
FM Lecture Series

Game Theory in Finance

Anjan V. Thakor

Anjan Thakor is the INB National Bank Professor of Finance at Indiana University

■ The purpose of this paper is to provide an overview of game theory, particularly as it relates to finance. My objective is to introduce the subject, so I will be illustrative rather than rigorous and complete. Since it would be impossible for me to review in this paper all of the potential applications of game theory in finance, I will focus on a subset of applications of game theory and offer a few observations on possible trends. For much of this paper, I'll work through a simple example to illustrate how the game-theoretic treatment of a "standard" asymmetric information problem differs from the approaches that were popular in finance prior to the surge of recent game-theoretic models.

This paper is an outgrowth of a talk prepared as part of a tutorial for the FMA meeting in Orlando, Florida, October 1990. I gratefully acknowledge helpful comments on an earlier draft by James Ang (the editor), Mark Bagnoli, Arnoud Boot, Jess Beltz, Michael Brennan, Sugato Chakravarty, Dan Indro, George Kanatas, Rene Stulz, and Patty Wilson. They are, of course, absolved of responsibility for any remaining infelicities.

It is generally believed that game theory began with the publication of von Neumann and Morgenstern's book, *The Theory of Games and Economic Behaviour*, in 1944. However, until the 1970s, game theory was not even an integral part of mainstream economics, let alone finance. It is not a coincidence that game theory became more widely adopted in economics in the 1970s, which is the time that Akerlof's [1] path-breaking work on adverse selection marked the beginning of enormous interest in the economics of information. Economists began to realize that limitations on the information possessed by individuals were important in understanding economic behavior, in part because such limitations induced people to alter their behavior. The resulting strategic interactions had potentially profound implications.

During the 1960s and the early part of the 1970s, finance was preoccupied with the notion that markets are efficient and that the details of institutional design are relatively unimportant in understanding the

functioning of markets. Thus, game theory and information economics did not have much appeal to finance. However, the work of Leland and Pyle [38], Ross [52], and Bhattacharya [7] brought information economics into the arena of finance research. These early models were attempts to break out of the domain of irrelevance results provided by the Modigliani and Miller [43, 44] work on capital structure and dividends. These papers were followed by numerous papers in the early 1980s which explained a variety of price reactions, institutional details, and contract features. This theoretical work dovetailed nicely with the voluminous empirical literature on event studies which suggested that corporate insiders such as managers have proprietary information not reflected in prices, and that prices do react to certain actions undertaken by them.

This early body of work in finance involving information economics did not have very much to do with game theory. The models belonged to a class I shall refer to as "Signalling Models with the Uninformed Moving First" (SMUF). I'll explain later what I mean by this. However, these models had several limitations which I'll discuss a little later. As we approached the end of the last decade, finance theorists began to turn increasingly to game-theoretic signalling models to avoid some of the nettlesome problems in SMUF. I shall refer to these game-theoretic models as "Signalling Models with the Informed Moving First" (SMIF). These models now occupy center stage in research involving signalling issues in finance.

What then is game theory? Simply put, it is a study of situations in which we make some primitive assumptions about the agents involved in a particular interaction (or game) and then try to figure out what happens when each agent acts to maximize his own expected utility subject to the constraints imposed by his information (and beliefs), endowments, and production function. The agents realize that their actions affect each other. These agents may play this game cooperatively (i.e., they can collude to implement some outcome that jointly maximizes their welfare) or non-cooperatively (i.e., each agent selfishly maximizes). Gaming behavior, which game theory attempts to study, is all around us. I am reminded of a movie I saw on television some years back. It was set in the cold-war years and dealt with the possibility of a nuclear war between the U.S. and the Soviet Union. In the movie, the U.S. is at war with the Soviet Union and each contemplates the use of nuclear weapons. Each side

recognizes that if it launches an all-out nuclear strike first, it can obliterate the other side with limited damage to itself from retaliation. If both strike simultaneously, both are obliterated. Of course, if neither side strikes, both can emerge unscathed with a peaceful resolution. This is a classic example of a noncooperative game. We can represent the payoffs in the following matrix, called the strategic form of the game.

		Soviet Union Strategy	
		Strike First	Do Not Strike Unless Struck First
U.S. Strategy	Strike First	$(-\infty, -\infty)$	$(-L, -\infty)$
	Do Not Strike Unless Struck First	$(-\infty, -L)$	(G_u, G_s)

Here the first number in each payoff pair is the payoff to the U.S. from its chosen strategy and the second number is the payoff to the Soviet Union. L , G_u and G_s are positive numbers. It is assumed that if either country strikes first, the other will retaliate and the loss to the first striker is L , i.e., its payoff is $-L$. If neither country strikes, the U.S. enjoys a utility of G_u and the Soviet Union enjoys a utility of G_s . What will happen?

In the movie, both countries push the button simultaneously and there is total destruction. As it turns out, this is a Nash equilibrium in this game. A Nash equilibrium is a set of strategies for ($n > 1$) agents involved in a noncooperative game. These equilibrium strategies have the property that, holding fixed the equilibrium strategies of all the other agents, no agent can do better than to choose his own equilibrium strategy. In our war game, suppose the U.S. assumes that the Soviets will strike first. Then they cannot do strictly better by deviating from their own equilibrium strategy of striking simultaneously. Of course, the Soviets can go through the same reasoning and thus conclude that they cannot do better than choosing a preemptive nuclear strike. Hence, one Nash equilibrium in this game is for both to strike simultaneously and annihilate each other as in the movie. Each country gets a payoff of $-\infty$. Note that this is a Nash equilibrium because, holding each country's equilibrium strategy fixed, neither can do better for itself by not striking.

Now, you might object and say, "Isn't it better for neither side to strike, since then each is better off than

if either side had struck?" True! Indeed, if each side were to believe that the other side will not strike first, then each is strictly better off not striking first. In that case, the Nash equilibrium is for neither side to strike.

Thus, in this game there are two Nash equilibria in pure strategies (each side's equilibrium strategy is to choose a particular move with probability one). A reasonable question to ask is: which Nash equilibrium will obtain? The answer really depends on what the agents involved in the game believe. If the Americans and the Russians trust each other, we get the nice Nash equilibrium involving a peaceful resolution. If they don't, we get nuclear holocaust.

Multiple Nash equilibria are common. For example, in the Diamond and Dybvig [22] bank runs model, there are two Nash equilibria, one in which there is no run on banks and one in which there is a run. They focus on the "bad" Nash equilibrium and show that governmental intervention with federal deposit insurance can help eliminate the bad Nash equilibrium.

As this discussion indicates, the information that people have and the beliefs that they form are important ingredients in games. Indeed, three important elements in game-theoretic models are: information and beliefs (which are often linked), rationality, and strategic behavior. Much of game theory is built on the assumption that individuals are rational, although rationality need not be pervasive, as we will see later. Often the mere suspicion that some player may be irrational has far-reaching consequences. I'll argue later that explicit attention to beliefs in SMIF and the relative suppression of the role of beliefs in SMUF may be the major factor in the potentially different implications yielded by these models in finance.

Thus far I've discussed game theory in very general terms. The rest of this paper will be devoted to specifics. I'll focus my attention largely on noncooperative game theory, and then remark briefly on cooperative game theory. In both these strands of game theory, individual agents are assumed, at least for the most part, to be Bayesian rational, i.e., they revise their beliefs according to Bayes' rule whenever possible. I will then have a little bit to say about a new branch of game theory that has not been used much in finance, but that might have interesting applications. It is the concept of evolutionary stable strategies (ESS). This branch of game theory allows agents to behave in a non-Bayesian fashion and attempts to understand the economic world as one in which strategies, institutions, markets, and even

preferences are those that survive an evolutionary process of "fitness maximization."

My discussion of noncooperative game theory will revolve around an example. I have taken a simple financial signalling model and computed the equilibrium we would get under SMUF. I then view it as a model belonging to the SMIF class and point out differences. I'll follow this with a discussion of selected applications. This discussion is not intended to be an exhaustive review. Finally, I'll conclude with a discussion of cooperative game theory and ESS.

I. Noncooperative Game Theory: Why It Matters Who Moves First

A. The Sequence of Moves

Because the early signalling models of Spence [60] and others were not game-theoretic models, they were not really concerned with who moved first in asymmetric information settings. We know now, however, that the sequencing of moves can matter. In most asymmetric information models, we can visualize two parties: the informed and the uninformed. The informed party has some private information that the uninformed is trying to infer through some action or set of actions taken by the informed. In Spence's model this action (or signal) is education, in Ross' [52] model it is debt, and in Bhattacharya's [7] model it is dividends.

In SMUF, the often-implicit assumption is that the uninformed party moves first, in the sense that it specifies a menu of its own responses to various possible values of the signal that the informed agent could emit. Thus, the informed agent knows what the uninformed agent will do in response to each signal. For example, in the Ross [52] model of capital structure, the firm's manager (the informed agent) knows exactly how the capital market will price his firm if he selects a particular debt level. We can view that model as one in which the capital market offers the manager a menu of choices, where each choice is a combination of a debt level and an accompanying market value of the firm. The manager then picks the combination that maximizes his own payoff function. Even though the market does not know the manager's private information about his firm's intrinsic worth, it knows the possible values the firm can have. Thus, it can design the menu of combinations it offers the manager in such a way that, regardless of the manager's private information, he chooses a debt level that correctly reveals his firm's

value to the market. This is the familiar incentive compatibility¹ condition in financial signalling models and we will see shortly how it works.

A key feature of these models is that the uninformed agent is assumed to precommit to the menu. That is, the uninformed agent cannot change its mind about its own response after observing the signal emitted by the informed agent. This can often be a problem. To see why, consider the Bhattacharya [7] dividend signalling model.² In that model, the firm's manager announces a dividend that signals his firm's value. This dividend is promised as a perpetuity since the manager's information is presumed to not change through time. The signalling equilibrium is one in which the higher the value of the firm (privately known by the manager), the higher is the promised dividend. What ensures incentive compatibility in the model is that the firm is forced to use external financing to make up the shortfall whenever its realized cash flow from operations is less than the promised dividend, and external financing is more costly than retained earnings. This delivers the intuitive result that you will promise a high dividend only if your expectation of future cash flows is high enough. That is, this cost structure serves to ensure that an intrinsically lower valued firm will not be tempted to mimic its higher valued counterpart by also promising a high dividend. In this model, however, once the manager makes his dividend announcement, all of the informational asymmetry is resolved as the market knows as much as the manager. Why then should we subject this poor manager to the tyranny of having to pay this dividend every period, possibly at the cost of having to resort to distress financing to make up cash flow shortfalls? The answer is that if we don't, the manager will not signal truthfully in the first place. The difficulty with this is that we are asking agents to abide *ex post* by strategies that were determined to be *ex ante*

efficient but are not necessarily *ex post* efficient. This is often difficult to ensure. In our dividend signalling example, the firm's value would be higher if the manager paid one dividend and then stopped paying dividends. Everybody would know the firm's true value, the market price would be correct, and no more value-dissipating dividends would need to be paid.³

Another way of viewing the precommitment assumption in SMUF is that the Nash equilibrium⁴ is an equilibrium for the overall game, but it is not an equilibrium over every properly defined subgame within this overall game. Strictly speaking, one cannot make this claim within the context of Bhattacharya's dividend signalling model because the firm's only strategy is a single dividend level promised in perpetuity. But a natural redefinition of that game would be to make it truly dynamic and allow the firm to choose a possibly different dividend level in each period. In that case, a subgame would be the game that would ensue after the first dividend is paid, for example. Since the market knows as much as the firm's manager now, it is not a Nash equilibrium for the manager to pay dividends henceforth. That is, in this subgame, if we assume that investors in the market are Bayesian rational, they should price the firm at its true value, even if the manager discontinues paying dividends. Holding fixed this market response, it is not a Nash equilibrium for the manager to pay any more dividends.

We have now implicitly extended our notion of what constitutes an equilibrium. We would like the overall Nash equilibrium strategies to also be Nash equilibrium strategies over properly defined⁵ components (or subgames) of the overall game. When a particular Nash equilibrium satisfies this requirement, it is called a subgame perfect Nash equilibrium. Many of the equilibria characterized by models in the SMUF class are not subgame perfect.

Another difficulty with many of the models in the SMUF class is that they implicitly study games involving strategic interaction among agents, but they don't

¹By incentive compatibility we mean that no agent has an incentive to misrepresent, regardless of his private information. The deployment of this condition can be justified on the basis of the revelation principle (see, for example, Myerson [46]). This principle says that if there exists a Nash equilibrium in a general reporting game that does not necessarily impose truth-telling constraints, then the same equilibrium outcome can be replicated in a simpler reporting game with truth-telling constraints.

²My comments apply to all SMUF. I do not mean to single out Bhattacharya [7]. In this, as well as many other dissipative signalling models, signalling costs are shared by the sender (the manager in this case) as well as the receiver (the investors) since taxes, transactions costs, etc., reduce the intrinsic value of investors' holdings.

