405	0.255	−0.279	0.193
State-independent	0	0.933			S_{23*}	0.317	0.408	0.271	−0.287	0.144
					S_{123}	0.933		0.651		0.303

most frequently adopted strategy by the firm and strategy ST_D also raises the most capital. When α rises to 0.5, the firm most frequently employs the splitting strategy ST_S , which generates the most positive investor response. The dominance of strategy ST_S is almost complete when α rises further to 0.9. In this case, strategy ST_S raises more than its intrinsic value.

A comparison of the winning bids under one-sided learning (Table III) and two-sided learning (Table IV) provides further insights. First, the winning bids for the most frequently issued securities are similar in the two cases. For example, in simulations in which $\alpha = 0.5$, the split two-state securities realize an average winning bid of 0.245 in the one-sided learning simulations and 0.238 in the two-sided learning simulations. This is not surprising, given that pricing in these simulations is a function of investor experience. That is, as is the case with the one-sided learning simulations, investors are likely to gain considerable experience in pricing the most frequently issued securities in the two-sided learning simulations. The second regularity that emerges from a comparison of the winning bids across the two cases relates to the pricing of the remaining securities. Specifically, these securities tend to command higher prices in the two-sided learning simulations. We believe this difference arises because, in the two-sided simulations, securities tend to be issued in the final round only if they receive favorable investor reception in earlier rounds. In contrast, in the one-sided simulations, the firm must issue securities even following poor investor response.

The results in Table IV highlight the fundamental insight of this paper. Both the intrinsic value and the level of underpricing vary across the menu of financing strategies. Mispricing tends to be greatest for the state-dependent strategy, ST_D , which has the highest intrinsic value. Conversely, even though the splitting strategy, ST_S , has an intrinsic value that exceeds only that of the state-independent financing strategy ST_I , it tends to be associated with the least amount of underpricing.

The difference between the intrinsic values of the state-dependent package and the other financing strategies is greatest when $\alpha = 0.1$. At the same time, the level of underpricing is relatively uniform across securities. Thus, the system tends to converge to state-dependent financing as predicted by Allen and Gale (1988). As the intensity of investor state preferences moderates, so does the variation in the intrinsic values of the financing strategies. This dilutes the advantage of state-dependent financing, as does a reduction in the level of mispricing of securities issued as part of the other financing strategies. Thus, once α rises to 0.5, state-dependent financing is no longer dominant.

The dominance of the splitting strategy over the state-independent financing strategy for $\alpha = 0.5$ and $\alpha = 0.9$ arises for similar reasons. First, the splitting strategy commands a higher intrinsic value because it ensures that a higher proportion of firm cash flows is directed toward investors who place the highest value on the cash flows. Further, the securities included in the package tend to be subject to less underpricing. The strategy of splitting cash flows dominates the hybrid strategies despite commanding a lower intrinsic value. The disadvantage of the hybrid strategies stems from the inclusion of a pure security in the financing mix. While the pure security secures the hybrid strategy a higher intrinsic value, its mispricing disadvantage more than outweighs this valuation advantage.

The simulations we report in Table III are stopped after 10,000 rounds. As Young (1996) demonstrates, in a system with mutation-provided background noise, convergence in shorter simulations may not represent the limiting behavior of the algorithm over the very long run. Table V presents outcomes of simulations run to capture the very long-run limiting behavior of our genetic algorithm. Each of these simulations lasts one million rounds. A comparison with the results from Table IV reveals no change in the relative rankings of financing strategies. However, the dominance of the dominant securities intensifies.

IV. Robustness of the Results

In this section we examine the robustness of the results above to changes in the selection mechanism that we employ in the genetic algorithm. First, we describe outcomes that correspond to rank-based selection. Next, we detail the effects of switching to a lottery selection mechanism that uses a nonlinear transformation of profits. In both instances we find that the changes do not alter the qualitative nature of our results. Then we extend the number of states from 3 to 10 and find that optimal security designs feature the bunching of losses on a single state of the world. Finally, we show that our results do not change in an economy with a constant population of investors.

A. Rank-Based Selection

To implement rank-based selection, after each round, we rank strategies on the basis of their profitability in the round. Ties are broken by randomizing across the tied strategies. We replace the four lowest-ranked (least profitable) strategies in investor chromosome pools with the four highest-ranked (most

Table V
Outcomes of Simulations with Two-Sided Learning, the Very Long Run

This table presents outcomes of the final rounds of three sets of simulations. Each panel of the table presents outcomes for a set of simulations that employ one calibration of the intensity of investor state preferences, α . Each set of simulations consists of 100 runs and every simulation lasts for 1,000,000 rounds. Both investors and firms learn during the simulations. The simulations employ crossover and mutation rates for investors of 0.6 and 0.003, respectively, and a mutation rate for the firm of 0.001. For each of the six financing strategies available to the firm, the table presents the frequency with which the strategy is adopted by the firm in the final round (FREQ), the intrinsic value of the strategy (STTV), average firm proceeds in the last round from the strategy (FGAIN), and the undervaluation of the strategy relative to its theoretical value (UVAL). For each security issued under each financing strategy, the table presents its intrinsic value (THEV), the average of winning bids (WIN), the average value of chromosome pools across investors (AVEBID), undervaluation of the winning bid relative to intrinsic value (UWIN), and bid shading as captured by the difference between the intrinsic value and chromosome pool average scaled by the intrinsic value (UAVE). The table also contains averages of this information across the three generically similar hybrid strategies.