³As mentioned earlier, this will not "work" if the manager anticipates this, because then he will misrepresent. Because of its infinite horizon structure, Bhattacharya's [7] model has two types of precommitment: the usual precommitment to a specific (price) response by the uninformed agent (the market) and precommitment by the informed agent (the manager) to keep paying the same dividend perpetually.

⁴Not all models of the SMUF class study Nash equilibria.

⁵This is the notion of "proper subgames" as in Kreps and Wilson [36]. Loosely speaking, a proper subgame is a component of the overall game that inherits all of the "essential structure" of the overall game.

explicitly specify a formal structure for the game, complete with a precisely defined sequence of moves, information sets at various stages of the game, and beliefs.⁶ This is a problem mainly because Nash equilibrium requires that we also specify what happens off the equilibrium path. That is, what if an agent does something he was not supposed to do in equilibrium? Due to the lack of an extensive form for the game, models in the SMUF class usually specify a set of ad hoc restrictions on reactions to out-of-equilibrium moves. Whether these restrictions would hold up if we could endow the game with the appropriate extensive form remains an open question.

This problem, in turn, leads to another potential difficulty. Quite often, SMUF do not have any Nash equilibria. Rothschild and Stiglitz [53], in their famous paper on insurance markets, showed that a Nash equilibrium may fail to exist when the (uninformed) insurance company moves first to specify a set of insurance contracts and the privately informed individual buying the insurance moves next, even in the case when the informed agent can be one of only two types. The reason for this nonexistence will become clear later on when I analyze a specific model. Later, Riley [51] established the even more powerful result that, under plausible conditions on the cross-sectional distribution of types, a Nash equilibrium never exists in the SMUF class with a continuum of possible types of informed agents.⁷ This problem of nonexistence of a Nash equilibrium led Riley [51] to define a non-Nash equilibrium concept called the reactive equilibrium (henceforth referred to as the Riley Reactive Equilibrium or RRE for short). Riley provided sufficient conditions for the existence of an RRE with a continuum of types.⁸

The picture changes colors when we turn to SMIF. Instead of assuming precommitment by the uninformed, we now assume that the informed agent moves first and emits the signal, based on his private information. The uninformed agents rationally interpret this

signal and revise their priors about the informed agent's type, using Bayes' rule whenever possible. What do I mean by whenever possible? Well, suppose there are n possible types of the informed agent, and the uninformed agents know that there are n possible types, but they cannot determine precisely a particular informed agent's type. Suppose that the common belief is that there will be a separating signalling equilibrium, with each type of informed agent emitting a distinct signal. Then there are n possible equilibrium values of the signal. So if an informed agent emits one of these n signals, the uninformed can use Bayes' rule to update their priors. But what if there is a defection from the conjectured equilibrium, with the informed agent choosing a signal that is outside the set of n equilibrium signals? The uninformed agents cannot use Bayes' rule now since the prior probability of the signal was zero. How the uninformed agents are supposed to formulate their beliefs in such a setting has occupied the attention of game theorists for quite some time now, and I'll return to this issue later.

This leads me to an interesting digression. In statistical decision theory, zero probability events are inconsequential because, by definition, the likelihood of their occurrence is zero, so that they don't affect expected utility. But in game theory, we can't ignore such events. The reason is that what has zero probability is endogenous to the equilibrium we are focusing on. A zero probability event in one equilibrium may have positive probability in another. Moreover, whether an outcome is an equilibrium depends on the zero probability events, not just the positive probability events.

Returning to SMIF, one major advantage of these models is that there is no precommitment requirement of the type encountered with SMUF. Another advantage is that everything is made precise: the sequence of moves, information sets, beliefs, etc. This eliminates much of the ad hoc structure inherent in SMUF. Finally, there is usually no paucity of equilibria. Indeed, we now have an embarrassment of riches, as there is often a large number of Nash equilibria, and the modeler may be required to choose between them. Although a final answer has yet to be provided, game theory has provided a way to think about this choice. This is an issue I turn to next. Rather than go through the literature on refinements of Nash equilibrium — which is quite extensive and growing — I'll do a simple example familiar to finance researchers and show how we can compute the various equilibria.

⁶This formal structure, with all of the dynamic features intact, is called the extensive form of the game. See Engers and Fernandez [24] for an attempt to provide a formal game-theoretic structure to some of the models in the SMUF class.

⁷By a Nash equilibrium here I mean the Rothschild and Stiglitz [53] version rather than the incompletely specified Nash equilibrium of Spence [60].

⁸Earlier, Wilson [66] proposed his "anticipatory equilibrium" as another non-Nash equilibrium concept. But Riley showed that, with a continuum of informed agents, an anticipatory equilibrium also does not exist.

Before getting to that, I would like to make a simple point. In any signalling game, we have to deal with defections from the equilibrium, i.e., agents possibly doing things they are not supposed to in equilibrium. Who can defect depends on how the moves are sequenced in the game. In signalling games with well-defined first and second movers, it is always the first mover who has the opportunity to defect. Thus, in SMUF it is the uninformed agent who may defect by offering the informed agent a contract that is outside the equilibrium set of contracts. In SMIF, it is the informed agent who may defect by emitting a signal that lies outside the set of equilibrium signals.

B. A Debt Signalling Example: SMUF, SMIF and Refinements

Consider a single period economy in which all agents are risk neutral. At $t = 0$ (the beginning of the period), firms announce debt levels and at $t = 1$, the debt is repaid out of cash flows if possible. There are three types of firms in the economy: the good, the bad and the ugly. The density function of the end-of-period cash flow for the good firm is $f_g(x)$, for the bad firm is $f_b(x)$ and for the ugly firm is $f_u(x)$. The support of $f_i(x)$ is $(-\infty, \infty) \forall i \in \{g, b, u\}$. Also, $f_i(x) > 0 \forall x \in (-\infty, \infty)$ and $\forall i \in \{g, b, u\}$. The riskless rate of interest is zero for simplicity. The utility function of the firm's manager is defined over the market value of the firm at $t=0$ (the manager's wage is aV , where a is a positive, finite, real valued constant and V is the market value of the firm at $t=0$) and whether or not the firm goes bankrupt at $t=1$. Bankruptcy is possible only if the firm is levered. In the event of bankruptcy, the manager is assessed a personal (nonpecuniary) penalty of $\pi \in (0, \infty)$, which is a constant and does not directly affect firm value. This penalty is purely dissipative, i.e., π does not accrue to the investors. Let F_i be the cumulative distribution function (CDF) corresponding to the density function f_i . The CDFs of the three types of firms are ordered by the following first-order stochastic dominance (FOSD) relationship: g FOSD b FOSD u . Investors' (common) prior belief is that a fraction θ_g of all firms are good, a fraction θ_b are bad and a fraction $\theta_u = 1 - \theta_g - \theta_b$ are ugly. However, although each firm's manager knows his own firm's type, investors are unable to a priori distinguish one type of firm from another. The manager makes decisions to maximize his own expected utility. There are no taxes or bankruptcy costs. We want to analyze this as a signalling game in

which the manager's choice of debt, D (chosen from some compact interval of real numbers), is a signal of his private information.⁹

1. Analysis of the Game as SMUF

We can write the expected utility of a manager choosing debt level D_j and managing a firm of type i as

$$EU(j | i) = aV(D_j) - \pi F_i(D_j) \quad (1)$$

where we recognize that the firm's current market value can be based only on the observed debt signal. I will derive the unique RRE in this game. In our context the RRE is a set of three distinct signals (debt levels) such that: (i) each firm is correctly priced (the expected return of investors who purchase each firm is zero) and (ii) there does not exist any alternative scheme (a pair consisting of a signal and an associated firm value) that could be offered to firms by any coalition of investors, say C_1 , such that C_1 could earn a nonnegative expected return on this scheme even after another coalition, say C_2 , reacts to the initial defection with yet another out-of-equilibrium scheme. The key is that C_2 should see a positive expected return in reacting to C_1 , and no further reactions can impose strict losses (negative expected return) on C_2 .

The RRE for this game is as follows. The worst type of firm, which is the ugly firm, sets its debt level $D_u = 0$. The bad firm signals with a debt level $D_b > 0$ which is set to satisfy

$$aV(D_b) - \pi F_u(D_b) = aV(0) \quad (2)$$

where the left hand side (LHS) is the expected utility of the ugly firm's manager from mimicking the bad firm's manager and the right hand side (RHS) is the expected utility of the ugly firm's manager from signalling truthfully. Note that competitive capital market pricing requires that, consistent with the conjectured RRE,

$$V(D_b) = \mu_b = \int_{-\infty}^{\infty} x f_b(x) dx \quad (3)$$

and

⁹Those familiar with Ross [52] will recognize that this is essentially the Ross model.

$$V(0) = \mu_u = \int_{-\infty}^{\infty} x f_u(x) dx, \quad (4)$$

where μ_b and μ_u are the mean cash flows of the bad and ugly firms, respectively.

Substituting (3) and (4) in (2) and solving for D_b gives

$$D_b = F_u^{-1}(a[\mu_b - \mu_u]/\pi). \quad (5)$$

Similarly, the good firm signals with a debt level $D_g > 0$ which satisfies

$$aV(D_g) - \pi F_b(D_g) = aV(D_b) - \pi F_b(D_b) \quad (6)$$

where the LHS is the expected utility of the bad firm from mimicking the good firm and the RHS is the expected utility of the bad firm from signalling truthfully. Again, competitive capital market pricing requires that, consistent with the conjectured RRE,

$$V(D_g) = \mu_g = \int_{-\infty}^{\infty} x f_g(x) dx. \quad (7)$$

Substituting (3) and (7) in (6) and solving for D_g gives us

$$D_g = F_b^{-1}(\{a[\mu_g - \mu_b] + \pi F_b(D_b)\}/\pi) \quad (8)$$

where D_b is given by (5). It is easy to verify that the ugly firm will not mimic the good firm and the good firm will not mimic either of its lower-valued counterparts. Also, the bad firm will not mimic either the ugly or the good firm. Moreover, $D_g > D_b > 0$.

It is appropriate to view this equilibrium as the outcome of a game in which the uninformed capital market is a single player that offers the following menu of choices to firms: (i) choose a debt level of zero and we will price your firm at μ_u , (ii) choose a debt level of D_b and we will price your firm at μ_b , and (iii) choose a debt level of D_g and we will price your firm at μ_g . The firm's manager then proceeds to choose a particular debt level and each firm is priced in equilibrium as if investors had the same information as managers.

Precommitment from the market to adhere to this pricing strategy is almost as crucial here, however, as it is in the dividend signalling model of Bhattacharya [7] that we discussed earlier. Recall that in that case the market was precommitted to a pricing response conditional on the manager paying a particular dividend perpetually; after the first period dividend, however, all

informational asymmetry is resolved and there is little reason for the manager to keep paying the dividend, except for the necessity of the precommitment to be adhered to for ex ante incentive compatibility reasons. We can view the nature of precommitment in our debt signalling model in a similar light. Once the firms have chosen their respective debt levels, suppose the manager of the firm that issued debt D_g were to announce to the market; "You know that I have chosen D_g , so my firm is good. There is no reason now for you to force me to be levered this way until the end. Could I not reduce my leverage by repurchasing some of my debt with an equity issue?" Clearly, in an ex post sense, investors are no worse off if the manager did this, and the manager is strictly better off. But precommitment in this game requires that the market's price response of μ_g be rigidly conditional on the manager not being able to do this, or else ex ante incentive compatibility is destroyed. In other words, this equilibrium hinges on there being precommitment to persist with an arrangement that is ex post efficient to dissolve.¹⁰ Models in the SMUF class lack the appropriate structure to deal satisfactorily with this complication.

Another point to note is that the equilibrium we have characterized is an RRE under our assumptions. However, suppose θ_g is very close to one. Then this equilibrium will not be a Nash equilibrium (in the Rothschild and Stiglitz [53] sense). In our context, a Rothschild and Stiglitz Nash equilibrium is a set of three distinct signals such that: (i) each firm is correctly priced in the equilibrium, and (ii) there does not exist an alternative scheme (a pair of debt signal and associated firm value outside the equilibrium set) such that the coalition of investors offering it could earn strictly positive profit and there is at least one type of firm that would prefer this scheme to its equilibrium allocation. Intuitively, the reason why a Nash equilibrium fails to exist when θ_g is close to 1 is that a group of uninformed investors can successfully defect from the equilibrium by offering a contract that prices all firms at a pooling price of $\bar{\mu} = \theta_g \mu_g + \theta_b \mu_b + [1 - \theta_g - \theta_b] \mu_u$ and a debt level of zero. If θ_g is sufficiently close to 1, it can be shown that

¹⁰One could argue that the RRE could cope with a possible debt repurchase decision by the manager by redefining the game to treat the initial debt issue and the subsequent debt repurchase as joint signals. However, we are still left with the problem of precommitment to the repurchase fraction. This precludes arrangements that are ex post more efficient than these stipulated in the equilibrium.

all firms will prefer this defection to their equilibrium allocations. For example, the manager of the good firm finds that the gain from signalling via debt, $a[\mu_g - \bar{\mu}]$, is insufficient to compensate him for π ; the bad and the ugly firms are, of course, happy to be subsidized. In this case, no Nash equilibrium exists, since a pooling outcome in such games is never a Nash equilibrium (see Rothschild and Stiglitz [53]).¹¹

The RRE characterized here has predictions that are generally in accord with some of the recent empirical evidence, namely that there is a positive association between debt and firm value.¹²

2. Analysis of the Game as SMIF

Now imagine that the informed manager is moving first, choosing his debt signal, and the a priori uninformed capital market moves next by setting a market value for the firm. An extensive form for this game is shown in Exhibit 1. This extensive form is drawn as if the firm had only four debt levels to choose from — the equilibrium levels $0, D_b^*$ and D_g^* (corresponding respectively to the debt choices of the ugly, bad and good firms in a perfectly separating equilibrium) and some other arbitrary out-of-equilibrium debt level \bar{D} . This is done merely for simplicity. In principle, the firm can choose from a continuum of debt levels. Moreover, for simplicity I also assume that the market can choose from only four valuation responses. In this extensive form, a vertical dotted line indicates that the market is unable to tell which type of firm made that move. Each circle (clear or full) is a node, and the box above the node indicates whose turn it is to move. The full black circles at the extreme left and extreme right of the figure are payoff nodes — the pair (S_1, S_2) indicates that S_1 is the manager's expected utility and S_2 is investors' expected return. (For more details on extensive forms, see Kreps and Wilson [36]).

¹¹Dasgupta and Maskin [20, 21] show that the nonexistence of Nash equilibrium is due to the restriction to pure strategies and that randomized strategies can often restore existence, thereby verifying an earlier conjecture by Ross [52].

¹²For example, Masulis [39, 40] reports that exchange offers that result in increases (decreases) in leverage are associated with positive (negative) abnormal common stock returns. More recently, Cornett and Travlos [19] specifically investigate the information effect caused by a firm's change in capital structure via debt-for-equity and equity-for-debt exchange offers. Their evidence indicates that debt-for-equity swaps lead to abnormal stock price increases, while equity-for-debt swaps lead to abnormal stock price decreases.

Suppose, for the moment, that $D_b^* = D_g^* = D^* > 0$. I'll now show that a completely pooling outcome in which each firm chooses an equilibrium debt level $D^* > 0$ satisfying

$$a\mu_u > a\bar{\mu} - \pi F_u(D^*) > 0 \quad (9)$$

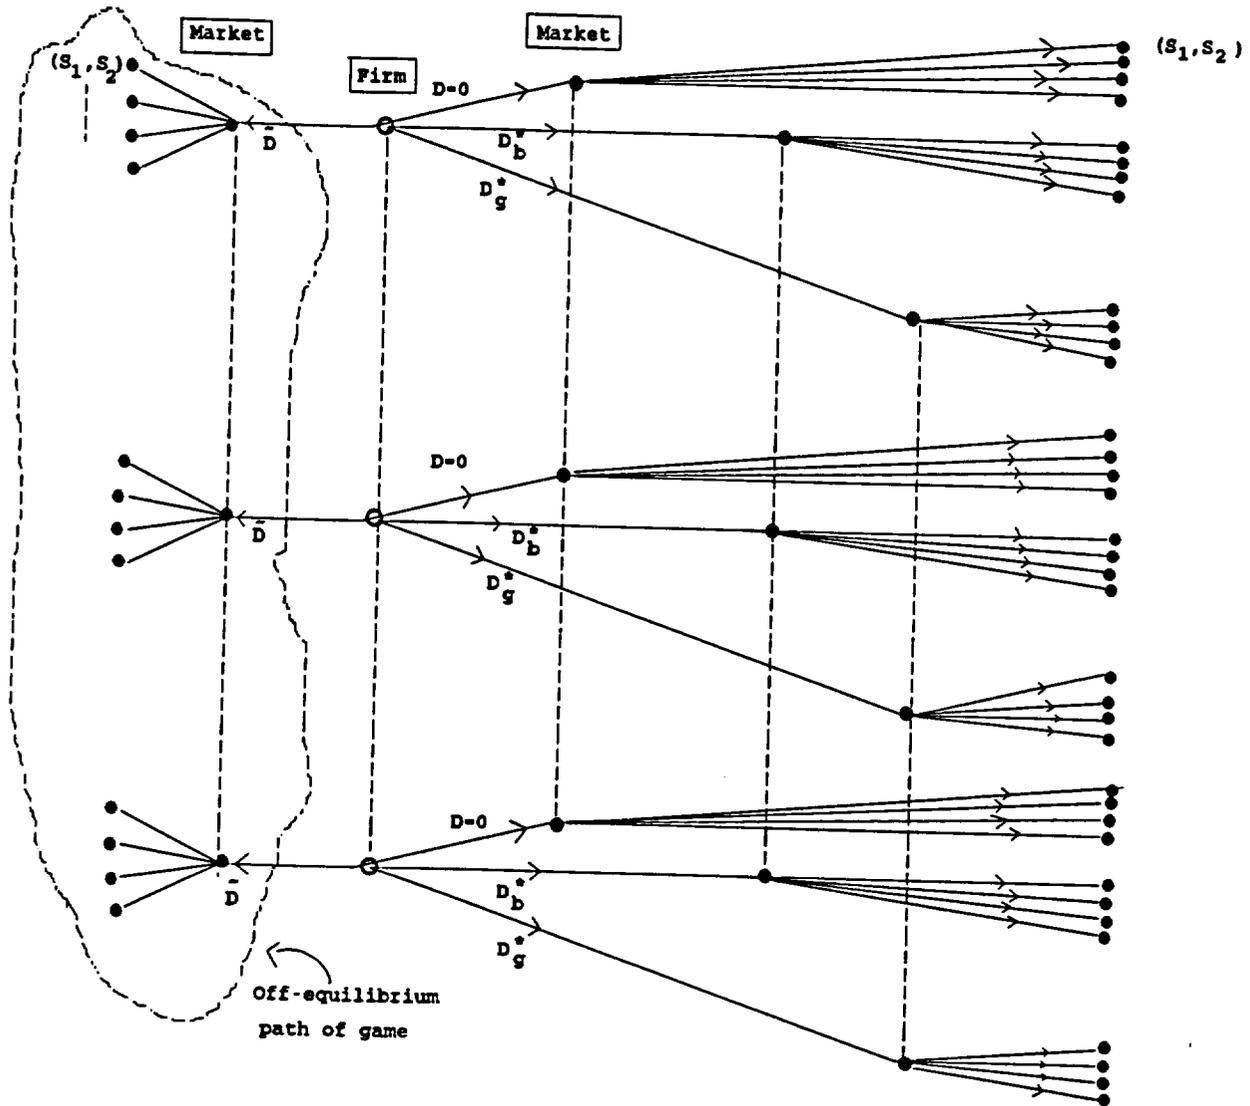
is a Nash equilibrium where $\bar{\mu}$ is the pooling mean defined earlier. To see why this is a Nash equilibrium, suppose a manager defects by choosing $\bar{D} \neq D^*$. This is an out-of-equilibrium move and the component of the game inside the dotted area represents the off-the-equilibrium path part of this game. How should the capital market react to this event which was supposed to have zero probability? Although the Nash equilibrium concept stipulates how the market should react to each event (including those that have zero probability), it places no restrictions on reactions to out-of-equilibrium moves, other than that such reactions help to sustain the equilibrium. So suppose we define a number $K > 0$ satisfying

$$a\mu_u/K < a\bar{\mu} - \pi F_u(D^*) \quad (9)'$$

and assume that whenever a manager chooses a debt level $\bar{D} \neq D^*$ (where \bar{D} could be zero), his firm will be priced at μ_u/K . That is, this response by the market is a part of the game that is off-the-equilibrium path and constitutes a threat used to support the Nash equilibrium. It is straightforward to verify that no manager will choose to defect from our conjectured Nash equilibrium.

This Nash equilibrium does not make much sense, however. All firms are being pooled at a common price and yet their managers are required to issue a positive amount of debt which is personally costly to them. The reason why this is a Nash equilibrium is that the equilibrium concept permits too much latitude in specifying reactions to out-of-equilibrium moves. This is where refinements of Nash equilibrium come in. The earliest of these refinements was subgame perfection which I referred to earlier. Unfortunately, subgame perfection admits too many unreasonable equilibria under asymmetric information, even though it prunes away some Nash equilibria. That is, there are many subgame perfect Nash equilibria that are unreasonable. A class of situations in which this happens is when the only proper subgame is the overall game

Exhibit 1. Illustrative Extensive Form for Debt Signalling Game



itself. In this case it is obvious that every Nash equilibrium is subgame perfect.

A further refinement was proposed by Kreps and Wilson [36]. It is called sequential equilibrium and it puts restrictions on strategies and beliefs. In a sequential equilibrium, we must not only specify strategies and beliefs of players at every node of the game that can be

reached in equilibrium, but also at nodes that cannot be reached in equilibrium. That is, a sequential equilibrium is a Nash equilibrium with the added requirements that, at any point in the game, the strategies of players are Nash, given the strategies and beliefs of the other players, and the beliefs at each information set in the game are formulated according to Bayes' rule

wherever possible.¹³ Simply put, in the context of our example, sequential equilibrium imposes two restrictions that Nash equilibrium does not. First, over any (properly defined) subgame of our overall game, strategies of all players should be Nash.¹⁴ Second, as in the case of Nash equilibrium, beliefs must be revised in accordance with Bayes' rule whenever an equilibrium move is observed, but when an out-of-equilibrium move is observed, there must be some belief on the part of the uninformed that makes their reaction to this out-of-equilibrium move rational on the basis of that belief. In game-theoretic terminology, the reaction of the uninformed to even an out-of-equilibrium move should be a best response given some belief. A key advantage of sequential equilibrium over subgame perfection is that, unlike subgame perfection, the sequential equilibrium concept allows us to eliminate some Nash equilibria even in games where the only proper subgame is the overall game itself.

We can see now that our Nash equilibrium is not necessarily sequential. There is no belief on the part of investors that would justify a (competitive) market value of μ_u/K as a best response for $K > 1$. In other words, the threat contained in a response of μ_u/K for $K > 1$ is not credible. Intuitively, the reason for this is that $K > 1$ means that the firm choosing an out-of-equilibrium move is priced at less than its lowest possible value. That is, whenever the K satisfying (9) exceeds 1, the Nash equilibrium described earlier is not sequential.¹⁵ However, a pooling equilibrium in which the strategy of each firm is to issue zero debt with probability one, and all firms are priced at $\bar{\mu}$ if they issue zero debt and at μ_u otherwise is a sequential equilibrium under the following structure of beliefs:

$\Pr(\text{firm is of type } i \mid \text{firm issues zero debt}) = \theta_i$, the prior probability that firm is of type i for all $i \in \{g, b, u\}$, and $\Pr(\text{firm is of type } i \mid \text{firm issues positive debt})$

¹³Strictly speaking, I have defined a Perfect Bayesian Nash equilibrium. Sequential equilibrium adds a third condition that the equilibrium strategies and beliefs should be the limit of some sequence of (Bayesian) rational beliefs and completely mixed strategies (see Kreps and Wilson [36] and Rasmusen [50]). However, the equilibria I study henceforth are sequential Nash equilibria. Note that a "strategy" in this context is defined as a probability distribution over all possible moves that the informed player can make, i.e., a strategy specifies the probability with which each move will be made.

¹⁴This is the subgame perfection requirement.

¹⁵On the other hand, if $K < 1$, then an equilibrium with $D > 0$ may be sequential but will not survive many of the refinements of sequential equilibrium.