Strategy	FREQ	STTV	FGAIN	UVAL	Security	THEV	WIN	AVEBID	UWIN	UAVE
Panel A: $\alpha = 0.1$										
State-dependent	84	1.000	0.356	0.644	S_1	0.333	0.110	0.058	0.669	0.827
					S_2	0.333	0.127	0.060	0.619	0.819
					S_3	0.333	0.119	0.060	0.643	0.820
Hybrid (ST_{H1})	3	0.700	0.268	0.616	S_1	0.333	0.098	0.058	0.705	0.827
					S_{23}	0.367	0.170	0.064	0.536	0.825
Hybrid (ST_{H2})	6	0.700	0.179	0.745	S_2	0.333	0.077	0.060	0.770	0.819
					S_{13}	0.367	0.102	0.064	0.722	0.824
Hybrid (ST_{H3})	7	0.700	0.207	0.704	S_3	0.333	0.092	0.060	0.725	0.820
					S_{12}	0.367	0.115	0.065	0.685	0.823
All hybrid (ST_H)	16	0.700	0.208	0.703	S_i	0.333	0.087	0.059	0.738	0.822
					S_{jk}	0.367	0.121	0.064	0.671	0.824
					S_{12*}	0.183		0.033		0.823
Splitting	0	0.550			S_{13*}	0.183		0.033		0.821
					S_{23*}	0.183		0.032		0.827
					S_{123}	0.400		0.073		0.817
Panel B: $\alpha = 0.5$										
State-dependent	0	1.000			S_1	0.333		0.053		0.840
					S_2	0.333		0.056		0.832
					S_3	0.333		0.056		0.833
Hybrid (ST_{H1})	1	0.833	0.647	0.223	S_1	0.333	0.151	0.053	0.548	0.840
					S_{23}	0.500	0.497	0.312	0.007	0.376
Hybrid (ST_{H2})	1	0.833	0.602	0.277	S_2	0.333	0.143	0.056	0.572	0.832
					S_{13}	0.500	0.459	0.317	0.081	0.366
Hybrid (ST_{H3})	1	0.833	0.573	0.313	S_3	0.333	0.080	0.056	0.760	0.833
					S_{12}	0.500	0.493	0.321	0.015	0.359
All hybrid (ST_H)	3	0.833	0.607	0.271	S_i	0.333	0.124	0.055	0.627	0.835
					S_{jk}	0.500	0.483	0.317	0.034	0.367
					S_{12*}	0.250	0.244	0.184	0.024	0.264
Splitting	97	0.750	0.736	0.019	S_{13*}	0.250	0.246	0.184	0.015	0.263
					S_{23*}	0.250	0.245	0.174	0.019	0.303
					S_{123}	0.667		0.384		0.424

(continued)

Table V—Continued

Strategy	FREQ	STTV	FGAIN	UVAL	Security	THEV	WIN	AVEBID	UWIN	UAVE
Panel C: $\alpha = 0.9$										
State-dependent	0	1.000			S_1	0.333		0.070		0.789
					S_2	0.333		0.070		0.791
					S_3	0.333		0.076		0.772
Hybrid (ST_{H1})	0	0.967			S_1	0.333		0.070		0.789
					S_{23}	0.633		0.434		0.315
Hybrid (ST_{H2})	0	0.967			S_2	0.333		0.070		0.791
					S_{13}	0.633		0.451		0.288
Hybrid (ST_{H3})	0	0.967			S_3	0.333		0.076		0.772
					S_{12}	0.633		0.443		0.300
All hybrid (ST_H)	0	0.967			S_i	0.333		0.072		0.784
					S_{jk}	0.633		0.443		0.301
					S_{12*}	0.317	0.403	0.265	-0.272	0.164
Splitting	98	0.950	1.213	-0.277	S_{13*}	0.317	0.408	0.265	-0.288	0.163
					S_{23*}	0.317	0.403	0.260	-0.271	0.180
State-independent	2	0.933	0.899	0.036	S_{123}	0.933	0.899	0.707	0.036	0.242

profitable) strategies.⁹ Because the firm chromosome pools contain fewer chromosome types, the pools are updated by replacing only the least profitable chromosome with the most profitable one. Following the replacement, the remainder of the process for generating new chromosome pools proceeds as in Section III above.

Table VI presents the outcomes of simulations in which only investors learn how to bid for securities while the firm employs the same strategy for 10,000 rounds. A comparison of resulting security prices with those from the original one-sided simulations, presented in Table III, indicates that pure security prices are uniformly lower following the switch to rank-based selection. In contrast, with the exception of simulations in which $\alpha = 0.1$, the remaining securities command uniformly higher prices given rank-based selection. Further, the modification in the selection process also reduces the difference between the prices of the two-state securities and the risk-free security.

There is a relatively straightforward explanation for these changes. As is the case with the learning dynamic that we describe earlier, when a security pays \$1 and investors value this cash flow at par, selection tends to replace the strategies below the previous winning bid with the four most profitable strategies. Note that these four strategies are the ones that exceed the winning bid by the least amount. When a security generates a loss for investors, rank-based selection replaces the strategies that exceed the winning bid by the largest margin with strategies that are lower than the winning bid. This occurs either when a security pays zero or security prices exceed investors' valuations of cash flows in their least preferred states. For the system to stabilize, it must be the

⁹ In unreported simulations, we obtain qualitatively similar results when the number of strategies replaced is two or eight.

Table VI
Outcomes of Simulations Fixing Firm Issue Strategies
with Rank-Based Selection

This table presents outcomes of the final rounds of three sets of simulations. Each set of simulations consists of 100 runs and every simulation lasts for 10,000 rounds. The firm's financing strategies are fixed for the duration of each simulation. The intensity of investor state preferences, α , varies across the three sets of simulations which employ crossover and mutation rates for investors of 0.6 and 0.003, respectively. The simulations also employ rank-based selection where the four best chromosomes for each investor replace the four worst chromosomes. For each of the four generic classes of securities we employ in our paper, the table presents an intrinsic value (THEV), the average of winning bids (WIN), the average value of chromosome pools across investors (AVEBID), undervaluation of the winning bid relative to intrinsic value (UWIN), and bid shading as captured by the difference between the intrinsic value and chromosome pool average scaled by the intrinsic value (UAVE).

Security	THEV	WIN	AVEBID	UWIN	UAVE
Panel A: $\alpha = 0.1$					
Pure	0.333	0.048	0.026	0.855	0.921
Two-state	0.367	0.111	0.077	0.698	0.790
Split two-state	0.183	0.056	0.036	0.693	0.805
Riskless	0.400	0.112	0.084	0.720	0.789
Panel B: $\alpha = 0.5$					
Pure	0.333	0.059	0.025	0.822	0.924
Two-state	0.500	0.499	0.377	0.001	0.246
Split two-state	0.250	0.248	0.197	0.010	0.212
Riskless	0.667	0.515	0.468	0.227	0.298
Panel C: $\alpha = 0.9$					
Pure	0.333	0.054	0.026	0.839	0.921
Two-state	0.633	0.893	0.621	-0.410	0.020
Split two-state	0.317	0.446	0.326	-0.410	-0.028
Riskless	0.933	0.899	0.877	0.036	0.061

case that the probability of losses and the probability of gains are similar. For two-state (split two-state) securities, this implies that winning bids must be roughly equal to investors' valuations of \$1 (\$0.5) in their less preferred states, α (0.5α). Similarly, in the case of the risk-free bond, this balance between the probabilities of gains and losses also approaches a price that approximates α . Thus, the two-state securities and the risk-free bond command similar prices despite the differences in their intrinsic values. The poor pricing of the pure securities follows because an approximate balance between the probability of gains and the probability of losses can be achieved only at prices close to zero.