$= \begin{cases} 1 & \text{if } i = u \\ 0 & \text{if } i \neq u \end{cases}$. The sequential equilibrium strategy of all firms is to issue zero debt with probability one regardless of type. Note that if a firm issues zero debt, the Bayesian posterior of investors about this firm's type coincides with the prior belief. Moreover, conditional on the out-of-equilibrium belief that any deviating firm is ugly, the market's reaction is a best response. Given this response, no firm defects from the equilibrium. Note that what I have really proved is that this is a Perfect Bayesian Nash equilibrium, since I have not checked that the equilibrium strategies and beliefs are the limit of some sequence of (Bayesian) rational beliefs and strategies. It is, however, straightforward to check that this third requirement is satisfied here.¹⁶

As is transparent, this no-signalling equilibrium rests on a somewhat objectionable specification of out-of-equilibrium beliefs. Game theorists have developed a literature on refinements of the sequential equilibrium itself, which attempts to provide ways of put-

¹⁶To see this, let $\sigma_i(D)$ represent the strategy of the manager of the type- i firm, with $i \in \{g, b, u\}$, so that the probability with which this manager chooses debt D is $\sigma_i(D)$. Define the sequence of strategies, $\{\sigma_i^n(D), n = 1, 2, \dots, \infty\}$, as follows:

$$\sigma_g^n(D) = \begin{cases} 1 & \text{if } D = 0 \\ 0 & \text{if } D > 0 \end{cases} \text{ for every } n$$

$$\sigma_b^n(D) = \begin{cases} 1 & \text{if } D = 0 \\ 0 & \text{if } D > 0 \end{cases} \text{ for every } n$$

and

$$\sigma_u^n(D) = \begin{cases} 1 - \frac{1}{n} & \text{if } D = 0 \\ f(D)/n & \text{if } D > 0 \end{cases}$$

where $f(D)$ is some (arbitrary) probability density function over $(0, \infty)$ that is finite and nonnegative over $(0, \infty)$. Since $f(D)$ is a probability density, $\int_0^{\infty} f(D)dD = 1$. Note that this sequence of strategies is such that $\lim_{n \rightarrow \infty} \sigma_i^n(D) = \sigma_i^*(D) \forall i \in \{g, b, u\}$, where $\sigma_i^*(D)$ is the equilibrium strategy of the type- i firm's manager, and it satisfies

$$\sigma_i^*(D) = \begin{cases} 1 & \text{if } D = 0 \\ 0 & \text{if } D > 0 \end{cases}$$

for every $i \in \{g, b, u\}$. With this sequence of strategies, we can also specify beliefs that satisfy Bayes' rule for out-of-equilibrium moves as well as equilibrium moves (recall that Bayes' rule cannot be applied to out-of-equilibrium moves, given just the equilibrium strategies). To see this, let $P^n(i \mid D)$ be the market's posterior belief that the firm

ting further sensible restrictions on beliefs off the equilibrium path. One such refinement is known as the intuitive criterion of Cho and Kreps [16]. An equilibrium passes the intuitive criterion if: (i) it is sequential and (ii) it satisfies the following additional restriction on beliefs and strategies off the equilibrium path.

Suppose the uninformed observe an out-of-equilibrium move, say m , in a game with N possible types of the informed agent. If there does not exist any belief on the part of the uninformed that would support a best response on their part such that a subset of player types, say $S \subset N$, would defect with move m , then the uninformed should assign zero probability (in their posteriors) to the event that the defector's type belongs to S . Then, for a sequential equilibrium to survive the Cho-Kreps intuitive criterion, for every type remaining in $M \setminus S$, there exists at least one probability distribution of beliefs of the uninformed, concentrated on the remaining types $M \setminus S$, such that the accompanying best response of the uninformed does not induce that informed agent type to strictly prefer to defect from the equilibrium.

This criterion is indeed very intuitive. It says that we should eliminate sequential equilibria that can only be supported by out-of-equilibrium beliefs that put positive probability weight on types that would never wish

is of type i , having observed move D as part of the n^{th} element of the sequence of strategies $\{\sigma^n(D)\}$. Then, using Bayes' rule

$$P_r^n(\text{firm is ugly} \mid \text{some specific } D > 0) = \frac{\{f(D)/n\} \{1 - \theta_g - \theta_b\}}{\{f(D)/n\} \{1 - \theta_g - \theta_b\} + 0} = 1 \quad \forall n.$$

$$P_r^n(\text{firm is ugly} \mid D=0) = \frac{[1 - \frac{1}{n}] [1 - \theta_g - \theta_b]}{\{[1 - \frac{1}{n}] [1 - \theta_g - \theta_b] + \theta_g + \theta_b\}}$$

$$= \frac{[1 - \frac{1}{n}] [1 - \theta_g - \theta_b]}{[1 - \frac{1}{n}] \{1 - \theta_g - \theta_b\}}$$

$$P_r^n(\text{firm is bad} \mid D=0) = \frac{\theta_b}{[1 - \frac{1}{n}] \{1 - \theta_g - \theta_b\}}$$

$$P_r^n(\text{firm is good} \mid D=0) = \frac{\theta_g}{[1 - \frac{1}{n}] \{1 - \theta_g - \theta_b\}}$$

It is apparent that

$$\lim_{n \rightarrow \infty} P_r^n(i \mid D > 0) = \begin{cases} 1 & \text{if } i = u \\ 0 & \text{if } i = g \text{ or } b \end{cases} \quad \text{and} \quad \lim_{n \rightarrow \infty} P_r^n(i \mid D=0) = \theta_i$$

for every i . Thus the sequence of Bayesian rational beliefs associated with the assumed sequence of strategies converges to the equilibrium beliefs.

to defect from the equilibrium, when reactions by the uninformed to defections by the informed are governed by the sequential equilibrium best response (rationality) requirement.¹⁷

We can now apply the Cho-Kreps intuitive criterion to eliminate the pooling sequential equilibrium. Define D^o as a debt level such that

$$a\bar{\mu} = a\mu_g - \pi F_b(D^o). \quad (10)$$

We continue to assume that a player will not defect from the equilibrium unless he/she is strictly better off for doing so. Thus, D^o is a debt level such that the bad firm will not wish to defect from the equilibrium even if the uninformed believe with probability one that the defector is good. Thus, there does not exist any belief on the part of the uninformed that would induce defection by the manager of the bad firm, since he'll do strictly worse with any belief that puts less than one probability weight on the defector being good. Because of the FOSD relationship among the types, it is easy to see that

$$a\bar{\mu} > a\mu_g - \pi F_u(D^o) \quad (11)$$

so that the manager of the ugly firm strictly prefers to stay with his pooling equilibrium strategy, and

$$a\bar{\mu} < a\mu_g - \pi F_g(D^o) \quad (12)$$

so that the manager of the good firm strictly prefers to defect from the pooling equilibrium.

We have thus established that there does not exist any belief such that either the ugly or the bad firm's manager will wish to defect with the debt level D^o . According to the Cho-Kreps intuitive criterion, we should set at zero the probability that the defector is either the ugly or the bad firm. Thus, the only admissible posterior belief is that the defector is the good firm with probability one. Given this belief, we now require that the (competitive) market's price response be a best response. That is, the defector should be priced as if it is a good firm. And if this is the case, we have already established that the manager of the good firm will indeed strictly prefer to defect by issuing debt

¹⁷That is, the uninformed must react to a defection in a manner rationalized by at least some belief on their part.

of D^0 . Hence, the pooling sequential equilibrium does not satisfy the Cho-Kreps intuitive criterion.

There is, however, another sequential equilibrium that does survive the Cho-Kreps intuitive criterion, given appropriate assumptions regarding the market's prior beliefs about types. This is a partly pooling equilibrium in which the good firm's manager signals with a debt level $D_g > 0$ and the managers of the bad and ugly firms pool at a debt level of zero.

For equilibrium moves, the a priori uninformed investors should revise their beliefs in accordance with Bayes' rule. Since they know that only the bad and the ugly firms are pooling at a zero debt level, when they see a firm that issues zero debt, they should believe that there is a probability $\theta_b/[1 - \theta_g]$ that this firm is bad and a probability $[1 - \theta_g - \theta_b]/[1 - \theta_g]$ that this firm is ugly. Thus, in equilibrium the common pooling price of firms that issue zero debt should be

$$\bar{V}_p = [\theta_b/(1 - \theta_g)]\mu_b + [(1 - \theta_g - \theta_b)/(1 - \theta_g)]\mu_u. \quad (13)$$

The equilibrium expected utility of the manager of either a bad or an ugly firm is

$$a\bar{V}_p. \quad (14)$$

A firm that issues debt of D_g will be priced at μ_g . So the equilibrium expected utility of the manager of the good firm is

$$a\mu_g - \pi F_g(D_g). \quad (15)$$

For incentive compatibility we need

$$a\bar{V}_p \geq a\mu_g - \pi F_b(D_g). \quad (16)$$

so that the manager of the bad firm does not mimic that of the good firm. As usual, (16) holds tightly in equilibrium, so that we obtain the equilibrium debt level

$$D_g^* = F_b^{-1}(\{a\mu_g - a\bar{V}_p\}/\pi). \quad (17)$$

Given (16) as an equality, it is straightforward to verify that

$$a\bar{V}_p > a\mu_g - \pi F_u(D_g^*) \quad (18)$$

and

$$a\bar{V}_p < a\mu_g - \pi F_g(D_g^*). \quad (19)$$

Thus, all the incentive compatibility conditions are satisfied.

I'll leave it up to the reader to verify that this is a Nash equilibrium. To check that this is a sequential equilibrium, consider a defection $D \in (0, D_g^*)$. Suppose investors assign the belief $\Pr(\text{defector is ugly} \mid D \in (0, D_g^*) \text{ observed}) = 1$. Then the defector is priced at μ_u . The expected utility of the defecting manager is $a\mu_u - \pi F_i(D) \quad \forall i \in \{g, b, u\}$. Note that since $\bar{V}_p > \mu_u$, $a\bar{V}_p > a\mu_u - \pi F_i(D) \quad \forall i \in \{b, u\}$ so that the managers of the bad and ugly firms don't defect. And since

$$aV_g - \pi F_g(D_g^*) > a\bar{V}_p > a\mu_u - \pi F_g(D),$$

the manager of the good firm doesn't defect. Thus, this is a sequential equilibrium. Clearly, no manager will defect with $D > D_g$ even if the firm is priced at μ_g .

Let us now verify that this equilibrium satisfies the Cho-Kreps intuitive criterion. Consider some defection $D \in (0, D_g^*)$. Is there any type that can be ruled out as a potential defector, regardless of the market's beliefs in response to the defection? Consider the manager of the ugly firm. Although (18) is satisfied, there may exist a debt level, say $D_u^* < D_g^*$, such that

$$a\bar{V}_p = a\mu_g - \pi F_u(D_u^*). \quad (20)$$

So if $D < D_u^*$, the manager of the ugly firm will defect, conditional on the market believing that the defector is the good firm with probability one. Thus, if the manager of the bad firm wishes to defect with a debt level such that he can convince the market that his firm is not ugly, he must set this debt level $D \geq D_u^*$; with this debt level, the manager of the ugly firm will never defect.

Suppose now that a defection $D \in [D_u^*, D_g^*)$ is observed. By the above logic, we can rule out the ugly firm as the defector. For the equilibrium to survive the intuitive criterion, for each type there must exist at least one belief concentrated over the bad and good firms such that the associated market best response does not cause that type of firm to strictly prefer to defect. Assume now that the prior beliefs of the market are such that

$$a\bar{V}_p > a\mu_b - \pi F_b(D_u^*). \quad (21)$$

So if we consider the following belief (admissible according to the intuitive criterion) $\Pr(\text{defector is the bad firm} \mid \text{defection } D \in [D_u^*, D_g^*] \text{ observed}) = 1$, then it does not pay for the bad firm to defect. This is because, given this belief, (21) guarantees that the expected utility of the bad firm's manager from defection is strictly less than his equilibrium expected utility. Moreover, it is transparent that neither the ugly nor the good firm will wish to defect.¹⁸ Hence, the equilibrium survives the Cho-Kreps intuitive criterion.

There are two points worth mentioning at this stage. First, the empirical content of this model is similar to that of the RRE identified earlier, even though the two equilibria are different. Debt issuance is good news.

Second, even with a refinement such as the Cho-Kreps intuitive criterion, it is possible to get some pooling in equilibrium.¹⁹ This is in sharp contrast to the predictions one gets using the RRE. At a common sense level, it is not obvious why equilibria should always be perfectly separating. For instance, in our illustration if the cross-section of firms is such that it leads to the prior beliefs that $\theta_g = 0.10$, $\theta_b = 0.85$ and $\theta_u = 0.05$, then it is not clear why the managers of bad firms would like to distinguish themselves from those of ugly firms by issuing personally costly debt. If the bad and the ugly firms issue no debt (while the good firm issues debt), the pooling price for these two types of firms will be very close to the first best price of the bad firm anyway. So why incur the cost of signalling? As far as evidence on this issue is concerned, Cadsby, Murray and Maksimovic [12] have recently reported the results of some experiments. They found that when the theory predicts a unique sequential equilibrium, the experimental evidence is consistent with the theory. When the theory allows for sequential equilibria that are pooling, partially pooling or separating, only the pooling equilibrium—which is the most efficient—is observed.

The partially pooling equilibrium we have just seen is not the only one that survives the Cho-Kreps intuitive criterion. There is a perfectly separating sequential equilibrium which also survives the Cho-Kreps intuitive criterion.

¹⁸It can be readily checked that if the inequality in (21) were replaced by $a\bar{V}_p < a\mu_b - \pi F_b(D_u^*)$, then this equilibrium will fail to survive the intuitive criterion.

¹⁹If we had considered only two types of firms in this illustration, the intuitive criterion would have ruled out any pooling in equilibrium.

tive criterion. In this sequential equilibrium, the good firm issues debt D_g , the bad firm issues debt \hat{D}_b , and the ugly firm stays unlevered, where $\hat{D}_g > \hat{D}_b > 0$. Since each type of firm chooses a distinctly different strategy, investors will be able to distinguish perfectly among firms in equilibrium. The equilibrium expected utility of the ugly firm's manager is

$$a\mu_u. \quad (22)$$

The equilibrium expected utility of the bad firm's manager is

$$a\mu_b - \pi F_b(D_b) \quad (23)$$

and that of the good firm's manager is

$$a\mu_g - \pi F_g(D_g). \quad (24)$$

For incentive compatibility we need

$$a\mu_u \geq a\mu_b - \pi F_u(D_b) \quad (25)$$

so that the manager of the ugly firm does not mimic that of the bad firm. Since (25) holds tightly in an equilibrium, we obtain

$$\hat{D}_b = F_u^{-1}(\{a\mu_b - a\mu_u\}/\pi). \quad (26)$$

Given (25) as an equality, it can be verified that

$$a\mu_b - \pi F_b(\hat{D}_b) > a\mu_u \quad (27)$$

so that the manager of the bad firm strictly prefers his own allocation to that of the ugly firm. For incentive compatibility we also need

$$a\mu_b - \pi F_b(\hat{D}_b) \geq a\mu_g - \pi F_b(D_g). \quad (28)$$

Since (28) holds as an equality in equilibrium, we have

$$\hat{D}_g = F_u^{-1}(\{a\mu_g - a\mu_b + \pi F_b(\hat{D}_b)\}/\pi). \quad (29)$$

It is easy to check that

$$a\mu_g - \pi F_g(\hat{D}_g) > a\mu_b - \pi F_g(\hat{D}_b) > a\mu_u \quad (30)$$

so that the manager of the good firm strictly prefers his own equilibrium allocation to any other.

I'll leave it up to the reader to verify that this is a sequential equilibrium. To check that this equilibrium satisfies the Cho-Kreps intuitive criterion, consider a defection $D \in (0, \hat{D}_b)$. Clearly, we cannot a priori rule out any type as a defector (regardless of beliefs) without restricting D . Suppose there exists some D , call it \bar{D}_u , such that the ugly firm's manager would not defect with any $D \geq \bar{D}_u$ even if his firm were (mistakenly) identified as a good firm. It is transparent that $\bar{D}_u > \hat{D}_b$. Moreover, since $\hat{D}_g > \hat{D}_b$ (this is easy to verify), it is also transparent that the bad firm cannot be ruled out as a potential defector with $D \in (0, \hat{D}_b)$ if it is possible that investors could view the defector as a good firm. Hence, no firm can be a priori eliminated as a potential defector when $D \in (0, \hat{D}_b)$ is observed, i.e., we cannot require that an arbitrarily high probability weight on the defector being of a particular type should be an a priori consideration of the equilibrium. This gives us just as much latitude in specifying out-of-equilibrium beliefs, in response to the observed defection $D \in (0, \hat{D}_b)$ as we have with sequential equilibrium. Thus, it is obvious that the sequential equilibrium survives Cho-Kreps.²⁰

Now consider a defection $D \in (\hat{D}_b, \hat{D}_g)$. To eliminate the ugly firm as a potential defector (or else we are back to the previous case), we need consider only defections $D \in (\bar{D}_u, \hat{D}_g)$. Investors' beliefs must now be concentrated on the bad and good firms. Is there any belief concentrated on these two types such that no type wishes to defect? Suppose the market believes that $\Pr(\text{defector is bad} \mid D \in (\bar{D}_u, \hat{D}_g)) = 1$. Then, it is clear that the ugly firm will not defect; it takes a debt level no more than \bar{D}_u and a belief that the defector is the good firm to induce it to defect. Moreover, the bad firm has no incentive to defect either, since it can be valued just as highly choosing $\hat{D}_b < \bar{D}_u$. Finally, the good firm strictly prefers not to be identified as a bad firm even at \hat{D}_b which is less than \bar{D}_u , so it too will not defect. Since no firm would ever

wish to defect with $D > \hat{D}_g$, we have shown that this separating equilibrium satisfies the Cho-Kreps intuitive criterion.

So we now have two equilibria that satisfy the intuitive criterion. Is there any way to choose? The answer is in the affirmative if we go to a stronger refinement, namely universal divinity. This refinement of sequential equilibrium, developed by Banks and Sobel [5] can be stated as follows:

Consider an out-of-equilibrium move m and let $\Omega_\tau(m)$ be the set of (mixed) best responses by the uninformed to this move m such that type τ strictly wishes to defect with m , given that set of best responses. Let T be the set of all types. Let $\Omega_\tau^*(m)$ be the set of (mixed) best responses such that type τ is indifferent between defecting and not defecting. Now suppose there are two types, τ and τ' , such that

$$\Omega_\tau \cup \Omega_\tau^* \subseteq \bigcup_{T \setminus \tau} \Omega_{\tau'} \quad (31)$$

then investors must believe that $\Pr(\text{defector is type } \tau \mid \text{defection } m \text{ observed}) = 0$, where by $T \setminus \tau$ we mean the set of all types with τ excluded. Apply this criterion pairwise to all types and then consider only the surviving types. Concentrate beliefs on the surviving types. If there exists at least one belief concentrated on the surviving types such that no type (survivors as well as non-survivors) wishes to defect given that belief, then the sequential equilibrium is universally divine.

This means that if it is more likely that the observed defection comes from type τ' than from type τ , then investors must believe that it is infinitely more likely that the defector is τ' . This is indeed a very strong refinement. A weaker version is divinity. It states that a sequential equilibrium is divine if there exists a belief of the uninformed in response to a defection m such that no type wishes to defect given the best response associated with that belief, and this belief belongs to the set of beliefs satisfying the rule that, given (31), the uninformed should not raise the probability (relative to their priors) that the defector is type τ rather than type τ' . While the set of universally divine sequential equilibria is nested within the set of sequential equilibria satisfying Cho-Kreps (i.e., every universally divine equilibrium passes the Cho-Kreps intuitive criterion, but the converse is not true), such a nesting relationship does not exist between divinity and Cho-

²⁰We can think of applying the Cho-Kreps intuitive criterion in a two-step process. First, we determine whether we can eliminate any type as a potential defector because that type would gain nothing from the defection regardless of the beliefs of the uninformed. Second, we see if there exists some belief concentrated on the remaining types such that no type would wish to defect, given the uninformed's best response associated with that belief. If we cannot eliminate any types in the first step, then the sequential equilibrium always survives Cho-Kreps.

Kreps. However, the set of universally divine equilibria is nested within the set of divine equilibria.

Let us now apply universal divinity to first eliminate the partially pooling equilibrium. Consider a defection $D \in (0, D_g^*)$. Given D , define V_u^* and V_b^* to satisfy

$$aV_u^* - \pi F_u(D) = a\bar{V}_p \quad (32)$$

$$aV_b^* - \pi F_b(D) = a\bar{V}_p \quad (33)$$

and define the sets of market value responses that induce the ugly and bad firms respectively to defect as $\Omega_u = (V_u^*, \mu_g]$ and $\Omega_b = (V_b^*, \mu_g]$. It is transparent from (32) and (33) that because of the FOSD relationship between the ugly and bad firms, $V_u^* > V_b^*$, so that $\Omega_u \cup \Omega_b$. Thus, investors must believe that $\Pr(\text{defector is ugly} \mid D \in (0, D_g^*)) = 0$. Thus, the lowest firm value that the defector can get is $V_b > \bar{V}_p$. Since $V_b > \bar{V}_p$, we know that for $D > 0$ small enough, $aV_b - \pi F_b(D) > a\bar{V}_p$, so that the bad firm defects. Hence, this sequential equilibrium is not universally divine.

We can verify that the separating equilibrium is universally divine. To do this we need to consider three possible types of defections: (i) $D < D_b$, (ii) $D \in (\hat{D}_b, \hat{D}_g)$, and (iii) $D > \hat{D}_g$. For any defection D , let $V_i^*(D)$ represent the valuation response of the market to the defection.

(i) $D < D_b$: Now, we know from the IC constraint that in equilibrium, the good firm strictly prefers its own allocation to that of the bad firm,

$$a\mu_u - \pi F_g(\hat{D}_g) > a\mu_b - \pi F_g(\hat{D}_b). \quad (34)$$

But, by the definition of $V_i^*(D)$, we have

$$aV_g^* - \pi F_g(D) = a\mu_g - \pi F_g(\hat{D}_g) \quad (35)$$

$$aV_b^* - \pi F_b(D) = a\mu_b - \pi F_b(\hat{D}_b) \quad (36)$$

where we have dropped the argument of $V_i^*(\cdot)$ for notational ease. Substituting (36) in (34) yields

$$a\mu_g - \pi F_g(\hat{D}_g) > aV_b^* - \pi F_b(D) + \pi F_b(\hat{D}_b) - \pi F_g(\hat{D}_b).$$

Now substituting (35) in the above inequality we get

$$a[V_g^* - V_b^*] > \pi\{F_b(\hat{D}_b) - F_g(\hat{D}_b)\} - \pi\{F_b(D) - F_g(D)\}. \quad (37)$$

Since F_g FOSD F_b and $D < \hat{D}_b$, we know that the RHS of (37) is strictly positive. Thus, $V_g^* > V_b^*$. This means that $\Omega_g \subset \Omega_b$, where $\Omega_g = (V_g^*, \mu_g]$ and $\Omega_b = (V_b^*, \mu_g]$. Thus, by universal divinity investors must believe that $\Pr(\text{defector is good} \mid D \in (0, D_b)) = 0$.

Similarly,

$$\begin{aligned} aV_u^* - \pi F_u(D) &= a\mu_u \\ &= a\mu_b - \pi F_u(\hat{D}_b) \\ &= aV_b^* - \pi F_b(D) + \pi F_b(\hat{D}_b) - \pi F_u(\hat{D}_b) \end{aligned}$$

This can be rearranged as

$$a[V_b^* - V_u^*] = \pi[F_u(\hat{D}_b) - F_b(\hat{D}_b)] - \pi[F_u(D) - F_b(D)] > 0.$$

Consequently, $V_b^* > V_u^*$, which implies $\Omega_b \subset \Omega_u = (V_u^*, \mu_g]$. By universal divinity then, investors must believe that $\Pr(\text{defector is bad} \mid D \in (0, \hat{D}_b)) = 0$. This means investors believe that the defector is ugly with probability one, in which case no firm wishes to defect with $D \in (0, D_b)$.

(ii) $D \in (\hat{D}_b, \hat{D}_g)$: From the definition of $V_g^*(D)$ and dropping the argument of the value function, we have

$$a[\mu_g - V_g^*] = \pi[F_g(\hat{D}_g) - F_g(D)]. \quad (38)$$

Similarly, using the definition of V_b^* , we have

$$a\mu_g = aV_b^* - \pi F_b(D) + \pi F_b(\hat{D}_g). \quad (39)$$

Substituting (39) in (38), since $D < \hat{D}_g$, we have

$$a[V_g^* - V_b^*] = \pi[F_b(\hat{D}_g) - F_g(\hat{D}_g)] - \pi[F_b(D) - F_g(D)] > 0.$$

Thus, $V_g^* > V_b^*$, implying that $\Omega_g \subset \Omega_b$. Investors must therefore believe that $\Pr(\text{defector is good} \mid D \in (D_b, D_g)) = 0$. Given these beliefs, however, no firm defects with $D \in (\hat{D}_b, \hat{D}_g)$.

(iii) $D > \hat{D}_g$: Even if investors believe with probability one that the defector is good, no firm will wish to defect with $D > D_g$, i.e., $\Omega_u = \Omega_b = \Omega_g = \emptyset$, the null set.

We have shown that no firm will wish to defect from the equilibrium. Hence, the equilibrium is universally divine. It is interesting that in this game, applying the strongest refinement we have used thus far leads to the perfectly separating equilibrium as the unique²¹ sequential equilibrium surviving that refinement. Recall that this is also the RRE if the uninformed agent moves first.