In Table VII, we present the outcomes of two-sided learning simulations that employ rank-based selection. For $\alpha = 0.5$ and $\alpha = 0.9$, the simulation results are qualitatively similar to the results from Table IV above, that is, the splitting

Table VII
Outcomes of Simulations with Two-Sided Learning
with Rank-Based Selection

This table presents outcomes of the final rounds of three sets of simulations. Each panel of the table presents outcomes for a set of simulations that employ one calibration of the intensity of investor state preferences, α . Each set of simulations consists of 100 runs and every simulation lasts for 10,000 rounds. Both investors and firms learn during the simulations. The simulations employ crossover and mutation rates for investors of 0.6 and 0.003, respectively, and a mutation rate for the firm of 0.001. The simulations also employ rank-based selection where the four best chromosomes for each investor replace the four worst chromosomes; the best chromosome replaces the worst chromosome for the firm. For each of the six financing strategies available to the firm, the table presents the frequency with which the strategy is adopted by the firm in the final round (FREQ), the intrinsic value of the strategy (STTV), average firm proceeds in the last round from the strategy (FGAIN), and the undervaluation of the strategy relative to its theoretical value (UVAL). For each security issued under each financing strategy, the table presents its intrinsic value (THEV), the average of winning bids (WIN), the average value of chromosome pools across investors (AVEBID), undervaluation of the winning bid relative to intrinsic value (UWIN), and bid shading as captured by the difference between the intrinsic value and chromosome pool average scaled by the intrinsic value (UAVE). The table also contains averages of this information across the three generically similar hybrid strategies.

Strategy	FREQ	STTV	FGAIN	UVAL	Security	THEV	WIN	AVEBID	UWIN	UAVE
Panel A: $\alpha = 0.1$										
State-dependent	13	1.000	0.226	0.774	S_1	0.333	0.063	0.031	0.811	0.906
					S_2	0.333	0.073	0.029	0.780	0.912
					S_3	0.333	0.090	0.034	0.731	0.898
Hybrid (ST_{H1})	17	0.700	0.193	0.725	S_1	0.333	0.081	0.031	0.756	0.906
Hybrid (ST_{H2})	9	0.700	0.200	0.715	S_{23}	0.367	0.111	0.061	0.696	0.834
					S_2	0.333	0.079	0.029	0.763	0.912
Hybrid (ST_{H3})	21	0.700	0.223	0.682	S_{13}	0.367	0.121	0.058	0.671	0.843
					S_3	0.333	0.096	0.034	0.712	0.898
All hybrid (ST_H)	47	0.700	0.208	0.704	S_{12}	0.367	0.127	0.064	0.654	0.825
					S_i	0.333	0.087	0.032	0.738	0.905
Splitting	32	0.550	0.194	0.647	S_{jk}	0.367	0.120	0.061	0.673	0.834
					S_{12*}	0.183	0.066	0.035	0.641	0.811
					S_{13*}	0.183	0.072	0.032	0.610	0.823
State-independent	8	0.400	0.160	0.600	S_{23*}	0.183	0.057	0.033	0.690	0.821
					S_{123}	0.400	0.160	0.080	0.600	0.799
Panel B: $\alpha = 0.5$										
State-dependent	0	1.000			S_1	0.333		0.072		0.783
					S_2	0.333		0.073		0.782
					S_3	0.333		0.072		0.785
Hybrid (ST_{H1})	1	0.833	0.542	0.350	S_1	0.333	0.043	0.072	0.871	0.783
Hybrid (ST_{H2})	2	0.833	0.604	0.275	S_{23}	0.500	0.499	0.303	0.003	0.393
					S_2	0.333	0.106	0.073	0.683	0.782
Hybrid (ST_{H3})	1	0.833	0.563	0.324	S_{13}	0.500	0.499	0.300	0.003	0.400
					S_3	0.333	0.065	0.072	0.806	0.785
All hybrid (ST_H)	4	0.833	0.578	0.306	S_{12}	0.500	0.499	0.307	0.003	0.387
					S_i	0.333	0.080	0.072	0.761	0.783
Splitting	96	0.750	0.740	0.013	S_{jk}	0.500	0.499	0.303	0.003	0.393
					S_{12*}	0.250	0.246	0.182	0.017	0.273
					S_{13*}	0.250	0.246	0.189	0.016	0.242
State-independent	0	0.667			S_{23*}	0.250	0.249	0.185	0.006	0.260
					S_{123}	0.667		0.337		0.494

(continued)

Table VII—Continued

Strategy	FREQ	STTV	FGAIN	UVAL	Security	THEV	WIN	AVEBID	UWIN	UAVE
Panel C: $\alpha = 0.9$										
State-dependent	0	1.000			S_1	0.333		0.112		0.663
					S_2	0.333		0.125		0.626
					S_3	0.333		0.124		0.627
Hybrid (ST_{H1})	0	0.967			S_1	0.333		0.112		0.663
					S_{23}	0.633		0.620		0.022
Hybrid (ST_{H2})	0	0.967			S_2	0.333		0.125		0.626
					S_{13}	0.633		0.603		0.047
Hybrid (ST_{H3})	0	0.967			S_3	0.333		0.124		0.627
					S_{12}	0.633		0.583		0.080
All hybrid (ST_H)	0	0.967			S_i	0.333		0.120		0.640
					S_{jk}	0.633		0.602		0.049
Splitting	100	0.950	1.340	−0.410	S_{12*}	0.317	0.446	0.329	−0.410	−0.038
					S_{13*}	0.317	0.447	0.328	−0.412	−0.035
					S_{23*}	0.317	0.446	0.327	−0.409	−0.034
State-independent	0	0.933			S_{123}	0.933		0.590		0.368

strategy dominates. Once again, this dominance arises because the set of securities issued as part of this strategy is priced more favorably than the other securities, and the superior pricing of this package is large enough to offset the disadvantage of a relatively low intrinsic value. The results from simulations in which $\alpha = 0.1$ differ from their counterparts in Table IV. With rank-based selection, hybrid strategies tend to dominate. The decrease in the usage of the state-dependent strategy follows from the relatively poor pricing of pure securities under rank-based selection.

B. Lottery Selection with a Nonlinear Transformation of Profits

In this section we describe outcomes from simulations in which the genetic algorithm employs a lottery mechanism based on a nonlinear transformation of profits to select chromosomes. To create a new pool of chromosomes for firms, selection begins with the computation of weights for each chromosome. To compute the weights, w_{ST} , for each financing strategy, we first transform profits by taking the exponential of the product of the profit from strategy ST, π_{ST} , and a coefficient, β_f , that is,

$$w_{ST} = e^{\beta_f \pi_{ST}}.$$

The coefficient β_f is set equal to two.¹⁰ According to the weights of each chromosome, we generate a new pool of chromosomes by randomly selecting chromosomes from the old pool. The probability that a chromosome n will be selected into the new pool is given by

¹⁰ In unreported simulations, we obtain qualitatively similar results when β_f and its investor counterpart, β_I , are both set to one or four.

$$\frac{w_n}{\sum_m w_m}.$$

This process continues until the new pool consists of the same number of chromosomes as the old one. We use the same nonlinear transformation of profits and the same lottery mechanism to create a new pool of chromosomes for investors. With the exception of the mechanism for the updating of chromosome pools, the genetic algorithm proceeds as in Section III.

Table VIII presents the outcomes of 10,000-round simulations in which only investors learn how to bid for securities while the firm employs the same strategy. Security prices are higher than in the original one-sided simulations we present in Table III. The price increase is most marked for pure securities.

Table VIII
Outcomes of Simulations Fixing Firm Strategies and Employing
a Nonlinear Transformation of Profits in Selection

This table presents outcomes of the final rounds of three sets of simulations. Each set of simulations consists of 100 runs and every simulation lasts for 10,000 rounds. The firm’s financing strategies are fixed for the duration of each simulation. The intensity of investor state preferences, α , varies across the three sets of simulations which employ crossover and mutation rates for investors of 0.6 and 0.003, respectively. The simulations also employ lottery-based selection where the probability of selection for each chromosome is dependent on a nonlinear transformation of its profit. For each of the four generic classes of securities that we employ in our paper, the table presents an intrinsic value (THEV), the average of winning bids (WIN), the average value of chromosome pools across investors (AVEBID), undervaluation of the winning bid relative to intrinsic value (UWIN), and bid shading as captured by the difference between the intrinsic value and chromosome pool average scaled by the intrinsic value (UAVE).

Security	THEV	WIN	AVEBID	UWIN	UAVE
Panel A: $\alpha = 0.1$					
Pure	0.333	0.250	0.147	0.251	0.560
Two-state	0.367	0.251	0.151	0.316	0.587
Split two-state	0.183	0.090	0.054	0.508	0.705
Riskless	0.400	0.275	0.163	0.313	0.591
Panel B: $\alpha = 0.5$					
Pure	0.333	0.241	0.139	0.278	0.582
Two-state	0.500	0.488	0.336	0.023	0.328
Split two-state	0.250	0.245	0.173	0.020	0.310
Riskless	0.667	0.503	0.393	0.246	0.411
Panel C: $\alpha = 0.9$					
Pure	0.333	0.235	0.140	0.295	0.579
Two-state	0.633	0.762	0.505	−0.203	0.203
Split two-state	0.317	0.403	0.272	−0.272	0.141
Riskless	0.933	0.899	0.792	0.037	0.152

However, for simulations in which $\alpha = 0.5$ and $\alpha = 0.9$, two-state and split two-state securities continue to enjoy the lowest underpricing. Consequently, and once again, the differences between the prices of two-state securities and the risk-free security are smaller than the differences between their intrinsic values. The explanation for these results is similar to that for the base-case lottery-selection simulations we describe in Section III. The nonlinear transformation of profits ensures that the probabilities of selection are convex in profits. Thus, the reward for a large positive profit is larger than the penalty for a large negative profit. This encourages higher bids by investors.

In Table IX, we present the outcomes of two-sided learning simulations. Once again, for simulations that employ $\alpha = 0.5$ and $\alpha = 0.9$, the results are qualitatively similar to the results from Table IV above, that is, the splitting strategy dominates. The forces underlying the dominance of the splitting strategy are identical to those we describe earlier. For simulations in which $\alpha = 0.1$, the dominance of the state-dependent strategy strengthens. This follows because of the improved pricing of pure securities following the change in the selection mechanism.

C. Extending the Number of States

In the previous sections, we observe that security designs that concentrate losses in one of the three states of nature dominate. This result raises the following question: Is it generally optimal to issue claims that feature a low probability of losses to the buyer? To answer this question we need to consider a richer state space. We therefore extend our analysis to a setting that features 10 states of nature instead of three. To focus on the issue of loss distribution, we consider 10 financing strategies for the firm, which we index from 1 to 10. Each financing strategy calls for the issuance of a set of 10 securities. Each security issued in financing strategy n pays $1/n$ in n states. For example security set 9 features 10 securities. Security 1 pays $1/9$ in states 1 through 9 and zero in state 10, security 2 pays zero in state 1 and $1/9$ in all other states, and so on. This design ensures that all cash flows are paid out.

To maintain a market structure that is comparable to the one we employ in the previous sections, 20 investors bid for each security. For each state j in which a security generates a positive payoff, the investor pool contains at least two investors of type j , where type- j investors value \$1 at par in state j and at α in every other state. For example, for the first security in strategy 9 the investor pool includes two investors each of type j , where $j \in \{1, 2, \dots, 9\}$, and two investors are selected at random from types 1 through 9. For the second security issued in strategy 9, the investor pools include two investors each of type j , where $j \in \{2, 3, \dots, 10\}$, and two investors are selected at random from types 2 through 10. With the exception of the change in the set of firm strategies and the composition of the investor pool, the rest of the structure of the genetic algorithm is identical to that in Section III.

Because of the increase in the number of securities that must be priced and the increase in the size of the state space, each run of the modified genetic

Table IX
Outcomes of Simulations with Two-Sided Learning Employing
a Nonlinear Transformation of Profits in Selection

This table presents outcomes of the final rounds of three sets of simulations. Each panel of the table presents outcomes for a set of simulations employing one calibration of the intensity of investor state preferences, α . Each set of simulations consists of 100 runs and every simulation lasts for 10,000 rounds. Both investors and firms learn during the simulations. The simulations employ crossover and mutation rates for investors of 0.6 and 0.003, respectively, and a mutation rate for the firm of 0.001. The simulations also employ lottery-based selection where the probability of selection for each chromosome is dependent on a nonlinear transformation of its profit. For each of the six financing strategies available to the firm, the table presents the frequency with which the strategy is adopted by the firm in the final round (FREQ), the intrinsic value of the strategy (STTV), average firm proceeds in the last round from the strategy (FGAIN), and the undervaluation of the strategy relative to its theoretical value (UVAL). For each security issued under each financing strategy, the table presents its intrinsic value (THEV), the average of winning bids (WIN), the average value of chromosome pools across investors (AVEBID), undervaluation of the winning bid relative to intrinsic value (UWIN), and bid shading as captured by the difference between the intrinsic value and chromosome pool average scaled by the intrinsic value (UAVE). The table also contains averages of this information across the three generically similar hybrid strategies.