There are other refinements as well. To mention a few, one could use the Cho [14] forward induction equilibrium (FIE), the Grossman and Perry [27] perfect sequential equilibrium (PSE), and the strategic stability refinement of Kohlberg and Mertens [34]. I will not go into these refinements here for a number of reasons. First, although both the FIE and the PSE are stronger than the intuitive criterion in the sense that the set of sequential equilibria surviving either the FIE or the PSE are nested within the set of sequential equilibria satisfying the intuitive criterion, there is no nesting relationship between universal divinity and either the FIE or the PSE. This makes the choice of one refinement from among these three refinements a somewhat arbitrary decision. In Exhibit 2, I have drawn a picture to illustrate how the different equilibria are nested. Second, although the PSE places a more intuitive restriction on out-of-equilibrium beliefs than either the FIE or universal divinity,²² there is no general existence theorem which assures us that in every extensive form game with perfect recall²³ there will be at least one PSE.²⁴ Indeed, Grossman and Perry provide examples of games in which no PSE exists. Third, in many of the standard signalling games in finance, universal divinity is strong enough to prune the set of sequential equilibria down to a singleton. Fourth, even though strategic stability is a stronger refinement than universal divinity and there is an existence theorem which says that there is at least one strategically stable equilibrium in every extensive form game with perfect recall, it is not formally defined for games which are not finite. That is, strategic stability cannot be used in a game in which either the number of informed types or their strategies are infinite. In financial signalling models, the informed agent is usually allowed to choose from a nondenumerable set

of signals, such as the possible values of debt each firm could choose in the example we just discussed.

The purpose of refinements, of course, is to eventually be able to identify a unique equilibrium that is sequential and imposes sensible restrictions on out-of-equilibrium beliefs. This is a daunting task, however, and I am not aware of any refinement that delivers a unique equilibrium in general extensive form games (with perfect recall), along with a guarantee that this unique equilibrium will always exist.

This is a somewhat troublesome juncture at which we are left at present. When a researcher studies a game in which there are multiple sequential equilibria, he has a choice of which refinement to use. Quite often, this choice is one of convenience. The researcher wishes to focus on a particular set of predictions, and therefore chooses the refinement that suits his purpose. A sensible way to deal with this issue is to require that there be ancillary predictions that go beyond known stylized facts.²⁵ We can then conduct a joint test of the suitability of the chosen refinement and its appropriateness in explaining the chosen stylized facts by confronting these ancillary predictions with the data. Another reasonable criterion to use is to ask how robust the equilibrium in question is with respect to the specification of out-of-equilibrium beliefs. Equilibria that survive stronger refinements are more robust than those that don't. Thus, a simple way to be persuasive about one's choice of equilibrium may be to subject it to the strongest refinement applicable.

My earlier criticism of models in the SMUF class notwithstanding, those models have provided extremely valuable insights into the nature of asymmetric information equilibria. The fact that the RRE outcome is the unique sequential equilibrium that survives the strongest refinements in our debt signalling model is not a fluke. Cho and Sobel [15] have shown that in a special class of games — our debt signalling model belongs to this class — the unique universally divine sequential equilibrium is the RRE outcome.

II. Selected Applications of Noncooperative Game Theory in Finance

I'll review a selected number of applications of non-cooperative game theory in finance. This is by no means

²¹Although I have not proved uniqueness, it is true that there is only universally divine sequential equilibrium in this game.

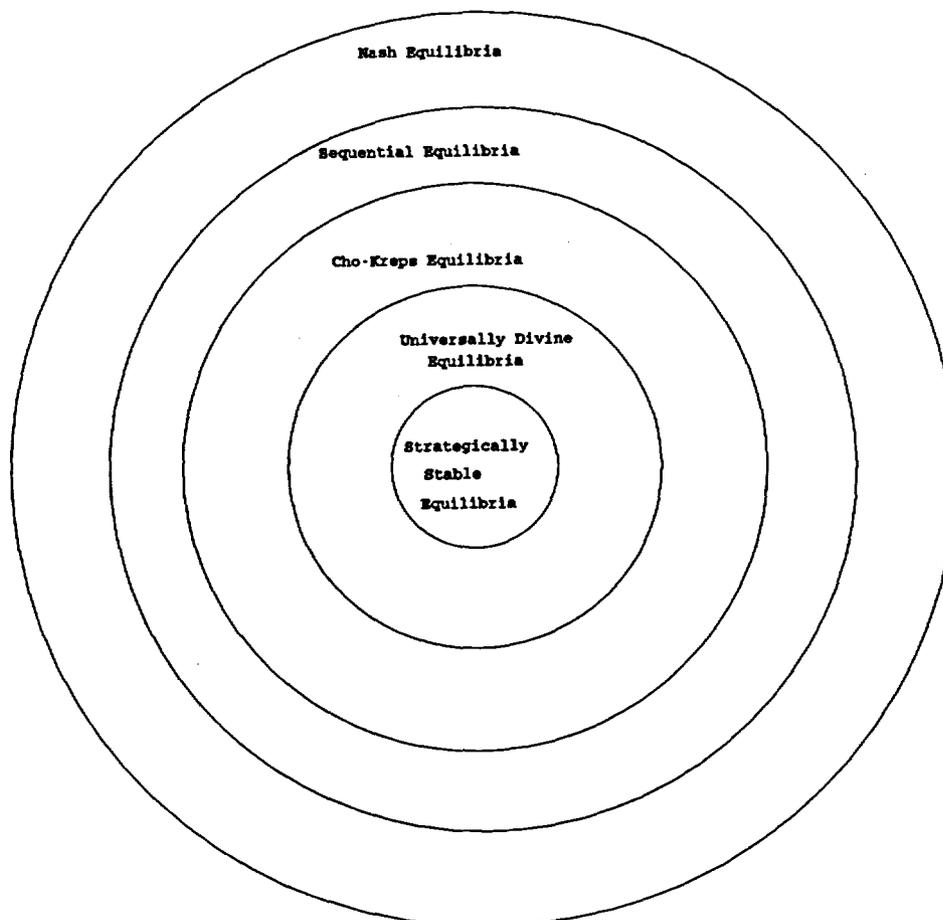
²²This is my personal opinion.

²³"Perfect recall" means that in a multistage game, players remember all of the previous moves.

²⁴See Grossman and Perry [27].

²⁵I owe this observation to Michael Brennan. See Thakor [62] for my discussion of Brennan's views on this.

Exhibit 2. Nesting Relationships of Refinements



intended to be an exhaustive review. It is merely intended to convey an idea of the wide range of issues in finance to which game theory can be applied.

A. Corporate Control and Takeovers

Issues in corporate control are particularly amenable to game-theoretic modeling because information is a crucial aspect of corporate control contests and because the sequence of moves is well defined. The traditional view was that takeovers improve efficiency by either displacing inefficient managers or by disciplining managers inclined to mismanage. Grossman

and Hart [26] argued that, if the target firm's shareholders are atomistic, then there is a free rider problem that interferes with the ability of takeovers to improve efficiency. In particular, the target firm's shareholders will decline any offer that is less than the post-takeover value of the firm. Thus, the bidder will not be able to profit from the takeover, and if it has to incur research and transactions costs prior to making its bid, it will lose money. This means there will be too few takeovers relative to the social optimum.

The Grossman and Hart solution to this free rider problem was to propose that firms write into their charters provisions that permit bidders to dilute the

value of the holdings of those shareholders who don't sell out to the bidder.²⁶ Shleifer and Vishny [59] showed, however, that dilution would be unnecessary as long as the bidder could — without tipping its hand— purchase a sufficiently large fraction of the target's shares prior to publicly announcing its offer.²⁷ Although the bidder does not necessarily profit on the shares it buys in the takeover, it does profit on the shares it acquired prior to the takeover. As long as this profit exceeds the bidder's research and transactions costs, the takeover is worthwhile. Apart from the other differences, a key assumption in Shleifer and Vishny that is different from Grossman and Hart is that not all shareholders are atomistic.

Shleifer and Vishny assume that the post-takeover value of the target is known privately to the bidder, but not to the target. They also restrict target shareholders to pure tendering strategies. With these specifications, they obtain a pooling sequential equilibrium in which the bidder offers a price that is independent of its private information and tender offers always succeed. Hirshleifer and Titman [30] relax these assumptions and permit target shareholders to adopt randomized tendering strategies. They show that the Shleifer-Vishny sequential equilibrium now does not survive a modified version of the Grossman and Perry [27] perfect sequential equilibrium,²⁸ even though there are beliefs that sustain the pooling equilibrium as a standard perfect sequential equilibrium. A sequential equilibrium that survives this criterion is separating. Shareholders who are indifferent between tendering and not tendering their shares do so randomly, making the success of the tender offer uncertain. There is a positive equilibrium relationship between the bid premium and the probability of acceptance of the bid, so that those bidders who have low valuations bid low and succeed less often on average, whereas those who have high valuations bid high and succeed more often on average.

Takeovers can be successful even without dilution or randomized tendering strategies. Bagnoli and Lip-

man [3] show that this can be achieved by assuming that the shareholders in the target firm are finite rather than atomistic. Intuitively, when shareholders are finite, each can be made pivotal to the success of the bid, and his incentive to free ride on the tendering of other shareholders can thus be eliminated. Bagnoli and Lipman characterize strong Nash equilibria — pure strategy Nash equilibria in which each player is employing his unique best response in equilibrium — in which the raider is tendered the number of shares needed for the takeover, and the bid succeeds with probability one.²⁹

As takeovers become more likely, we also observe increasing use of attempts by target management to resist takeovers. One popular method is greenmail which involves the target firm's management buying out selected shareholders at an inflated stock price to prevent the takeover. There are many who argue that this is simply a manifestation of management entrenchment and doesn't do the shareholders any good. However, Shleifer and Vishny [58] show that greenmail may benefit target shareholders. The basic idea in their model is that greenmail encourages potential bidders to investigate the firm, which possibly results in a takeover at a higher price than that offered in the initial takeover attempt that triggered the greenmail. That is, the manager of the target firm offers the initial bidder greenmail in the hope of attracting a white knight.³⁰

Another way for target management to resist a takeover is through a stock repurchase. A game-theoretic treatment of this appears in Bagnoli, Gordon and Lipman [4]. They assume that the manager cares both about retaining his job and about the value of the firm. Thus, the amount that the privately informed manager is willing to overpay in repurchasing stock depends upon what he knows about the value of the firm under his control. The higher this privately known value, the more he is willing to pay to repurchase stock. This way a repurchase offer acts as a positive signal of firm value (under current management) to the target firm's shareholders and convinces them not to sell out to the potential acquiror. Bagnoli, Gordon and Lipman show that this game has multiple sequential equilibria. They then use a variant of the Grossman and Perry [27]

²⁶One way to dilute the holdings of target shareholders who don't sell out is with a "freeze-out" merger in which the bidder buys 51% of the target shares and merges it with another firm he owns, at a price below the post-acquisition value of the target.

²⁷In the U.S., the Williams Act restricts this.

²⁸A perfect sequential equilibrium can be viewed as a refinement of sequential equilibria that survive the Cho-Kreps intuitive criterion, i.e., every perfect sequential equilibrium survives the intuitive criterion, but the converse is not true.

²⁹A property of a strong Nash equilibrium is that one can add a "small" amount of uncertainty to the original equilibrium and find a Nash equilibrium "close" to the original one.

³⁰The empirical evidence in Mikkelsen and Ruback [42] shows that greenmail has a negative effect on the target firm's stock price.

perfect sequential equilibrium to characterize the equilibrium they focus on as the unique one.

Although this paper was the first game-theoretic treatment of stock repurchase as a signal, Vermaelen [64] as well as Ofer and Thakor [49] had earlier modeled stock repurchases as signals. In fact, one could reverse the sequence of moves in Ofer and Thakor and show that the unique RRE in their model is a sequential equilibrium that survives the Banks and Sobel [5] universal divinity refinement.

An issue not yet well understood is why such huge premia are paid in takeovers. After all, we do not observe any profound restructuring of production processes subsequent to takeovers that would justify these premia. The most common practice is for acquirors to divest money-losing divisions of the target. But if this accounts for the premia, why don't the managers of target firms divest themselves?

Boot [8] provides an interesting answer. He shows that if managerial ability to identify positive NPV projects is privately known only to the manager, then managers will be concerned about their reputations. There is then a sequential equilibrium in which a manager may not divest a money-losing project that he initiated himself. Moreover, this sequential equilibrium survives many refinements, such as the Cho-Kreps intuitive criterion, universal divinity, and even strategic stability. The intuition is that divesting a project you picked yourself is an admission of a previous error. Therefore, you may hang on to such a project at the shareholders' expense, just to suppress the release of bad news about yourself. This makes it tempting for an acquiror to take over the firm and divest it of the money-losing segments of its business. The key is that the acquiror's reputation is unaffected by the divestiture because the divested projects were initiated by someone else. This explains how a takeover, when successful, can result in a substantial increase in shareholder wealth. Perhaps even more significantly, the threat of a takeover induces target firms' managers to be less reluctant to divest and hence improves shareholder wealth.

B. Capital Structure

There is an enormous literature on optimal capital structure. The contemporary literature, however, has used the tools of game theory to improve our understanding of capital structure.