Strategy	FREQ	STTV	FGAIN	UVAL	Security	THEV	WIN	AVEBID	UWIN	UAVE
Panel A: $\alpha = 0.1$										
State-dependent	85	1.000	0.745	0.255	S_1	0.333	0.242	0.131	0.273	0.606
					S_2	0.333	0.258	0.132	0.227	0.605
					S_3	0.333	0.245	0.138	0.265	0.587
Hybrid (ST_{H1})	6	0.700	0.538	0.232	S_1	0.333	0.276	0.131	0.172	0.606
					S_{23}	0.367	0.262	0.153	0.286	0.581
Hybrid (ST_{H2})	4	0.700	0.504	0.280	S_2	0.333	0.221	0.132	0.337	0.605
					S_{13}	0.367	0.283	0.153	0.228	0.583
Hybrid (ST_{H3})	3	0.700	0.522	0.254	S_3	0.333	0.196	0.138	0.413	0.587
					S_{12}	0.367	0.326	0.148	0.110	0.596
All hybrid (ST_H)	13	0.700	0.524	0.252	S_i	0.333	0.240	0.134	0.279	0.599
					S_{jk}	0.367	0.283	0.151	0.228	0.587
Splitting	2	0.550	0.500	0.092	S_{12*}	0.183	0.130	0.092	0.291	0.500
					S_{13*}	0.183	0.196	0.095	-0.066	0.484
					S_{23*}	0.183	0.174	0.087	0.051	0.526
State-independent	0	0.400			S_{123}	0.400		0.235		0.413
Panel B: $\alpha = 0.5$										
State-dependent	1	1.000	0.719	0.281	S_1	0.333	0.250	0.108	0.249	0.677
					S_2	0.333	0.250	0.110	0.249	0.669
					S_3	0.333	0.219	0.113	0.343	0.660
Hybrid (ST_{H1})	10	0.833	0.735	0.118	S_1	0.333	0.240	0.108	0.279	0.677
					S_{23}	0.500	0.494	0.276	0.011	0.448
Hybrid (ST_{H2})	19	0.833	0.729	0.125	S_2	0.333	0.238	0.110	0.286	0.669
					S_{13}	0.500	0.491	0.274	0.018	0.451
Hybrid (ST_{H3})	16	0.833	0.767	0.080	S_3	0.333	0.253	0.113	0.240	0.660
					S_{12}	0.500	0.513	0.275	-0.027	0.449
All hybrid (ST_H)	45	0.833	0.744	0.108	S_i	0.333	0.244	0.110	0.268	0.669
					S_{jk}	0.500	0.500	0.275	0.001	0.449
Splitting	54	0.750	0.723	0.036	S_{12*}	0.250	0.241	0.152	0.037	0.392
					S_{13*}	0.250	0.240	0.152	0.039	0.394
					S_{23*}	0.250	0.242	0.150	0.033	0.400
State-independent	0	0.667			S_{123}	0.667		0.325		0.512

(continued)

Table IX—Continued

Strategy	FREQ	STTV	FGAIN	UVAL	Security	THEV	WIN	AVEBID	UWIN	UAVE
Panel C: $\alpha = 0.9$										
State-dependent	0	1.000			S_1	0.333		0.144		0.568
					S_2	0.333		0.146		0.562
					S_3	0.333		0.153		0.540
Hybrid (ST_{H1})	4	0.967	1.035	−0.071	S_1	0.333	0.266	0.144	0.201	0.568
Hybrid (ST_{H2})	3	0.967	0.957	0.010	S_{23}	0.633	0.769	0.442	−0.214	0.302
					S_2	0.333	0.235	0.146	0.296	0.562
Hybrid (ST_{H3})	2	0.967	1.005	−0.040	S_{13}	0.633	0.722	0.431	−0.140	0.320
					S_3	0.333	0.169	0.153	0.493	0.540
All hybrid (ST_{H3})	9	0.967	1.002	−0.037	S_{12}	0.633	0.836	0.434	−0.320	0.315
					S_i	0.333	0.234	0.148	0.297	0.556
Splitting	89	0.950	1.201	−0.264	S_{jk}	0.633	0.768	0.435	−0.213	0.313
					S_{12*}	0.317	0.397	0.265	−0.255	0.164
					S_{13*}	0.317	0.403	0.259	−0.271	0.182
State-independent	2	0.933	0.899	0.036	S_{23*}	0.317	0.401	0.259	−0.267	0.182
					S_{123}	0.933	0.899	0.623	0.036	0.333

algorithm lasts 100,000 rounds. Table X reports the outcomes of the final rounds of these runs. As the table illustrates, when state preferences are marked ($\alpha = 0.1$), the firm employs all the strategies at its disposal but displays a weak preference for strategies in which securities tend to pay off in the majority of states. When $\alpha = 0.5$ or $\alpha = 0.9$, the firm primarily chooses strategy 9, which results in the issuance of 10 securities, each of which produces positive payoffs in nine states. As is the case with the outcomes we report earlier, there appears to be little difference in the pricing of riskless securities issued under strategy 10 and the securities issued under strategy 9. Further, underbidding, as captured by the difference between investor chromosome pool averages and the intrinsic value of securities, increases with the number of states in which a security makes a payoff of zero. This confirms our earlier conjecture that underpricing increases with an increase in the downward pressure on security prices that results from an increased frequency of investor losses from purchasing the security.

Our results support the hypothesis that optimal security designs feature the bunching of losses on a single state of the world. Such securities, even when correctly valued, produce gains for the purchasing investors most of the time. These securities are the mirror image of pure securities, as they pay off most of the time but generate large losses in a few states of the world. The optimality of these claims seems to follow from the trade-off we describe earlier. First, underpricing is reduced by reducing the number of states in which the security produces losses. However, because investors' state valuations are not identical, reducing riskiness reduces state variation in payoffs, which in turn can lower the marginal value of the securities being issued. Optimal designs therefore trade off better valuations given state-dependent payoffs against adverse learning effects from infrequent payoffs. The interaction of these two effects

Table X
Outcomes of Simulations with Two-Sided Learning When There
Are 10 States

This table presents outcomes of the final rounds of three sets of simulations. Each panel of the table presents outcomes for a set of simulations that employ one calibration of the intensity of investor state preferences, α . Each set of simulations consists of 100 runs and every simulation lasts for 100,000 rounds. Both investors and firms learn during the simulations. The simulations employ crossover and mutation rates for investors of 0.6 and 0.003, respectively, and a mutation rate for the firm of 0.001. The firm chooses from 10 financing strategies. For Strategy j , the firm issues 10 securities with each security paying $1/j$ in j states and 0 in other states. For each of the 10 financing strategies, the table presents the frequency with which the strategy is adopted by the firm in the final round (FREQ), the intrinsic value of the strategy (STTV), average firm proceeds in the last round from the strategy (FGAIN), and the undervaluation of the strategy relative to its theoretical value (UVAL). For securities issued under each financing strategy, the table presents its intrinsic value (THEV), the average of winning bids (WIN), the average value of chromosome pools across investors (AVEBID), undervaluation of the winning bid relative to intrinsic value (UWIN), and bid shading as captured by the difference between the intrinsic value and chromosome pool average scaled by the intrinsic value (UAVE).

Strategy	FREQ	STTV	FGAIN	UVAL	THEV	WIN	AVEBID	UWIN	UAVE
Panel A: $\alpha = 0.1$									
1	7	1.000	0.142	0.858	0.100	0.014	0.007	0.858	0.931
2	6	0.550	0.149	0.730	0.055	0.015	0.007	0.730	0.866
3	12	0.400	0.149	0.629	0.040	0.015	0.007	0.629	0.819
4	7	0.325	0.142	0.564	0.033	0.014	0.007	0.564	0.783
5	5	0.280	0.155	0.445	0.028	0.016	0.007	0.445	0.743
6	7	0.250	0.152	0.393	0.025	0.015	0.007	0.393	0.715
7	13	0.229	0.141	0.384	0.023	0.014	0.007	0.384	0.683
8	13	0.213	0.142	0.330	0.021	0.014	0.007	0.330	0.668
9	18	0.200	0.155	0.227	0.020	0.015	0.007	0.227	0.650
10	12	0.190	0.153	0.195	0.019	0.015	0.007	0.195	0.613
Panel B: $\alpha = 0.5$									
1	0	1.000			0.100		0.038		0.624
2	1	0.750	0.433	0.423	0.075	0.043	0.038	0.423	0.499
3	0	0.667			0.067		0.037		0.438
4	0	0.625			0.063		0.036		0.419
5	0	0.600			0.060		0.037		0.378
6	1	0.583	0.429	0.264	0.058	0.043	0.037	0.264	0.360
7	0	0.571			0.057		0.037		0.344
8	0	0.563			0.056		0.037		0.340
9	94	0.556	0.547	0.015	0.056	0.055	0.037	0.015	0.333
10	4	0.550	0.499	0.094	0.055	0.050	0.037	0.094	0.324
Panel C: $\alpha = 0.9$									
1	0	1.000			0.100		0.054		0.461
2	0	0.950			0.095		0.055		0.421
3	0	0.933			0.093		0.055		0.411
4	0	0.925			0.093		0.055		0.400
5	0	0.920			0.092		0.055		0.404
6	0	0.917			0.092		0.056		0.389
7	0	0.914			0.091		0.054		0.410
8	0	0.913			0.091		0.054		0.409
9	98	0.911	0.997	-0.094	0.091	0.100	0.054	-0.094	0.403
10	2	0.910	0.899	0.012	0.091	0.090	0.055	0.012	0.399

leads to the optimality of the securities that offer constant payoffs across all states save one, in which case they pay nothing.

D. Security Design with a Constant Population of Investors

We now examine the robustness of our results to a change in the structure of the investor pool. In all our simulations thus far, each security has its own investor pool, where investor pool is tailored to ensure that it only contains investor types that have the highest preference for payoffs in the state in which the security generates positive payoffs. The advantage of this design is that it equates the competitiveness of the markets for each of the securities, thereby isolating the effects of learning on security choice. However, it is more realistic to assume that the investor pool is invariant. Thus, we now turn to results from simulations that employ investor pools that do not vary with the identity of the securities being issued by the firm.

In Table XI, we present results from simulations in which only investors learn to price securities while the firm's issuance strategy is held constant. The investor pool consists of six investors. Two investors have a preference for payoffs in each of the three possible states. A comparison of the results in Table XI with those from similar simulations in which the investor pool varies with the identity of the security being issued (Table III) highlights the role of the composition of the investor pool. When the investor pool is held constant, prices for pure securities are 53% lower when $\alpha = 0.1$ and 31% lower when $\alpha = 0.5$. These price changes are statistically significant. The two-state and split two-state securities incur much lower price declines of 23% and 16% when $\alpha = 0.1$. These price changes are also statistically significant. However, the prices of these two securities are not statistically different when $\alpha = 0.5$. When state preferences are slight ($\alpha = 0.9$), there is no significant change in the price of any security. We attribute this change in prices to a decrease in the level of competitiveness of the securities market, as proxied by the number of investors with a preference for payoffs in states in which securities generate positive payoffs, given the investor pool is held constant. From Table XI, as we increase the preference for income in each investor's most favored state, the investors with the highest preference for income in the given state are more likely to win the auction for a security that pays off in that state. In the extreme case in which $\alpha = 0.1$, there are effectively only two investors bidding for pure securities when the investor pool is held constant, compared with our earlier simulations in which there are effectively six investors bidding for each pure security. The reduction in the level of competitiveness is less marked in the case of either the two-state or the split two-state securities. Consequently, these securities' prices do not fall as steeply when the investor pool is held constant. The pricing of the risk-free security is unaffected as the composition of the investor pool in our earlier simulations is identical to that in the simulations we report in Table XI.

In Table XII we assess how a change in the investor pool can affect the design of securities when both firms and investors learn over the course of the

Table XI
Outcomes of Simulations Fixing Firm Issue Strategies
When the Investor Population Is Fixed

This table presents outcomes of the final rounds of three sets of simulations. Each set of simulations consists of 100 runs and every simulation lasts for 10,000 rounds. We fix the firm's financing strategies for the duration of each simulation. The intensity of investor state preferences, α , varies across the three sets of simulations which employ crossover and mutation rates for investors of 0.6 and 0.003, respectively. For each of the four generic classes of securities that we employ in our paper, the table presents an intrinsic value (THEV), the average of winning bids (WIN), the average value of chromosome pools across investors (AVEBID), undervaluation of the winning bid relative to intrinsic value (UWIN), bid shading as captured by the difference between the intrinsic value and chromosome pool average scaled by the intrinsic value (UAVE), and the percentage of auctions that are won by investors who have the highest preference for payoffs in one of the states in which the security generates positive payoffs (WINP).