Noe [48] models the financing decisions of a firm as a sequential signaling game. One of his objectives is to

explain why firms issue equity so infrequently. For example, Kalay and Shimrat [33] find that in a randomly selected sample of unregulated firms during the period 1964–1983, the average firm had issued equity only once in 50 years. Noe shows that, when the manager has perfect information regarding his own firm's cash flows that outsiders don't, there is a sequential equilibrium in which debt dominates equity. He also subjects this sequential equilibrium to a modified version of the Cho-Kreps intuitive criterion and proves that it survives this refinement. Thus, as in Myers and Majluf [45], a negative announcement effect is predicted for an equity issue.

A model that links capital structure to managerial reputational considerations and takeovers is Hirshleifer and Thakor [31]. As in Boot [8], it is assumed that managerial ability to identify good projects is unknown. They show that in a subgame perfect (Nash) reputational equilibrium, managers may opt for too much safety relative to the shareholders' optimum. When the firm issues debt, this incentive coincidentally aligns the manager's interests with those of the bondholders and reduces the agency costs of debt for shareholders. This leads to higher optimal leverage when the manager is concerned about his personal reputation than when he is not. This implies that, during periods of escalated takeover activity, managers — who will be more concerned about their reputations due to the increased likelihood of being displaced through a takeover — will optimally issue more debt than they would otherwise. Note that this prediction is different from that of models such as Harris and Raviv [28] and Stulz [61] which predict that managers may use debt as a takeover defense.

C. Dividends and Stock Repurchases

Another well-researched area, dividend policy, has also been reexamined with game-theoretic models. For example, Kumar [37] assumes that the manager has superior information about his firm's future cash flows and can choose investment policy as well as dividend policy for his firm. He shows that there is a sequential equilibrium in which dividends act as a coarse signal of future cash flows. That is, in equilibrium dividends show less variability than cash flows, as has been documented by Shiller [57] and others. Kumar also shows that in his equilibrium, dividends are poor predictors of future cash flows. This coarse sequential signalling equilibrium cannot be eliminated using the Cho-Kreps intuitive criterion.

As Ofer and Thakor [49] showed, since stock repurchases and dividends can both act as signals of future cash flows, it is important to treat them as simultaneous signals, if one wishes to understand relative price responses to these signals. Constantinides and Grundy [17] make a similar case for simultaneously studying stock purchases and investment choices as signals. They study a model in which the manager is privately informed about his firm's future prospects and moves first with a signal. They obtain a perfectly separating sequential equilibrium in which the investment level is the first best and management repurchases some of the outsiders' stock. Management is assumed to own some equity itself.

The intuition is as follows: As in Myers and Majluf [45], the manager wishes to sell overpriced securities. In the Constantinides-Grundy model, management has an incentive to overstate the value of the issued claim when raising funds to finance an investment. Thus, management should have an offsetting incentive to understate the value of the issued claim so that incentive compatibility is assured in a separating sequential equilibrium. This is achieved by management offering to repurchase some stock and excluding itself in the repurchase. The more management inflates the perceived value of the claim it issues, the more capital it raises and the larger is the number of shares that must be repurchased at the inflated price. This sequential equilibrium passes the Cho-Kreps intuitive criterion.

D. External Financing and Project Choice

In an influential paper, Myers and Majluf [45] showed that in a pooling equilibrium, equity issues are bad news. Giammarino and Lewis [25] show that one can obtain a separating sequential equilibrium by assuming that there is an underwriter intermediating between the issuing firm and the capital market. This sequential equilibrium survives the Cho-Kreps intuitive criterion and works as follows: The issuing firm (which is the informed agent) moves first by asking for a given price. The uninformed underwriter then responds by either accepting or rejecting the offer. If the offer is rejected, the firm foregoes its investment opportunity.

Giammarino and Lewis show that there is a sequential equilibrium in which the high-valued firm asks for a high price with probability one and the low-valued firm randomizes between asking for a high price and asking for a low price. The underwriter always accepts the low price offer, but randomizes its acceptance of

the high price offer, sometimes accepting it and sometimes rejecting it. This scheme is designed to ensure that the underwriter exactly breaks even and neither type of firm has an incentive to misrepresent.³¹ The virtue of this scheme is that one does not encounter as extreme a distortion due to asymmetric information as one does in Myers and Majluf [45] where positive NPV projects are sometimes eschewed because of the potential price decline accompanying equity issues.

In another paper that examines the implications of asymmetric information for the external financing decisions of firms, I (see Thakor [63]) consider a dynamic, game-theoretic model in which the informed manager moves first and the uninformed capital market responds. The manager has a project choice, and the market receives a noisy signal of this choice. I show that in a sequential equilibrium which survives the strategic stability refinement, the initially undervalued firm invests myopically even when the manager is maximizing the wealth of current shareholders. The reason is that an undervalued firm is interested in minimizing its future access to the equity market for investment funds. This is achieved by investing in a project that generates early cash flows, even though its intrinsic first best value is lower.

With this structure, I obtain numerous results about price reactions to equity issues, cash stockpiling attempts, etc. For example, the Myers and Majluf [45] analysis suggests that firms would benefit from having surplus cash to invest in projects. Does this mean that cash stockpiling is an optimal strategy? I show that the answer is no. Any attempt to stockpile cash leads to a negative stock price reaction. Throughout I assume that management is motivated to maximize current shareholder wealth. But if it has an entrenchment motive and a takeover is possible, preference for investment myopia is strengthened.

Related to this is recent work by Hirshleifer and Chordia [29] who examine a manager's preference for the timing of cash flow uncertainty resolution. Since a manager's ability may be better known to him than to others, perceptions of his ability are likely to be governed by observable proxies such as cash flows. Behaving strategically, the manager may influence these perceptions by manipulating cash flows through his investment choices. Within this framework, Hirshleifer and Chordia obtain numerous interesting results. When outsiders cannot observe whether the

³¹See Thakor [62] for a discussion of the intuition.

manager advanced or deferred resolution of cash flows, the unique Perfect Bayesian equilibrium involves the higher ability manager advancing resolution with a higher probability than one of lower ability. However, when outsiders can observe the manager's resolution decision, the unique universally divine sequential equilibrium involves both types of managers advancing resolution.

E. Financial Intermediation

Asymmetric information is at the heart of why financial intermediaries exist (see Allen [2], and the papers cited therein). So the potential for applying game theory to problems in intermediation is vast. I think a particularly fruitful area is multiperiod contracting involving intermediaries. Reputational issues become important in these contexts.

Reputational issues on the part of borrowers were first analyzed by John and Nachman [32]. In a two-period model, they depicted a sequential equilibrium in which the agency problem of debt is attenuated relative to that encountered in a single-period setting. Diamond [23] has a similar model, but one in which borrowers deal with banks over more than two time periods and have an incentive to develop reputations for repaying loans. As in John and Nachman, this provides a partial amelioration of the agency problem in single-shot games in which the borrower prefers riskier investments than the lender would like. This paper shows that in the early stages of a borrower's career, when reputation has yet to be developed, risk-taking incentives are the strongest. Once the borrower has developed a good reputation, it faces a future stream of lower interest rates than it did in the past. Thus, the present value of future loans is greater and the cost of bankruptcy (or credit exclusion) is higher. Risk-taking incentives are curtailed as a consequence.

Multiperiod contracting also creates a value for bank-customer relationships. In particular, information on a given borrower is likely to be reusable through time (see Chan, Greenbaum and Thakor [13]), and banks will have more information about their older borrowers than other banks do. This allows them to capture some rents in later stages of their relationships with customers, and has implications for initial contracting as well as capital allocation. Sharpe [56] has recently analyzed this problem.

Game-theoretic issues come alive in banking relationships where the supply, demand and pricing considerations related to financial contracts are in-

fluenced by potential strategic behavior. Boot, Thakor and Udell [9] have looked at bank loan commitments in which banks have incentives to renege on their promises to lend at below-market interest rates. They show that the bank's decision of whether or not to honor the commitment depends on its beliefs about some hidden action choice of the commitment customer, and this action choice depends on the customer's beliefs about whether or not the bank will honor the commitment. They find that it takes a sufficiently large bank to obtain a perfect sequential equilibrium in which the customer takes the appropriate action (which affects its financial condition at the time of drawing down on the loan commitment) and the bank honors the commitment.

III. Concluding Remarks

I have devoted my attention to noncooperative game theory thus far. This is for two reasons. One is my own research interest. The other is that this is the part of game theory that has been used the most in finance. However, there are other areas in game theory that may gain prominence in future financial applications. One is cooperative game theory and the other is evolutionary stability.

In noncooperative game theory, we do not permit players to collude and implement an outcome that could make them all better off. The commonly invoked impediments to collusion are communications difficulties, lack of trust, etc. In many settings, this is a reasonable assumption. However, there are also situations in which it is more useful to view the contracting outcome with the tools of cooperative game theory. One example is bankruptcy proceedings and reorganization. The eventual outcomes of such games involve considerable communication among the parties involved. Moreover, the relative bargaining strengths of the various parties significantly influence the final outcome.

An axiomatic approach to deal with such situations was proposed by Nash [47]. The solution concept he proposed is now known as the Nash bargaining solution. Its axiomatic requirements are that: (i) the solution should be independent of the units in which we measure utility, (ii) the solution should be Pareto optimal, so that no player can be made better off while leaving others no worse off, (iii) there should be independence of irrelevant alternatives (i.e., if we drop from the set of alternatives one that was not part of the optimal solution over that set of alternatives, then the

optimal solution does not change), and (iv) symmetry (i.e., switching labels on the players does not change the outcome). In these games, one also defines a threat point, which is the default outcome if the players do not reach an agreement. It can be shown that the optimal solution can be reached by maximizing the product of incremental utilities of the players involved, subject to the constraint that no player get less than his/her threat point utility, where incremental utility is defined as the difference between the utility in the optimal solution and the threat point utility.

What are some of the potential applications of cooperative game theory? One is the bankruptcy proceedings application I mentioned earlier. The important point in studying this application is not just to examine how the ex post outcome may be determined, but to ask how the ex ante use of debt and equity is affected by the anticipation of ex post bargaining. That is, there may be interesting capital structure ramifications of studying the game this way. Another fruitful application may be to loan restructuring in the sovereign debt market. We are all familiar with the LDC problems of banks. Renegotiations are common in these situations, and the relative bargaining positions of the coalition of banks on the one hand and the sovereign debtor on the other hand are crucial in how the renegotiation proceeds. Examples of papers on the LDC debt renegotiation problem are Bulow and Rogoff [10, 11] and Schwartz and Zurita [54].

Recently, there has also been some interest in yet another emerging branch of game theory called evolutionary stability. The game theory I have talked about thus far assumes that individuals are Bayesian rational, i.e., they revise their beliefs in accordance with Bayes' rule and maximize expected utility. Many believe, however, that human behavior is not characterized by ideal rationality. Rather, individuals behave in an adaptive manner that facilitates survival. This is the premise of evolutionary game theory, started by Maynard Smith and Price [41] and applied successfully in biology. It is curious that game theory, which has been created as a theory of rational behavior, is now applied to plants and animals. The idea is that a given species evolves to adapt or maximize fitness, where fitness is defined as reproductive success (roughly speaking, this is the expected number of offspring in the next generation). Applied to economics, this would mean that even preferences can now be explained as a result of evolution. Individuals like what is good for their fitness.

An interesting application ESS in finance appears in Cornell and Roll [18]. They show that an equilibrium in the stock market may involve some investors accepting the market price on some occasions and conducting security analysis on others. Their analysis is simple and appealing, although more generality is probably needed before their implications can be successfully tested. It is unclear at this stage, at least to me, how far evolutionary game theory³² will go in changing the way we think about financial economics. It is possible that we might explain some types of anomalies in prices or firm behavior by building models based on evolutionary game theory principles. The ultimate success of the theory will depend on how well its predictions stand up to empirical scrutiny.

While on the subject of rational behavior, it is worth noting that it is dangerous to view apparent violations of Bayesian rationality in everyday human behavior as evidence against Bayesian rationality. For example, there is a well-known problem in game theory — the finitely repeated prisoner's dilemma. This game has only one equilibrium outcome predicted by the theory: noncooperation in every period. Yet, in experiments one observes cooperation until shortly before the end. Does this refute the rationality assumptions of game theory? The answer is, not necessarily. In an interesting paper, Kreps, Milgrom, Roberts and Wilson [35] introduced a small amount of incomplete information about the payoffs of the other players and found that, in this slightly modified version of the usual game, the observed pattern of behavior is the equilibrium outcome.

In closing, the potential for applying game theory in finance seems substantial. I've probably omitted mentioning many things that people consider important on this subject. It is not necessarily intended to be a reflection of the importance of those topics relative to those I have discussed here. Rather, such possibly inadvertent omissions point to the ever-expanding scope of game theory and the consequent inability to do justice to the topic in a single paper.