S	THEV	WIN	AVEBID	UWIN	UAVE	WINP (%)
Panel A: $\alpha = 0.1$						
S_1	0.333	0.060 ^a	0.032	0.819	0.904	75
S_{23}	0.367	0.108 ^a	0.066	0.705	0.820	94
S_{23*}	0.183	0.059 ^a	0.035	0.677	0.812	93
S_{123}	0.400	0.149	0.087	0.629	0.783	100
Panel B: $\alpha = 0.5$						
S_1	0.333	0.088 ^a	0.046	0.736	0.863	60
S_{23}	0.500	0.481	0.328	0.038	0.343	82
S_{23*}	0.250	0.240	0.169	0.040	0.325	73
S_{123}	0.667	0.501	0.433	0.249	0.350	100
Panel C: $\alpha = 0.9$						
S_1	0.333	0.111	0.059	0.666	0.822	35
S_{23}	0.633	0.732	0.472	-0.156	0.254	70
S_{23*}	0.317	0.404	0.266	-0.276	0.161	74
S_{123}	0.933	0.899	0.787	0.036	0.157	100

^aThe average of winning bids is significantly less than the corresponding value in the base case (Table III) at the 5% level based on a *t*-test.

simulations. The results correspond to 1,000,000-round three-state simulations that are similar to those summarized earlier in Table V. The only difference is that the investor pool is held constant to generate the results reported in Table XII. For each of the three states, the pool includes two investors with a preference for payoffs in that state. The only perceptible change in the results occurs when state preferences are most marked ($\alpha = 0.1$). While the state-dependent strategy that calls for the issuance of only pure securities appears to dominate in our earlier simulations, in the case of a fixed investor pool the hybrid issuance strategies that call for the issuance of a two-state security and one pure security dominate. This result is consistent with the deterioration in the pricing of pure securities that we document in Table XI.

Table XII
Outcomes of Simulations with Two-Sided Learning in the Very Long Run
When the Investor Population Is Fixed

This table presents outcomes of the final rounds of three sets of simulations. Each panel of the table presents outcomes for a set of simulations that employ one calibration of the intensity of investor state preferences, α . Each set of simulations consists of 100 runs and every simulation lasts for 1,000,000 rounds. Both investors and firms learn during the simulations. The simulations employ crossover and mutation rates for investors of 0.6 and 0.003, respectively, and a mutation rate for the firm of 0.001. For each of the six financing strategies available to the firm, the table presents the frequency with which the strategy is adopted by the firm in the final round (FREQ), the intrinsic value of the strategy (STTV), average firm proceeds in the last round from the strategy (FGAIN), and the undervaluation of the strategy relative to its theoretical value (UVAL). For each security issued under each financing strategy, the table presents its intrinsic value (THEV), the average of winning bids (WIN), the average value of chromosome pools across investors (AVEBID), undervaluation of the winning bid relative to intrinsic value (UWIN), bid shading as captured by the difference between the intrinsic value and chromosome pool average scaled by the intrinsic value (UAVE), and the percentage of auctions that are won by investors who have the highest preference for payoffs in one of the states in which the security generates positive payoffs (WINP). The table also contains averages of this information across the three generically similar hybrid strategies.

Strategy	FREQ	STTV	FGAIN	UVAL	Security	THEV	WIN	AVEBID	UWIN	UAVE	WINP (%)
Panel A: $\alpha = 0.1$											
ST_D	5	1.000	0.116 ^a	0.884	S_1	0.333	0.040	0.016	0.880	0.951	90
					S_2	0.333	0.032	0.016	0.905	0.951	80
					S_3	0.333	0.044	0.017	0.867	0.949	100
ST_{H1}	18	0.700	0.157 ^a	0.776	S_1	0.333	0.036	0.016	0.893	0.951	92
					S_{23}	0.367	0.121	0.061	0.670	0.833	97
ST_{H2}	14	0.700	0.175	0.750	S_2	0.333	0.053	0.016	0.840	0.951	79
					S_{13}	0.367	0.122	0.061	0.668	0.834	100
ST_{H3}	14	0.700	0.181	0.741	S_3	0.333	0.049	0.017	0.854	0.949	64
					S_{12}	0.367	0.132	0.062	0.639	0.830	100
ST_H	46	0.700	0.170 ^a	0.758	S_i	0.333	0.045	0.017	0.865	0.950	79
					S_{jk}	0.367	0.125	0.061	0.660	0.832	99
ST_S	44	0.550	0.161	0.707	S_{23*}	0.183	0.055	0.032	0.700	0.825	91
					S_{13*}	0.183	0.054	0.032	0.705	0.823	95
					S_{12*}	0.183	0.052	0.031	0.717	0.828	89
ST_I	5	0.400	0.127	0.683	S_{123}	0.400	0.127	0.072	0.683	0.820	100
Panel B: $\alpha = 0.5$											
ST_D	0	1.000			S_1	0.333		0.042		0.873	
					S_2	0.333		0.046		0.861	
					S_3	0.333		0.044		0.867	
ST_{H1}	1	0.833	0.495	0.406	S_1	0.333	0.088	0.042	0.736	0.873	0
					S_{23}	0.500	0.407	0.321	0.187	0.359	100
ST_{H2}	0	0.833			S_2	0.333		0.046		0.861	
					S_{13}	0.500		0.312		0.377	
ST_{H3}	1	0.833	0.559	0.329	S_3	0.333	0.125	0.044	0.625	0.867	100
					S_{12}	0.500	0.434	0.299	0.132	0.401	100
ST_H	2	0.833	0.527 ^a	0.368	S_i	0.333	0.107	0.044	0.680	0.867	50
					S_{jk}	0.500	0.420	0.311	0.159	0.379	100
ST_S	98	0.750	0.738	0.017	S_{23*}	0.250	0.246	0.183	0.017	0.269	70
					S_{13*}	0.250	0.246	0.190	0.015	0.241	77
					S_{12*}	0.250	0.246	0.185	0.018	0.260	72
ST_I	0	0.667			S_{123}	0.667		0.389		0.416	

(continued)