References

1. G. Akerlof, "The Market for 'Lemons': Qualitative Uncertainty and the Market Mechanism," *Quarterly Journal of Economics* (February 1970), pp. 488-500.
2. F. Allen, "The Market for Information and the Origin of Financial Intermediation," *Journal of Financial Intermediation* (March 1990), pp. 3-30.

³²For a sparkling discussion of evolutionary game theory, see Selten [55]. Also see Becker [6].

3. M. Bagnoli and B.L. Lipman, "Successful Takeovers Without Exclusion," *Review of Financial Studies* 1-1, 1988, pp. 89-110.
4. M. Bagnoli, R. Gordon and B.L. Lipman, "Stock Repurchase as a Takeover Defense," *Review of Financial Studies* 2-3, 1989, pp. 423-443.
5. J. Banks and J. Sobel, "Equilibrium Selection in Signaling Games," *Econometrica* (May 1987), pp. 647-661.
6. G. Becker, "Altruism, Egoism and Genetic Fitness: Economics and Sociology," *Journal of Economic Literature* 14, 1976, pp. 817-825.
7. S. Bhattacharya, "Imperfect Information, Dividend Policy, and the 'Bird in the Hand' Fallacy," *Bell Journal of Economics* (Spring 1979), pp. 259-270.
8. A. Boot, "Why Hang on to Losers: Divestitures and Takeovers," Working Paper, Northwestern University, revised, October 1990.
9. A. Boot, A.V. Thakor and G.F. Udell, "Credible Commitments, Contract Enforcement Problems and Banks: Intermediation as Credibility Assurance," Working Paper, Indiana University, (forthcoming), *Journal of Banking and Finance*, 1991.
10. J. Bulow and K. Rogoff, "A Constant Recontracting Model of Sovereign Debt," *Journal of Political Economy* 97-1, 1989, pp. 155-177.
11. J. Bulow and K. Rogoff, "Sovereign Debt: Is to Forgive to Forget?" *American Economic Review* 79-1, 1989, pp. 43-50.
12. C.B. Cadsby, F. Murray and V. Maksimovic, "Pooling, Separating, and Semiseparating Equilibria in Financial Markets: Some Experimental Evidence," *The Review of Financial Studies* 3-3, 1990, pp. 315-342.
13. Y. Chan, S.I. Greenbaum and A.V. Thakor, "Information Reusability, Competition, and Bank Asset Quality," *Journal of Banking and Finance* 10, 1986, pp. 243-254.
14. I. Cho, "A Refinement of Sequential Equilibrium," *Econometrica* (November 1987), pp. 1367-1389.
15. I. Cho and J. Sobel, "Strategic Stability and Uniqueness in Signaling Games," *Journal of Economic Theory* (April 1990), pp. 381-413.
16. I. Cho and D. Kreps, "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics* (May 1987), pp. 179-221.
17. G. Constantinides and B. Grundy, "Optimal Investment with Stock Repurchase and Financing as Signals," *Review of Financial Studies* 2-4, 1989, pp. 445-465.
18. B. Cornell and R. Roll, "Strategies for Pairwise Competitions in Markets and Organizations," *Bell Journal of Economics* (Spring 1981), pp. 201-216.
19. M.M. Cornett and N.G. Travlos, "Information Effects Associated with Debt-for-Equity and Equity-for-Debt Exchange Offers," *Journal of Finance* (June 1989), pp. 451-468.
20. P. Dasgupta and E. Maskin, "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory," *Review of Economic Studies* (January 1986), pp. 1-26.
21. P. Dasgupta and E. Maskin, "The Existence of Equilibrium in Discontinuous Economic Games, II: Applications," *Review of Economic Studies* (January 1986), pp. 27-41.
22. D. Diamond and P. Dybvig, "Bank Runs, Deposit Insurance and Liquidity," *Journal of Political Economy* (June 1983), pp. 401-419.
23. D. Diamond, "Reputation Acquisition in Debt Markets," *Journal of Political Economy* (August 1989), pp. 828-862.
24. M. Engers and L. Fernandez, "Market Equilibrium with Hidden Knowledge and Self-Selection," *Econometrica* (March 1987), pp. 425-440.
25. R. Giammarino and T. Lewis, "A Theory of Negotiated Equity Financing," *Review of Financial Studies* (Fall 1988), pp. 265-288.
26. S. Grossman and O. Hart, "Takeover Bids, the Free-Rider Problem, and the Theory of the Corporation," *Bell Journal of Economics* (Spring 1980), pp. 42-64.
27. S. Grossman and M. Perry, "Perfect Sequential Equilibrium," *Journal of Economic Theory* (June 1986), pp. 97-119.
28. M. Harris and A. Raviv, "Corporate Control Contests and Capital Structure," *Journal of Financial Economics* (January/March 1988), pp. 55-86.
29. D. Hirshleifer and T. Chordia, "Resolution Preference and Project Choice," Working Paper, University of California at Los Angeles, November 1990.
30. D. Hirshleifer and S. Titman, "Share Tendering Strategies and the Success of Hostile Takeover Bids," *Journal of Political Economy* (April 1990), pp. 295-324.
31. D. Hirshleifer and A.V. Thakor, "Managerial Conservatism, Project Choice and Debt," Working Paper, revised, University of California at Los Angeles, October 1990.
32. K. John and D. Nachman, "Risky Debt, Investment Incentives, and Reputation in a Sequential Equilibrium," *Journal of Finance* (July 1985), pp. 863-878.
33. A. Kalay and A. Shimrat, "On the Payment of Equity Financed Dividends," Working Paper, New York University, 1986.
34. E. Kohlberg and J. Mertens, "On the Strategic Stability of Equilibria," *Econometrica* (September 1986), pp. 1003-1038.
35. D. Kreps, P. Milgrom, J. Roberts and R. Wilson, "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma," *Journal of Economic Theory* (August 1982), pp. 245-252.
36. D. Kreps and R. Wilson, "Sequential Equilibrium," *Econometrica* (July 1982), pp. 863-894.
37. P. Kumar, "Shareholder-Manager Conflict and the Information Content of Dividends," *The Review of Financial Studies* (Summer 1988), pp. 111-136.
38. H. Leland and D. Pyle, "Information Asymmetries, Financial Structure and Financial Intermediaries," *Journal of Finance* (May 1977), pp. 371-387.
39. R.W. Masulis, "The Effects of Capital Structure Changes on Security Prices," *Journal of Financial Economics* 8, 1980, pp. 139-178.
40. R.W. Masulis, "The Impact of Capital Structure Change on Firm Value: Some Estimates," *Journal of Finance* 38, 1983, pp. 107-126.
41. J. Maynard Smith and G.R. Price, "The Logic of Animal Conflict," *Nature* 246, 1973, pp. 15-18.
42. W. Mikkelsen and R. Ruback, "Targeted Repurchases and Common Stock Returns," Working Paper, Massachusetts Institute of Technology, 1986.
43. M. Miller and F. Modigliani, "Dividend Policy, Growth and the Valuation of Shares," *Journal of Business* (October 1961), pp. 411-433.
44. F. Modigliani and M. Miller, "The Cost of Capital, Corporation Finance, and the Theory of Investment," *American Economic Review* (June 1958), pp. 261-297.
45. S. Myers and N. Majluf, "Corporate Financing and Investment Decisions When Firms Have Information Investors Do Not Have," *Journal of Financial Economics* (June 1984), pp. 187-221.

46. R.B. Myerson, "Optimal Auction Design," *Mathematics of Operations Research* 6, 1981, pp. 58-73.
47. J. Nash, "The Bargaining Problem," *Econometrica* (January 1950), pp. 155-162.
48. T.H. Noe, "Capital Structure and Signaling Game Equilibria," *Review of Financial Studies* (Winter 1988), pp. 331-355.
49. A. Ofer and A.V. Thakor, "A Theory of Stock Price Responses to Alternative Corporate Cash Disbursement Mechanisms: Stock Repurchases and Dividends," *Journal of Finance* (June 1987), pp. 365-394.
50. E. Rasmusen, *Games and Information*, University Press, Cambridge, 1989.
51. J. Riley, "Informational Equilibrium," *Econometrica* (March 1979), pp. 331-359.
52. Ross, Stephen, "The Determination of Financial Structure: The Incentive Signalling Approach," *The Bell Journal of Economics* 40, 1977, pp. 23-40.
53. M. Rothschild and J.E. Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics* 90, 1976.
54. E.S. Schwartz and S. Zurita, "Sovereign Debt: Optimal Contract, Underinvestment and Forgiveness," UCLA Working Paper #15-90, University of California at Los Angeles, February 1991.
55. R. Selten, "Evolution, Learning and Economic Behavior," 1989 Nancy L. Schwartz Memorial Lecture, J.L. Kellogg Graduate School of Management, Northwestern University.
56. S.A. Sharpe, "Asymmetric Information, Bank Lending, and Implicit Contracts: A Stylized Model of Customer Relationships," *Journal of Finance* XLV-4, (September 1990), pp. 1069-1087.
57. R. Shiller, "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review* 71, 1981, pp. 421-436.
58. A. Shleifer and R. Vishny, "Greenmail, White Knights, and Shareholders' Interest," *Rand Journal of Economics* (Autumn 1985), pp. 293-309.
59. A. Shleifer and R. Vishny, "Large Shareholders and Corporate Control," *Journal of Political Economy* (June 1986), pp. 461-488.
60. M. Spence, "Competitive and Optimal Responses to Signals: Analysis of Efficiency and Distribution," *Journal of Economic Theory* (March 1974), pp. 296-332.
61. R. Stulz, "Managerial Control of Voting Rights: Financing Policies and the Market for Corporate Control," *Journal of Financial Economics* (January/March 1988), pp. 25-54.
62. A.V. Thakor, "Strategic Issues in Financial Contracting: An Overview," *Financial Management* (Summer 1989), pp. 39-58.
63. A.V. Thakor, "Information, Investment and Price Reactions," Working Paper, Indiana University, revised October 1990.
64. T. Vermaelen, "Repurchase Tender Offers, Signaling and Managerial Incentives," *Journal of Financial and Quantitative Analysis* 19, 1984, pp. 163-181.
65. J. von Neumann and O. Morgenstern, *The Theory of Games and Economic Behaviour*, New York, Wiley, 1944.
66. C. Wilson, "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory* 16, 1978.

1990 AMERICAN REAL ESTATE SOCIETY MANUSCRIPT PRIZE WINNERS

John Wiley & Sons

"Regional Economic Diversification and Residential Mortgage Default Risk," by Terrance M. Clauretie (University of Nevada at Las Vegas) in Volume 3, Number 1, pp. 87-97, *The Journal of Real Estate Research*, for the John Wiley & Sons manuscript prize (\$1,000). This prize is for the best paper to appear in *The Journal of Real Estate Research* in the year preceding the American Real Estate Society's (ARES) annual meeting. "Best" is determined by a mail ballot of the ARES membership.

Industrial Development Research Council

"Corporate Real Estate's Impact on the Takeover Market," by Brent W. Ambrose (University of Wisconsin at Milwaukee), for the Industrial Development Research Council manuscript prize (\$1,000). This prize is for the best paper presented on corporate real estate at the American Real Estate Society annual meeting.

Grubb & Ellis

"A Four-Square Design for Relating the Two Essential Dimensions of Real Estate Market Studies," by Dowell Myers (University of Southern California) and Kenneth Beck (University of Southern California), for the Grubb & Ellis manuscript prize (\$500). This prize is for the best paper presented on real estate market analysis at the American Real Estate Society annual meeting.

F M

Journal of The
Financial
Management
Association

Volume 20/Number 1
Spring 1991
\$10 per single copy

Financial Management

AGENCY AND CONTROL ISSUES

18 *INSIDER TRADING BEFORE NEW ISSUE ANNOUNCEMENTS*

Jonathan M. Karpoff and Daniel Lee

27 *AN AGENCY PERSPECTIVE ON FRANCHISING*

James A. Brickley, Frederick H. Dark and Michael S. Weisbach

**36 *THE ASSOCIATION BETWEEN EXECUTIVE STOCK OPTION PLAN CHANGES
AND MANAGERIAL DECISION MAKING***

Richard A. DeFusco, Thomas S. Zorn and Robert R. Johnson

.....

CORPORATE CASH DISTRIBUTIONS

44 *SHARE REPURCHASE MECHANISMS:*

***A COMPARATIVE ANALYSIS OF EFFICACY, SHAREHOLDER WEALTH
AND CORPORATE CONTROL EFFECTS***

Gerald D. Gay, Jayant R. Kale and Thomas H. Noe

60 *THE IMPACT OF THE 1986 TAX REFORM ACT ON EX-DIVIDEND DAY RETURNS*

Ashok J. Robin

.....

FM LECTURE SERIES

71 *GAME THEORY IN FINANCE*

Anjan V. Thakor

.....

SPECIAL TOPIC REVIEW

95 *THE JAPANESE COST OF FINANCE: A SURVEY*

Jeffrey A. Frankel