Table XII—Continued

Strategy	FREQ	STTV	FGAIN	UVAL	Security	THEV	WIN	AVEBID	UWIN	UAVE	WINP (%)
Panel C: $\alpha = 0.9$											
ST_D	0	1.000			S_1	0.333		0.068		0.795	
					S_2	0.333		0.073		0.780	
					S_3	0.333		0.069		0.793	
ST_{H1}	0	0.967			S_1	0.333		0.068		0.795	
					S_{23}	0.633		0.413		0.348	
ST_{H2}	0	0.967			S_2	0.333		0.073		0.780	
					S_{13}	0.633		0.417		0.342	
ST_{H3}	0	0.967			S_3	0.333		0.069		0.793	
					S_{12}	0.633		0.416		0.343	
ST_H	0	0.967			S_i	0.333		0.070			
					S_{jk}	0.633		0.415			
ST_S	99	0.950	1.221	−0.285	S_{23*}	0.317	0.412	0.273	−0.302	0.139	72
					S_{13*}	0.317	0.407	0.268	−0.285	0.154	69
					S_{12*}	0.317	0.401	0.263	−0.267	0.170	71
ST_I	1	0.933	0.899	0.036	S_{123}	0.933	0.899	0.745	0.036	0.202	100

^aThe value is significantly different from the corresponding value in the base case (Table V) at the 5% level based on a t -test.

V. Conclusion

In this paper, we consider security design from an evolutionary learning perspective. Our main result is that securities with narrow investor appeal tend to be undervalued by financial markets. In turn, undervaluation encourages firms to issue claims that have broad appeal. These claims produce gains most of the time, isolating large losses to a few states. This result *prima facie* is consistent with stylized facts. Most corporate bonds are associated with small but positive default probabilities and the stocks of most firms earn positive returns most of the time. Moreover, while firms could raise capital by issuing a very large number of claims such that each pays off only in a few select events, they usually do not follow this path. We can rationalize this regularity in actual corporate security designs without resorting to transactions or marketing costs of security issuance. This rationalization is appealing because both issuance and marketing costs display considerable variation cross-sectionally and over time. Hence, if these costs are the only impediments to firms raising capital with narrow payoff securities, one would expect that, as costs decrease, firms would move toward financing based on such narrowly tailored securities. Such movement is not apparent.

While we view this result as an important first step to developing an adaptive learning theory of security design, our analysis is both preliminary and exploratory and more work is required to achieve a thorough understanding of the key issues. For instance, one issue that merits exploration is the linearity of the firm's payoff in funds raised. This assumption, which follows Allen and Gale (1988), is consistent with a firm's issuing securities solely to capture the economic surplus from outside investors. However, securities are frequently

issued to finance projects and these projects generate returns that are nonlinear in the level of investment. In the extreme case, a project may have a fixed minimum threshold of investment. In this situation, firms face a considerable risk of losing investment opportunities if the price response of outside investors to a given claim is variable. Hence, such firms are concerned with both the expected proceeds of an issue as well as the variance of proceeds. This concern for downside risk plays an important part in the dominance of debt finance that obtains in the adaptive learning analysis of Noe, et al. (2003). An interesting question is whether this effect would still be important in the current setting, which features a much richer specification of investor preferences, or whether it would be swamped by the underpricing versus investor surplus dynamic that drives our analysis of security design. Only by incorporating nonlinearity into the firm's objective function can we answer this question.

Another important issue that merits consideration is the effect of the number of bidders in the security market. Under Bertrand competition with rational expectations, two identical agents are sufficient to force prices to their competitive level. However, in our adaptive setting, the number of agents with identical preferences that bid for a given claim will affect pricing. On the one hand, with many agents bidding, the winning bid is much more likely to be the product of mutation in the strategy pool and to therefore be suboptimally high. On the other hand, because bidding sufficiently high to top such mutant bids will frequently result in large losses, adaptation may select in very conservative bidding strategies. Only by varying the number of bidders can we determine which of these two effects dominates.

Our analysis can also be extended by incorporating aggregate risk into the simulated economy. In the analysis here, the total payoff on all of the securities is constant across states. By allowing variation in total payoffs it would be possible to identify "good" and "bad" states of nature and a more realistic pricing kernel for claims under rational expectations. The pricing kernel produced by learning could then be compared to the rational expectations pricing kernel, and we could try to explain problems such as the equity premium puzzle.

REFERENCES

- Allen, Franklin, and Douglas Gale, 1988, Optimal security design, *Review of Financial Studies* 1, 229–263.
- Allen, Franklin, and Risto Karjalainen, 1999, Using genetic algorithms to find technical trading rules, *Journal of Financial Economics* 51, 245–271.
- Arifovic, Jasmina, 1994, Genetic algorithm learning and the cobweb model, *Journal of Economic Dynamics and Control* 18, 3–28.
- Arifovic, Jasmina, 1996, The behavior of the exchange rate in the genetic algorithm and experimental economies, *Journal of Political Economy* 104, 510–541.
- Chen, Shu-Heng, and Chia-Hsuan Yeh, 1996, On the coordination and adaptability of the large economy: An application of genetic programming to the cobweb model, in P. Angeline and K. Kinnear Jr., eds. *Advances in Genetic Programming II* (MIT Press, Cambridge, MA).
- Gale, Douglas, 1992, Standard securities, *Review of Economic Studies* 59, 731–755.
- Hirshleifer, David, 2001, Investor psychology and asset pricing, *Journal of Finance* 56, 1533–1598.

- Holland, John H., 1962, Outline for a logical theory of adaptive systems, *Journal of the Association for Computing Machinery* 3, 297–314.
- Holland, John H., 1975, *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (Ann Arbor, MI: University of Michigan Press).
- Lo, Andrew, 2002, Bubble, rubble, finance in trouble? *Journal of Psychology and Financial Markets* 3, 76–86.
- Madan, Dilip, and Badih Soubra, 1991, Marketing of financial products, *Review of Financial Studies* 4, 362–384.
- Noe, Thomas, Michael Rebello, and Jun Wang, 2003, Corporate financing: An artificial agent-based analysis, *Journal of Finance* 63, 943–973.
- Routledge, Bryan, 2001, Genetic algorithm learning to choose and use information, *Macroeconomic Dynamics* 5, 303–325.
- Spector, Lee, and Erik Goodman, eds. 2001, *GECCO 2001: Proceedings of the Genetic and Evolutionary Computation Conference* (Morgan Kaufmann, San Fransisco, CA).
- von Hayek, Friedrich, 1945, The use of knowledge in society, *American Economic Review* 35, 519–530.
- Young, Peyton, 1996, The economics of convention, *Journal of Economic Perspectives* 10, 105–122.
- Young, Peyton, 1998, *Individual Strategy and Social Structure* (Princeton University Press Princeton, NJ).

Copyright of Journal of Finance is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.